CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS OSPFv2 Authentication

Task

- Configure OSPF on R1, R2, and R3.
- The link between R1 and R3 should be in OSPF area 0.
- The link between R2 and R3, and R3's link to the Test PC, should be in OSPF area 1.
- Advertise the routers' Loopback networks into OSPF.
- Authenticate the OSPF exchange between R1 and R3 using the clear-text string **CLEARKEY**.
- Authenticate the OSPF exchange between R2 and R3 using an MD5 hash of the string **MD5HASH**.
- R1 and R2 should enable authentication at the interface level, and R3 should enable authentication globally under the OSPF process.
- Ensure that R3 does not have authentication enabled on its VLAN11 interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

R1:

```
interface Loopback0
 ip ospf 1 area 0
!
interface GigabitEthernet0/0
 ip ospf authentication
```

https://t.me/learningnets

```
 ip ospf authentication-key CLEARKEY
 ip ospf 1 area 0
R2:
interface Loopback0
 ip ospf 1 area 1
!
interface GigabitEthernet0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 MD5HASH
 ip ospf 1 area 1
R3:

interface Loopback0
 ip ospf 1 area 1
!
interface FastEthernet0/0.11
 ip ospf 1 area 1
 ip ospf authentication null
!
interface FastEthernet0/0.13
 ip ospf authentication-key CLEARKEY
 ip ospf 1 area 0
!
interface FastEthernet0/0.23
 ip ospf message-digest-key 1 md5 MD5HASH
 ip ospf 1 area 1
!
router ospf 1
 area 0 authentication
 area 1 authentication message-digest
```

Verification

OSPF authentication is configured in two distinct steps:

1. Enable authentication (clear-text or MD5) at the interface level or global level.
2. Configure the string used for authentication, as well as the key number only for MD5 authentication.

The common confusion with OSPF authentication comes from the fact that the first of these steps, enabling the authentication, can be done in two ways: at the interface level and under the global OSPF process. When OSPF authentication is enabled under the global process, it applies to all interfaces in that area. In this

https://t.me/learningnets

particular example, R3 has the commands `area 0 authentication` and `area 1 authentication message-digest` configured under the OSPF process. This means that all interfaces in area 0 run clear-text authentication, and all interfaces in area 1 run MD5 authentication.

When OSPF authentication is enabled at the interface level, it overrides the configuration of the global process, but only for that particular interface. For example, while R3 has the command `area 0 authentication` configured under the global process, the command `ip ospf authentication message-digest` or `ip ospf authentication null` configured on an interface in area 0 would override the global clear-text authentication configuration.

Regardless how OSPF authentication is enabled, either under the global process or at the interface level, the packet format for the authentication remains the same. This means that the difference in configuration is only cosmetic, because both configurations have the same result.

The authentication key, however, is always configured at the interface level. For clear text authentication, this is the `ip ospf authentication-key` command, whereas for MD5 authentication, this is the `ip ospf message-digest-key` command. Note that just like RIP and EIGRP, the MD5 key number must match between the neighbors, because it is a salt for the MD5 hash.

OSPF authentication, like EIGRP authentication, can be verified simply based on the fact that an adjacency has formed, as seen below:

```
R3#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.1.1         1   FULL/DR         00:00:35    136.1.13.1      FastEthernet0/0.13
150.1.2.2         1   FULL/BDR        00:00:36    136.1.23.2      FastEthernet0/0.23
```

The specific authentication configuration can be verified as seen below.

```
R3#show ip ospf interface FastEthernet0/0.13
FastEthernet0/0.13 is up, line protocol is up
  Internet Address 136.1.13.3/24, Area 0
  Process ID 1, Router ID 150.1.3.3, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.1.1, Interface address 136.1.13.1
  Backup Designated router (ID) 150.1.3.3, Interface address 136.1.13.3
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

https://t.me/learningnets

```
    oob-resync timeout 40
    Hello due in 00:00:03
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.1.1  (Designated Router)
  Suppress hello for 0 neighbor(s) Simple password authentication enabled
R3#show ip ospf interface FastEthernet0/0.23
FastEthernet0/0.23 is up, line protocol is up
  Internet Address 136.1.23.3/24, Area 1
  Process ID 1, Router ID 150.1.3.3, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1         no          no            Base
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 150.1.3.3, Interface address 136.1.23.3
  Backup Designated router (ID) 150.1.2.2, Interface address 136.1.23.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:00
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 3/4, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 2, maximum is 2
  Last flood scan time is 0 msec, maximum is 4 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.2.2  (Backup Designated Router)
  Suppress hello for 0 neighbor(s) Message digest authentication enabled
 Youngest key id is 1
```

If authentication has been enabled at the global level, it can be determined by using the command `show ip ospf`：

```
R3#  show ip ospf | section Area
Area BACKBONE(0)
    Number of interfaces in this area is 1 Area has simple password authentication
```

https://t.me/learningnets

```
    SPF algorithm last executed 00:09:29.632 ago
    SPF algorithm executed 3 times
    Area ranges are
Area 1
    Number of interfaces in this area is 3 (1 loopback) Area has message digest authentication

    SPF algorithm last executed 00:04:16.836 ago
    SPF algorithm executed 6 times
    Area ranges are
```

Failures in OSPF authentication can be verified as a failure to establish adjacency. The below process-level debug shows a mismatch in authentication type, where R3 has MD5 authentication configured to R1, and R1 has clear-text authentication configured to R3. From an OSPF packet level, type 0 authentication means NULL - or no authentication, type 1 is clear-text, and type 2 is MD5.

```
R3#debug ip ospf adj
OSPF adjacency events debugging is on
R3#config t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#interface FastEthernet0/0.13R3(config-subif)#ip ospf authentication message-digest
R3(config-subif)#
OSPF: Send with youngest Key 0 OSPF: Rcv pkt from 136.1.13.1, FastEthernet0/0.13 :
Mismatch Authentication type. Input packet specified type 1, we use type 2
```

A mismatch in the password between two neighbors (or hash in the case of MD5) will result in a different output in this debug, as seen below on R2:

```
R2#debug ip ospf adj
OSPF adjacency debugging is on
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#interface GigabitEthernet0/0
R2(config-if)#no ip ospf message-digest-key 1
R2(config-if)#ip ospf message-digest-key 1 md5 WRONG
OSPF-1 ADJ   Gi0/0: Rcv pkt from 136.1.23.3 :Mismatched Authentication Key - Message Digest Key 1
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS OSPFv3 Authentication and Encryption

Task

- Configure OSPFv3 on R1, R2, and R3.
- The link between R1 and R3 should be in OSPF area 0.
- The link between R2 and R3, and R3's link to the Test PC should be in OSPF area 1.
- Advertise the routers' Loopback networks into OSPF.
- Authenticate the OSPF exchange between R1 and R3 as follows:
  - Use IPsec AH with an SPI of 1000.
  - Use SHA1 for authentication, and a key of **0123456789012345678901234567890123456789**.
- Authenticate and encrypt the OSPF exchange between R2 and R3 as follows:
  - Use IPsec ESP with an SPI of 2000.
  - Use MD5 for authentication, and a key of **0123456789ABCDEF0123456789ABCDEF**.
  - Use AES 256 for encryption, and a key of **0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF**.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

R1:

https://t.me/learningnets

```
interface Loopback0
 ipv6 ospf 1 area 0
!
interface GigabitEthernet0/0
 ipv6 ospf 1 area 0
 ipv6 ospf authentication ipsec spi 1000 sha1 01234567890123456789012345678901234567890123456789
R2:
interface Loopback0
 ipv6 ospf 1 area 1
!
interface GigabitEthernet0/0
 ipv6 ospf 1 area 1
 ipv6 ospf encryption ipsec spi 2000 esp aes-cbc 256 0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDE
R3:

interface Loopback0
 ipv6 ospf 1 area 1
!
interface FastEthernet0/0.11
 ipv6 ospf 1 area 1
!
interface FastEthernet0/0.13
 ipv6 ospf 1 area 0
 ipv6 ospf authentication ipsec spi 1000 sha1 01234567890123456789012345678901234567890123456789
!
interface FastEthernet0/0.23
 ipv6 ospf 1 area 1
 ipv6 ospf encryption ipsec spi 2000 esp aes-cbc 256 0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDE
```

Verification

OSPFv3 uses IPsec for both authentication and encryption, as IPsec is already embedded into the IPv6 specification. Like regular LAN-to-LAN IPsec tunnels for IPv4, IPv6 IPsec supports both Authentication Header (AH) for authentication, and Encapsulating Security Payload (ESP) for authentication and encryption. Because OSPFv3 doesn't use ISAKMP for dynamic key negotiation in Phase 1, the Phase 2 authentication and encryption keys must be manually configured. This is why the strings in the above configuration are so long; these values are normally automatically negotiated by the routers through a process such as the Diffie Hellman (DH) exchange.

Similar to OSPFv2 authentication, a failure in OSPFv3 authentication or encryption will result in a failed adjacency. This means that a simple final verification of the

https://t.me/learningnets

above configuration is to check the adjacency status, as seen below:

```
R3#show ipv6 ospf neighbor

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
150.1.1.1         1   FULL/DR         00:00:35    4               FastEthernet0/0.13
150.1.2.2         1   FULL/DR         00:00:35    3               FastEthernet0/0.23
```

Verification of the configuration can be seen at the interface level output below:

```
R3#show ipv6 ospf interface FastEthernet0/0.13
FastEthernet0/0.13 is up, line protocol is up
  Link Local Address FE80::215:C6FF:FE78:8A40, Interface ID 7
  Area 0, Process ID 1, Instance ID 0, Router ID 150.1.3.3
  Network Type BROADCAST, Cost: 1 SHA-1 authentication SPI 1000, secure socket UP (errors: 0)
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.1.1, local address FE80::207:7DFF:FE3D:94E8
  Backup Designated router (ID) 150.1.3.3, local address FE80::215:C6FF:FE78:8A40
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Graceful restart helper support enabled
  Index 1/1/3, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 4, maximum is 6
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.1.1  (Designated Router)
  Suppress hello for 0 neighbor(s)
R3#show ipv6 ospf interface FastEthernet0/0.23
FastEthernet0/0.23 is up, line protocol is up
  Link Local Address FE80::215:C6FF:FE78:8A40, Interface ID 8
  Area 1, Process ID 1, Instance ID 0, Router ID 150.1.3.3
  Network Type BROADCAST, Cost: 1 AES-CBC-256 encryption MD5 auth SPI 2000, secure socket UP (errors: 0)

  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 150.1.2.2, local address FE80::207:7DFF:FECE:F00
  Backup Designated router (ID) 150.1.3.3, local address FE80::215:C6FF:FE78:8A40
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Graceful restart helper support enabled
  Index 1/3/4, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 2, maximum is 6
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
```

https://t.me/learningnets

```
    Adjacent with neighbor 150.1.2.2  (Designated Router)
  Suppress hello for 0 neighbor(s)
```

Because IPsec is the protocol doing the authentication and encryption, a normal Phase 2 Security Association (IPsec SA) is formed between the adjacent neighbors. This can be verified similar to any other LAN-to-LAN or Remote Access IPsec VPN configuration on the router, as seen below:

```
R3#show crypto ipsec sa
 interface: FastEthernet0/0.13
     Crypto map tag: (none), local addr FE80::215:C6FF:FE78:8A40

   IPsecv6 policy name: OSPFv3-1-1000
   IPsecv6-created ACL name: FastEthernet0/0.13-ipsecv6-ACL

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (FE80::/10/89/0)
   remote ident (addr/mask/prot/port): (::/0/89/0)
   current_peer :: port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 147, #pkts encrypt: 147, #pkts digest: 147
 #pkts decaps: 145, #pkts decrypt: 145, #pkts verify: 145
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: FE80::215:C6FF:FE78:8A40,
     remote crypto endpt.: ::
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.13
     current outbound spi: 0x3E8(1000)
     PFS (Y/N): N, DH group: none

     inbound esp sas:
 inbound ah sas:
 spi: 0x3E8(1000)
 transform: ah-sha-hmac
 ,
        in use settings ={Transport, }
        conn id: 2001, flow_id: NETGX:1, sibling_flags 80000001, crypto map: (none)
        no sa timing
        replay detection support: N
        Status: ACTIVE

     inbound pcp sas:
```

https://t.me/learningnets

```
     outbound esp sas:
 outbound ah sas:
spi: 0x3E8(1000)
transform: ah-sha-hmac
,
        in use settings ={Transport, }
        conn id: 2002, flow_id: NETGX:2, sibling_flags 80000001, crypto map: (none)
        no sa timing
        replay detection support: N
        Status: ACTIVE

     outbound pcp sas:
interface: FastEthernet0/0.23
    Crypto map tag: (none), local addr FE80::215:C6FF:FE78:8A40

   IPsecv6 policy name: OSPFv3-1-2000
   IPsecv6-created ACL name: FastEthernet0/0.23-ipsecv6-ACL

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (FE80::/10/89/0)
   remote ident (addr/mask/prot/port): (::/0/89/0)
   current_peer :: port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 134, #pkts encrypt: 134, #pkts digest: 134
#pkts decaps: 135, #pkts decrypt: 135, #pkts verify: 135
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: FE80::215:C6FF:FE78:8A40,
     remote crypto endpt.: ::
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.23
     current outbound spi: 0x7D0(2000)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
spi: 0x7D0(2000)
transform: esp-256-aes esp-md5-hmac
,
        in use settings ={Transport, }
        conn id: 2003, flow_id: NETGX:3, sibling_flags 80000001, crypto map: (none)
        no sa timing
        IV size: 0 bytes
        replay detection support: N
        Status: ACTIVE

     inbound ah sas:
```

https://t.me/learningnets

```
     inbound pcp sas:
 outbound esp sas:
spi: 0x7D0(2000)
transform: esp-256-aes esp-md5-hmac
,
        in use settings ={Transport, }
        conn id: 2004, flow_id: NETGX:4, sibling_flags 80000001, crypto map: (none)
        no sa timing
        IV size: 0 bytes
        replay detection support: N
        Status: ACTIVE

     outbound ah sas:

     outbound pcp sas:


R3#show crypto sockets


Number of Crypto Socket connections 2

Fa0/0. Peers: (local): ::
               (remote): ::
       Local Ident (addr/plen/port/prot): (FE80::/10/0/89)
       Remote Ident (addr/plen/port/prot): (::/0/0/89)
       IPSec Profile: "OSPFv3-1-1000"
       Socket State: OpenClient: "OSPFv3 1" (Client State: Active)
Fa0/0. Peers: (local): ::
               (remote): ::
       Local Ident (addr/plen/port/prot): (FE80::/10/0/89)
       Remote Ident (addr/plen/port/prot): (::/0/0/89)
       IPSec Profile: "OSPFv3-1-2000"
       Socket State: OpenClient: "OSPFv3 1" (Client State: Active)


Crypto Sockets in Listen state:
```

A failure in authentication or encryption will result in a log message in the crypto process, as seen below:

```
R2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#interface GigabitEthernet0/0
R2(config-if)#no ipv6 ospf encryption ipsec spi 2000
R2(config-if)#ipv6 ospf encryption ipsec spi 2000 esp aes-cbc 256 1111111111111111111111111111111111111111111111111111111111111111
```

https://t.me/learningnets

```
(ipsec_license_release) IPSEC license not requested - nothing to release
R2(config-if)#
%CRYPTO-6-ISAKMP_ON_OFF: ISAKMP is ON
%OSPFv3-3-UNEXPECTED_MESSAGE: OSPFv3 has received an unexpected message: 9 / 0
%CRYPTO-4-RECVD_PKT_MAC_ERR: decrypt: mac verify failed for connection id=2003 spi=000007D0 seqno=000000A0


R3#
%CRYPTO-4-RECVD_PKT_MAC_ERR: decrypt: mac verify failed for connection id=2003 spi=000007D0 seqno=00000001
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS EIGRP Authentication

Task

- Configure EIGRP AS 1 on R1, R2, and R3.
- Advertise the routers' Ethernet links and Loopbacks into EIGRP.
- Authenticate the EIGRP exchange between R1 and R3 using an MD5 hash of the string **MD5HASH**.
- Authenticate the EIGRP exchange between R2 and R3 as follows:
  - Use a key chain with three keys.
  - Use the key-strings **KEY1**, **KEY2**, and **KEY3**, respectively.
  - KEY1 should be sent between 00:00:00 Jan 1 2013 and 00:00:00 Jan 1 2014.
  - KEY1 should be accepted between 00:00:00 Jan 1 2013 and 01:00:00 Jan 1 2014.
  - KEY2 should be sent between 00:00:00 Jan 1 2014 and 00:00:00 Jan 1 2015.
  - KEY2 should be accepted between 23:00:00 Dec 31 2013 and 01:00:00 Jan 1 2015.
  - KEY3 should be sent between 00:00:00 Jan 1 2015 and 00:00:00 Jan 1 2016.
  - KEY3 should be accepted between 23:00:00 Dec 31 2014 and 01:00:00 Jan 1 2016.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of

https://t.me/learningnets

reloading the devices.

Configuration

R1:

```
key chain MD5CHAIN
 key 1
  key-string MD5HASH
!
interface GigabitEthernet0/0
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 MD5CHAIN
!
router eigrp 1
 network 136.1.0.0
 network 150.1.0.0
```

R2:

```
key chain ROLLOVER
 key 1
  key-string KEY1
  accept-lifetime 00:00:00 Jan 1 2013 01:00:00 Jan 1 2014
  send-lifetime 00:00:00 Jan 1 2013 00:00:00 Jan 1 2014
 key 2
  key-string KEY2
  accept-lifetime 23:00:00 Dec 31 2013 01:00:00 Jan 1 2015
  send-lifetime 00:00:00 Jan 1 2014 00:00:00 Jan 1 2015
 key 3
  key-string KEY3
  accept-lifetime 23:00:00 Dec 31 2014 01:00:00 Jan 1 2016
  send-lifetime 00:00:00 Jan 1 2015 00:00:00 Jan 1 2016
!
interface GigabitEthernet0/0
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 ROLLOVER
!
router eigrp 1
 network 136.1.0.0
 network 150.1.0.0
```

R3:

```
key chain MD5CHAIN
 key 1
```

https://t.me/learningnets

```
  key-string MD5HASH
!
key chain ROLLOVER
 key 1
  key-string KEY1
  accept-lifetime 00:00:00 Jan 1 2013 01:00:00 Jan 1 2014
  send-lifetime 00:00:00 Jan 1 2013 01:00:00 Jan 1 2014
 key 2
  key-string KEY2
  accept-lifetime 23:00:00 Dec 31 2013 01:00:00 Jan 1 2015
  send-lifetime 23:00:00 Dec 31 2013 01:00:00 Jan 1 2015
 key 3
  key-string KEY3
  accept-lifetime 23:00:00 Dec 31 2014 01:00:00 Jan 1 2016
  send-lifetime 23:00:00 Dec 31 2014 01:00:00 Jan 1 2016
!
interface FastEthernet0/0.13
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 MD5CHAIN
!
interface FastEthernet0/0.23
 ip authentication mode eigrp 1 md5
 ip authentication key-chain eigrp 1 ROLLOVER
!
router eigrp 1
 network 136.1.0.0
 network 150.1.0.0
```

Verification

EIGRP authentication works similar to RIP and IS-IS authentication in that it is key-chain based. Unlike RIP and IS-IS authentication, though, EIGRP only supports MD5, not clear-text. This example shows EIGRP being used with a single authentication key between R1 and R3, but with automatic key rotation between R2 and R3.

The accept and send lifetimes configured in the key chains allow you to pre-stage the routers to periodically change their passwords used for authentication. The potential problem with this configuration is that if the routers' time is not NTP synchronized, key rollover could happen at different time intervals, which could potentially cause a loss of the EIGRP adjacencies. To prevent this, there is typically a "fudge factor" used with the key rotation, where the multiple authentication keys accept lifetimes overlap. In this example, keys 2 and 3 are valid if received for a

https://t.me/learningnets

year, plus or minus one hour.

To verify whether EIGRP-enabled interfaces have authentication configured and which key chain is being used, use below command:

```
R3#show ip eigrp interfaces detail fastEthernet0/0.13
EIGRP-IPv4 Interfaces for AS(1)
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0.13           1        0/0         4       0/1           50           0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Un/reliable mcasts: 0/2  Un/reliable ucasts: 2/1
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 1
  Topology-ids on interface - 0 Authentication mode is md5,  key-chain is "MD5CHAIN"

R3#show ip eigrp interfaces detail fastEthernet0/0.23
EIGRP-IPv4 Interfaces for AS(1)
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0.23           1        0/0         4       0/1           50           0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Un/reliable mcasts: 0/2  Un/reliable ucasts: 2/1
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 1
  Topology-ids on interface - 0 Authentication mode is md5,  key-chain is "ROLLOVER"
```

The key that is currently being used can be verified as shown here:

```
R3#show clock
18:08:46.763 UTC WedAug 7 2013
R3#show key chain
Key-chain MD5CHAIN:
    key 1 -- text "MD5HASH"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
Key-chain ROLLOVER:
    key 1 -- text "KEY1"        accept lifetime (00:00:00 UTC Jan 1 2013) - (01:00:00 UTC Jan 1 2014)
[valid now]
```

https://t.me/learningnets

```
        send lifetime (00:00:00 UTC Jan 1 2013) - (01:00:00 UTC Jan 1 2014) [valid now]

    key 2 -- text "KEY2"
        accept lifetime (23:00:00 UTC Dec 31 2013) - (01:00:00 UTC Jan 1 2015)
        send lifetime (23:00:00 UTC Dec 31 2013) - (01:00:00 UTC Jan 1 2015)
    key 3 -- text "KEY3"
        accept lifetime (23:00:00 UTC Dec 31 2014) - (01:00:00 UTC Jan 1 2016)
        send lifetime (23:00:00 UTC Dec 31 2014) - (01:00:00 UTC Jan 1 2016)
```

The automatic rollover can be seen by changing the clock on the router, and then making sure that the EIGRP adjacency doesn't go down. As an example of this, R2 and R3's clocks are updated to be near Jan 1st 2014, but R3's clock is about 1 minute faster than R2's. This means that R3 will perform rollover first, and will drop the adjacency to R2 if the accept and send lifetimes aren't properly configured:

```
R2#clock set 23:58:00 dec 31 2013
%SYS-6-CLOCKUPDATE: System clock has been updated from 18:20:21 UTC Wed Aug 7 2013 to 23:58:00 UTC Tue Dec 31 2013,
R3#clock set 23:59:00 dec 31 2013

%SYS-6-CLOCKUPDATE: System clock has been updated from 18:20:18 UTC Wed Aug 7 2013 to 23:59:00 UTC Tue Dec 31 2013,
```

After the clock changes, both R2 and R3 initially agree that key 1 is being sent, but both keys 1 and 2 can be received:

```
R2#show key chain
Key-chain ROLLOVER:
    key 1 -- text "KEY1"        accept lifetime (00:00:00 UTC Jan 1 2013) - (01:00:00 UTC Jan 1 2014)
[valid now]
        send lifetime (00:00:00 UTC Jan 1 2013) - (00:00:00 UTC Jan 1 2014) [valid now]
    key 2 -- text "KEY2"        accept lifetime (23:00:00 UTC Dec 31 2013) - (01:00:00 UTC Jan 1 2015)
[valid now]
        send lifetime (00:00:00 UTC Jan 1 2014) - (00:00:00 UTC Jan 1 2015)
    key 3 -- text "KEY3"
        accept lifetime (23:00:00 UTC Dec 31 2014) - (01:00:00 UTC Jan 1 2016)
        send lifetime (00:00:00 UTC Jan 1 2015) - (00:00:00 UTC Jan 1 2016)
R3#show key chain
Key-chain MD5CHAIN:
    key 1 -- text "MD5HASH"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
Key-chain ROLLOVER:
    key 1 -- text "KEY1"        accept lifetime (00:00:00 UTC Jan 1 2013) - (01:00:00 UTC Jan 1 2014)
[valid now]
```

https://t.me/learningnets

```
        send lifetime (00:00:00 UTC Jan 1 2013) - (00:00:00 UTC Jan 1 2014) [valid now]

    key 2 -- text "KEY2"        accept lifetime (23:00:00 UTC Dec 31 2013) - (01:00:00 UTC Jan 1 2015)
 [valid now]

        send lifetime (00:00:00 UTC Jan 1 2014) - (00:00:00 UTC Jan 1 2015)
    key 3 -- text "KEY3"
        accept lifetime (23:00:00 UTC Dec 31 2014) - (01:00:00 UTC Jan 1 2016)
        send lifetime (00:00:00 UTC Jan 1 2015) - (00:00:00 UTC Jan 1 2016)
```

Because R3 reaches 2014 first, key 1 is no longer valid to be sent, and key 2 is sent instead:

```
R3#show key chain
 Key-chain MD5CHAIN:
    key 1 -- text "MD5HASH"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
 Key-chain ROLLOVER:
    key 1 -- text "KEY1"        accept lifetime (00:00:00 UTC Jan 1 2013) - (01:00:00 UTC Jan 1 2014)
 [valid now]
        send lifetime (00:00:00 UTC Jan 1 2013) - (00:00:00 UTC Jan 1 2014)
    key 2 -- text "KEY2"        accept lifetime (23:00:00 UTC Dec 31 2013) - (01:00:00 UTC Jan 1 2015)
 [valid now]
        send lifetime (00:00:00 UTC Jan 1 2014) - (00:00:00 UTC Jan 1 2015) [valid now]

    key 3 -- text "KEY3"
        accept lifetime (23:00:00 UTC Dec 31 2014) - (01:00:00 UTC Jan 1 2016)
        send lifetime (00:00:00 UTC Jan 1 2015) - (00:00:00 UTC Jan 1 2016)
```

Although only one key (the lowest-numbered key) can be sent at a time, multiple keys can be received. Therefore, even though the routers performed key rollover about a minute apart, the adjacency was not lost. Eventually R2 catches up, and they then agree that only key 2 is being sent.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS BGP Authentication

Task

- Configure BGP on R1, R2, and R3.
- Use BGP ASN's 1, 2, and 3, respectively.
- Advertise the Loopback networks of the routers into BGP.
- R3 should peer with both R1 and R2.
- Authenticate these peerings with an MD5 hash of the password **CISCO**.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

```
R1:
router bgp 1
 network 150.1.1.1 mask 255.255.255.255
 neighbor 136.1.13.3 remote-as 3
 neighbor 136.1.13.3 password CISCO
R2:
router bgp 2
 network 150.1.2.2 mask 255.255.255.255
 neighbor 136.1.23.3 remote-as 3
 neighbor 136.1.23.3 password CISCO
R3:

router bgp 3
```

https://t.me/learningnets

```
network 150.1.3.3 mask 255.255.255.255
neighbor 136.1.13.1 remote-as 1
neighbor 136.1.13.1 password CISCO
neighbor 136.1.23.2 remote-as 2
neighbor 136.1.23.2 password CISCO
```

Verification

BGP uses TCP for transport, and because TCP already has a specification for authentication, BGP uses TCP authentication instead of a separate internal mechanism. Specifically, it uses TCP Option 19, which is the MD5 Signature Option. Like OSPF or EIGRP, this means that if BGP authentication fails, a peering relationship will not be established. Therefore, the most basic way to verify that BGP authentication is working is to look at the **show ip bgp summary** output and ensure that the peers are up.

```
R3#show ip bgp summary
BGP router identifier 150.1.3.3, local AS number 3
BGP table version is 4, main routing table version 4
3 network entries using 408 bytes of memory
3 path entries using 156 bytes of memory
3/3 BGP path/bestpath attribute entries using 372 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 984 total bytes of memory
BGP activity 3/0 prefixes, 3/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
136.1.13.1      4            1       6       7        4    0    0 00:02:001
136.1.23.2      4            2       6       7        4    0    0 00:01:591
```

The fact that prefixes are being received implies that the TCP session is up, but this could be further verified by viewing the BGP table, as seen below:

```
R3#show ip bgp

BGP table version is 4, local router ID is 150.1.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete
```

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 150.1.1.1/32     136.1.13.1               0             0 1 i
*> 150.1.2.2/32     136.1.23.2               0             0 2 i
*> 150.1.3.3/32     0.0.0.0                  0         32768 i
```

Outside of simply looking at the BGP configuration, authentication can be verified on a per-neighbor basis, as seen below:

```
R3#show ip bgp neighbors 136.1.13.1
BGP neighbor is 136.1.13.1,  remote AS 1, external link
  BGP version 4, remote router ID 150.1.1.1
  BGP state = Established, up for 00:02:09
  Last read 00:00:15, last write 00:00:43, hold time is 180, keepalive interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Multisession Capability: advertised
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                4          2
    Keepalives:             2          3
    Route Refresh:          0          0
    Total:                  7          6
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  Session: 136.1.13.1
  BGP table version 4, neighbor version 4/0
  Output queue size : 0
  Index 1
  1 update-group member
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                                Sent       Rcvd
  Prefix activity:              ----       ----
    Prefixes Current:              3          1 (Consumes 52 bytes)
    Prefixes Total:                3          1
```

https://t.me/learningnets

```
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          1
    Used as multipath:            n/a          0

                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    Total:                               0          0
  Number of NLRIs in the update sent: max 1, min 0
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never

  Address tracking is enabled, the RIB does have a route to 136.1.13.1
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 1
Local host: 136.1.13.3, Local port: 179
Foreign host: 136.1.13.1, Foreign port: 59545
Connection tableid (VRF): 0
Maximum output segment queue size: 50

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0xB4F2D0):
Timer          Starts    Wakeups            Next
Retrans             5          0             0x0
TimeWait            0          0             0x0
AckHold             5          4             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
ProcessQ            0          0             0x0

iss: 3705844654  snduna: 3705844933  sndnxt: 3705844933
irs: 3937982015  rcvnxt: 3937982208

sndwnd:  16106  scale:      0  maxrcvwnd:  16384
rcvwnd:  16192  scale:      0  delrcvwnd:    192

SRTT: 487 ms, RTTO: 3168 ms, RTV: 2681 ms, KRTT: 0 ms
```

https://t.me/learningnets

```
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
Status Flags: passive open, gen tcbs Option Flags: nagle, path mtu capable, md5
, 0x1000000
IP Precedence value : 6

Datagrams (max data segment is 1460 bytes):
Rcvd: 10 (out of order: 0), with data: 5, total data bytes: 192
Sent: 11 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 0), with data: 5, total data bytes: 27
 Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
```

A failure in BGP authentication will trigger a syslog message:

```
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#router bgp 1R1(config-router)#neighbor 136.1.13.3 password WRONG
%TCP-6-BADAUTH: Invalid MD5 digest from 136.1.13.3(179) to 136.1.13.1(59545)

R3#%TCP-6-BADAUTH: Invalid MD5 digest from 136.1.13.1(59545) to 136.1.13.3(179)
```

Note that this is a different error message than if authentication is simply not enabled on the peering, such as seen below:

```
R1(config-router)#no neighbor 136.1.13.3 password


R3#%TCP-6-BADAUTH: No MD5 digest from 136.1.13.1(179) to 136.1.13.3(37874)
```

This is an important difference that we will see later when BGP authenticated sessions pass through either the ASA firewall or the IPS sensor, because the modifications they do to the IP header and payload will affect the TCP MD5 signature.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Routing Protocol Filtering

Task

- Configure RIPv2 on R1, R2, and R3.
- Advertise their Ethernet links and Loopback networks into RIP.
- Configure control plane filtering for RIPv2 as follows:
  - R1 and R3 should only accept RIPv2 multicast packets from each other, and discard RIP packets from any other devices on the VLAN 13 segment.
  - R2 and R3 should only exchange unicast RIPv2 updates with each other, and discard RIP packets from any other devices on the VLAN 23 segment.
- Configure route filtering for RIPv2 as follows:
  - Configure R3 so that it does not advertise the VLAN 11 segment to R1.
  - Configure R2 so that it does not accept any routes with a subnet mask of /28 or greater.
- R3 should not send RIPv2 updates out the VLAN 11 segment.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

R1:

```
interface GigabitEthernet0/0
 ip access-group RIPV2_MULTICAST in
!
```

https://t.me/learningnets

```
router rip
 version 2
 network 136.1.0.0
 network 150.1.0.0
 no auto-summary
!
ip access-list extended RIPV2_MULTICAST
 permit udp host 136.1.13.3 eq rip host 224.0.0.9 eq rip
 deny udp any any eq rip
 deny udp any eq rip any
 permit ip any any
```

R2:

```
interface GigabitEthernet0/0
 ip access-group RIPV2_UNICAST in
!
router rip
 version 2
 passive-interface GigabitEthernet0/0
 network 136.1.0.0
 network 150.1.0.0
 neighbor 136.1.23.3
 distribute-list prefix PERMIT_27_OR_LOWER in GigabitEthernet0/0
 no auto-summary
!
ip access-list extended RIPV2_UNICAST
 permit udp host 136.1.23.3 eq rip host 136.1.23.2 eq rip
 deny udp any any eq rip
 deny udp any eq rip any
 permit ip any any
!
ip prefix-list PERMIT_27_OR_LOWER permit 0.0.0.0/0 le 27
```

R3:

```
interface FastEthernet0/0.13
 ip access-group RIPV2_MULTICAST in
!
interface FastEthernet0/0.23
 ip access-group RIPV2_UNICAST in
!
router rip
 version 2
 passive-interface FastEthernet0/0.11
 passive-interface FastEthernet0/0.23
 network 136.1.0.0
 network 150.1.0.0
 network 10.1.0.0
```

https://t.me/learningnets

```
 neighbor 136.1.23.2
 distribute-list prefix DENY_VLAN11 out FastEthernet0/0.13
 no auto-summary
!
ip access-list extended RIPV2_MULTICAST
 permit udp host 136.1.13.1 eq rip host 224.0.0.9 eq rip
 deny udp any any eq rip
 deny udp any eq rip any
 permit ip any any
!
ip access-list extended RIPV2_UNICAST
 permit udp host 136.1.23.2 eq rip host 136.1.23.3 eq rip
 deny udp any any eq rip
 deny udp any eq rip any
 permit ip any any
!
ip prefix-list DENY_VLAN11 deny 10.1.0.0/24
ip prefix-list DENY_VLAN11 permit 0.0.0.0/0 le 32
```

Verification

This task illustrates two types of filtering, control plane filtering and route filtering. In the control plane, RIPv2 uses UDP source and destination port 520 multicasts to the address 224.0.0.9 to exchange routing information by default, as follows:

- RIP Request Messages are sent with a random source port or a port of 520, but the destination port is always 520.
- RIP Response Messages are sent with a source port of 520 and a destination port equal to the source port of the receiving RIP Request Message.
- RIP Unsolicited Messages are sent with both source and destination of 520.

https://t.me/learningnets

The destination address can be changed to either broadcast or unicast if desired. In the case of unicast updates, the `neighbor` command specifies the unicast destination, whereas the `passive-interface` command is used to suppress the default multicast updates sent outbound. Because RIP does not have the concept of neighbors formed through hello messages, and because when RIP is enabled on an interface it starts sending and listenting for routing updates, a passive-interfaces still allows the receiving of routing updates. Therefore, beyond authentication, a simple way to secure the RIPv2 control plane is to configure an access-list that permits or denies UDP port 520 traffic to and from the specific addresses that are relevant to the network design. This access-list is then applied inbound at the interface level as a data plane filter with the `ip access-group` command.

> For other routing protocols such as EIGRP, OSPF, or BGP, the logic of these filters is exactly the same, but the layer 3 and layer 4 header information is different. EIGRP and OSPF use their own transports, as IP protocol numbers 88 and 89, respectively. BGP uses TCP for transport and port 179. EIGRP and OSPF can use both multicast or unicast depending on configuration, whereas BGP uses only unicast. Knowledge of the layer 3 and layer 4 header information of the different control plane protocols is critical within the scope of CCIE Security.

For route filtering, the `distribute-list` command under the routing process is used to filter the actual payload of the RIP updates. The distribute-list can then call a standard access-list, an extended access-list, or a prefix-list, to specify which particular networks (prefixes) are permitted to be advertised to or received from a neighbor. In the case of inbound filters, you can also match on the source of the routing update, for example, to say permit network X but only from source Y. Not all variations of route filtering are covered in this example, simply because within the scope of CCIE Security there is not a very large emphasis on routing; however, you can use the following as a guideline for RIP:

- Using standard access-lists allows filtering based on the prefix, e.g. 10.0.0.0.
- Using extended access-lists allows filtering based on the prefix and the update source/neighbor.
- Using prefix-lists allows filtering based on the prefix and prefix-length, e.g. 10.0.0.0/24.

In terms of verification, control plane filter can be checked simply by ensuring that updates are being installed from the correct sources, and that hits are occuring on the correct access-list entry, as seen below:

```
R3#show ip route rip
```

https://t.me/learningnets

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 3 subnets
R        150.1.1.1 [120/1] via 136.1.13.1
, 00:00:27, FastEthernet0/0.13
R        150.1.2.2 [120/1] via 136.1.23.2
, 00:00:00, FastEthernet0/0.23
R3#show ip access-list
Extended IP access list RIPV2_MULTICAST
    10 permit udp host 136.1.13.1 eq rip host 224.0.0.9 eq rip (2 matches)
    20 deny udp any any eq rip (2 matches)
    30 deny udp any eq rip any
    40 permit ip any any
Extended IP access list RIPV2_UNICAST
    10 permit udp host 136.1.23.2 eq rip host 136.1.23.3 eq rip (2 matches)

    20 deny udp any any eq rip
    30 deny udp any eq rip any
    40 permit ip any any
```

The route filters can be verified by either checking the resulting routes installed in the routing table (or the RIP database), or through the process-level debug shown below. Note that the 10.1.0.0/24 network is advertised to R2, but not R1, and that the destination address of RIP updates to R1 is multicast, whereas to R2 it is unicast.

```
R3#debug ip rip
RIP protocol debugging is on
RIP: sending v2 update to 224.0.0.9
 via FastEthernet0/0.13 (136.1.13.3)
RIP: build update entries
        136.1.23.0/24 via 0.0.0.0, metric 1, tag 0
        150.1.2.2/32 via 0.0.0.0, metric 2, tag 0
        150.1.3.3/32 via 0.0.0.0, metric 1, tag 0
RIP: sending v2 update to 136.1.23.2
 via FastEthernet0/0.23 (136.1.23.3)
RIP: build update entries
10.1.0.0/24 via 0.0.0.0, metric 1, tag 0

        136.1.13.0/24 via 0.0.0.0, metric 1, tag 0
        150.1.1.1/32 via 0.0.0.0, metric 2, tag 0
```

https://t.me/learningnets

```
        150.1.3.3/32 via 0.0.0.0, metric 1, tag 0
```

Passive interfaces as well as applied route filters can be determined using command `show ip protocols` , as follows:

```
R3#show ip protocols

*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
FastEthernet0/0.13 filtered by (prefix-list) DENY_VLAN11 (per-user), default is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 27 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Neighbor(s):
    136.1.23.2
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0.13    2     2
    Loopback0             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
    136.1.0.0
    150.1.0.0 Passive Interface(s):
FastEthernet0/0.11
FastEthernet0/0.23

  Routing Information Sources:
    Gateway         Distance      Last Update
    136.1.13.1           120      00:00:13
    136.1.23.2           120      00:00:24
  Distance: (default is 120)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Control Plane Policing

Task

- Configure Control Plane Policing on R3 to prevent it from being flooded with UDP DoS attacks.
- UDP packets should be limited to 16 Kbps.
- Other traffic to the router should not be affected by this policy.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

https://t.me/learningnets

```
R3:

ip access-list extended UDP_ACL
 permit udp any any
!
class-map match-all UDP_CLASS
 match access-group name UDP_ACL
!
policy-map POLICE_UDP
 class UDP_CLASS
  police 16000
!
control-plane
 service-policy input POLICE_UDP
```

Verification

Packets destined to an IP address of the router itself, such as IGP routing protocols, BGP, Telnet, SSH, SNMP, etc., must be processed by the general-purpose CPU of the control plane. This is different from transit traffic in the data plane, such as the router moving a web browsing flow between interfaces, because data plane traffic can be hardware-accelerated with dedicated ASICs, depending on the platform.

When the router receives a packet destined to itself, it is punted from the hardware switching process to the software process, which is referred to as "process switching." Packets that are process switched can result in high CPU utilization on the router, which can ultimately result in packet loss in the data plane or potentially even dropped control plane sessions, such as OSPF adjacencies.

This also means that the router itself is susceptible to a Denial of Service attack when packets that result in process switching are flooded at or near line rate. To prevent this type of attack, the Control Plane Policing (CoPP) feature can be used to deny traffic or limit the rate of traffic that is punted from the hardware switching to the software switching process.

To verify this feature in action, first view the router's current and historic CPU utilization:

```
R3#show processes cpu sorted
CPU utilization for five seconds: 1%/0%; one minute: 13%; five minutes: 21%
 PID Runtime(ms)     Invoked      uSecs   5Sec   1Min   5Min TTY Process
 114         432      233337          1  0.63%  0.83%  0.98%   0 Ethernet Msec Ti
   2         212         369        574  0.07%  0.03%  0.02%   0 Load Meter
```

https://t.me/learningnets

```
 359           32         7137           4  0.07%  0.03%  0.02%   0 MLD
  66           20         1848          10  0.07%  0.02%  0.00%   0 Per-Second Jobs
 161           12         7209           1  0.07%  0.01%  0.00%   0 SSS Feature Time
 <snip>
R3#show processes cpu history


     111111111222221111111111111112222211111111111111111111222221
  100
   90
   80
   70
   60
   50
   40
   30
   20
   10
     0....5....1....1....2....2....3....3....4....4....5....5....6
               0    5    0    5    0    5    0    5    0    5    0
               CPU% per second (last 60 seconds)
```

The output above shows that the CPU is currently under very little load. Currently there is no CoPP policy applied. To verify the result of a DoS attack without CoPP protection enabled, we launch a UDP flooding attack from the Test PC against the router's IP address.

https://t.me/learningnets

Image not readable or empty
/uploads/workbooks/images/originals/cdb41e123fd8b967524e21e73f039f419810b57f.png



Now that the flood attack is in progress, check the router's CPU utilization again. The `sorted` argument is similar to the linux `top` command, where it shows the processes with highest CPU utilization first.

```
R3#show processes cpu sorted
CPU utilization for five seconds: 78%/31%; one minute: 39%; five minutes: 21%
```

https://t.me/learningnets

```
 PID Runtime(ms)     Invoked      uSecs   5Sec   1Min   5Min TTY Process
 142      123836       79625       1555 43.59% 21.50%  5.93%   0 IP Input

 114         508      258919          1  0.95%  0.80%  0.84%   0 Ethernet Msec Ti
  12        4552        2471       1842  0.95%  0.36%  0.17%   0 Exec
   6        3412         333      10246  0.87%  0.19%  0.13%   0 Check heaps
 138         180       63681          2  0.15%  0.16%  0.16%   0 IPAM Manager
<snip>
```

The process **IP Input** means process switched packets. For more information on this see, the white paper [Troubleshooting High CPU Utilization on Cisco Routers](). Next, apply the CoPP policy to R3, and check for matches against the policy:

```
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#control-plane
R3(config-cp)#service-policy input POLICE_UDP
R3(config-cp)#end
R3#R3#show policy-map control-plane
 Control Plane

  Service-policy input: POLICE_UDP

    Class-map: UDP_CLASS (match-all) 18300 packets, 23578306 bytes
      5 minute offered rate 94000 bps, drop rate 94000 bps
      Match: access-group name UDP_ACL
      police:
          cir 16000 bps, bc 1500 bytes
        conformed 20 packets, 21842 bytes; actions:
          transmit  exceeded 18286 packets, 23564132 bytes; actions:
drop

        conformed 2000 bps, exceed 784000 bps

    Class-map: class-default (match-any)
      3 packets, 206 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

https://t.me/learningnets

Note in the above output that most of the UDP traffic is exceeding the policy, and is therefore getting dropped. If we view the CPU utilization on the router, we can see that it has dropped dramatically. From the `show processes cpu history` output, we can see exactly where the CoPP policy was applied, as the CPU dropped from 80% utilization down to 30%:

```
R3#show processes cpu sorted
CPU utilization for five seconds: 25%/23%
; one minute: 65%; five minutes: 56%
 PID Runtime(ms)     Invoked      uSecs   5Sec   1Min   5Min TTY Process
 114          996      301049          3  1.03%  1.01%  0.96%   0 Ethernet Msec Ti
 138          312       74188          4  0.23%  0.18%  0.16%   0 IPAM Manager
  12         4708        2609       1804  0.23%  0.08%  0.07%   0 Exec
 142       279556      172029       1625  0.15% 34.40% 29.58%   0 IP Input
<snip>
R3#show processes cpu history


R3   09:33:27 AM Saturday Aug 10 2013 UTC




     2222222222222222222222227777777777777777777777777777777777
     6666666665555555555888886666666666666667777666666666688888
 100
  90
  80                          ***********************************
  70                          ***********************************
  60                          ***********************************
  50                          ***********************************
  40                          ***********************************
  30 ************************************************************
  20 ************************************************************
  10 ************************************************************
    0....5....1....1....2....2....3....3....4....4....5....5....6
              0    5    0    5    0    5    0    5    0    5    0
              CPU% per second (last 60 seconds)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Control Plane Protection

Task

- Configure RIPv2 routing between R1, R2, and R3.
- Configure Control Plane Protection on R3 as follows:
  - Drop packets going to all closed ports.
  - Ensure that this does not affect connections made on TCP ports 3020 and 3040 or UDP 520.
  - To reduce the effect of broadcast storms, limit the rate of non-IP input packets to 100 per second.
  - Set the queue-limit for input HTTP packets to 100 packets and limit the packet rate to 10 per second.
  - Ensure that all fragmented packets transiting the router are limited to 1 Mpps.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

```
R1:
router rip
 version 2
 network 0.0.0.0
 no auto-summary
R2:
```

https://t.me/learningnets

```
router rip
 version 2
 network 0.0.0.0
 no auto-summary
```

R3:

```
ip access-list extended FRAGMENTS
 permit ip any any fragment
!
class-map FRAGMENTS
 match access-group name FRAGMENTS
!
ip access-list extended HTTP
 permit tcp any any eq www
!
class-map HTTP
 match access-group name HTTP
!
class-map type port-filter match-all CLOSED_PORTS
 match closed-ports
 match not port TCP 3020
 match not port TCP 3040
 match not port UDP 520
!
class-map type queue-threshold QUEUE_HTTP
 match protocol http
!
policy-map type port-filter HOST_PORT_FILTER
 class CLOSED_PORTS
  drop
!
policy-map type queue-threshold HOST_QUEUE_THRESHOLD
  class QUEUE_HTTP
    queue-limit 100
!
policy-map HOST_RATE_LIMIT
 class HTTP
  police rate 10 pps
!
policy-map TRANSIT_RATE_LIMIT
 class FRAGMENTS
  police rate 1000000 pps
!
policy-map CEF_EXCEPTION_RATE_LIMIT
 class class-default
  police rate 100 pps
```

https://t.me/learningnets

```
!
control-plane host
  service-policy input HOST_RATE_LIMIT
  service-policy type queue-threshold input HOST_QUEUE_THRESHOLD
  service-policy type port-filter input HOST_PORT_FILTER
!
control-plane transit
  service-policy input TRANSIT_RATE_LIMIT
!
control-plane cef-exception
 service-policy input CEF_EXCEPTION_RATE_LIMIT
!
router rip
 version 2
 network 0.0.0.0
 no auto-summary
```

Verification

Control Plane Protection (CPPr) is an extension of the Control Plane Policing (CoPP/CPP) feature. Both the CoPP and CPPr features are intended to prevent infrastructure attacks, such as attacks that target the router itself. CoPP implements this through a single aggregate policer that is applied to the CPU—for example, allowing you to rate-limit OSPF, BGP, SNMP, etc. that is targeted at the router itself. CPPr extends the CoPP feature to make it more granular in its policers.

Specifically, CPPr does this by treating the Route Processor (RP) as a virtual interface attached to the router. All packets redirected to the RP are classified into one of three possible categories, which in turn correspond to one of three subinterfaces of the virtual interface. These categories are the **control-plane host** subinterface, the **control-plane transit** subinterface, and the **control-plane CEF-exception** subinterface.

The **control-plane host** subinterface receives all TCP and UDP traffic that is directly destined for one of the router's interfaces. Examples of control-plane host traffic include management traffic and routing protocols. Most control plane protection features operate on this subinterface, and thus this subinterface provides the most features, such as policing, port filtering, and per-protocol queue thresholds. Note, however, that non-TCP/UDP control traffic does not terminate on this subinterface, but instead terminates on the CEF-exception subinterface.

The **control-plane CEF-exception** subinterface is where packets that cause an exception in CEF switching path land. This includes non-IP router-destined traffic,

https://t.me/learningnets

such as CDP, L2 keepalive messages (such as Frame-Relay LMI keepalives), or ARP packets and IP packets with options or TTL <=1. Additionally, non-UDP local multicast traffic destined to the router (such as OSPF updates) falls into this category as well.

The **control-plane transit subinterface** is intended to handle transit IP packets that for some reason can't be switched via CEF. This might happen, for instance, when a packet must be routed out an Ethernet interface and no ARP lookup has been performed yet for the next-hop, which makes the CEF adjacency incomplete.

In this task, we are first asked to drop traffic for all closed ports on the router, with the exception of TCP 3020 and 3040, as well as UDP 520. This is accomplished by configuring a policy of type port-filter on the control-plane host interface, because the router can automatically detect open ports, and then drop traffic destined for the closed ports. The current open ports can be verified as follows:

```
R3#show control-plane host open-ports

Active internet connections (servers and established)
Prot                Local Address            Foreign Address                  Service    State
 tcp                        *:23                       *:0                      Telnet   LISTEN
 udp                        *:67                       *:0           DHCPD Receive   LISTEN
 udp                      *:1975                       *:0                     IPC   LISTEN
```

> In the output above, UDP 520 is not listed as an open port, even though the RIP routing process is enabled. This means that the router will drop RIP packets due to the `match closed-ports` in the `class-map type port-filter`. This is why the extra `match not port UDP 520` is required so that RIP packets aren't discarded. For more information on this problem, see the white paper Understanding Control Plane Protection.

The rest of this configuration can be verified as follows:

```
R3#show policy-map type port-filter control-plane all

 Control Plane Host

  Service-policy port-filter input: HOST_PORT_FILTER

    Class-map: CLOSED_PORTS (match-all)
      305 packets, 29980 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match:  closed-ports
```

https://t.me/learningnets

```
      Match: not  port tcp 3020
      Match: not  port tcp 3040
      Match: not  port udp 520 drop

    Class-map: class-default (match-any)
      37 packets, 4850 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
R3#show policy-map type queue-threshold control-plane all
 queue-limit 100
      queue-count 0     packets allowed/dropped 0/0
 Control Plane Host

  Service-policy queue-threshold input: HOST_QUEUE_THRESHOLD

    Class-map: QUEUE_HTTP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match:  protocol http

    Class-map: class-default (match-any)
      43 packets, 5296 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
R3#show policy-map control-plane host

 Control Plane Host

  Service-policy input: HOST_RATE_LIMIT

    Class-map: HTTP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name HTTP
      police: rate 10 pps, burst 2 packets
        conformed 0 packets; actions:
          transmit
        exceeded 0 packets; actions:
          drop
        conformed 0 pps, exceed 0 pps

    Class-map: class-default (match-any)
      318 packets, 32054 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
R3#show policy-map control-plane transit
```

https://t.me/learningnets

```
 Control Plane Transit

  Service-policy input: TRANSIT_RATE_LIMIT

    Class-map: FRAGMENTS (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name FRAGMENTS
      police: rate 1000000 pps, burst 244140 packets
        conformed 0 packets; actions:
          transmit
        exceeded 0 packets; actions:
          drop
        conformed 0 pps, exceed 0 pps

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
R3#show policy-map control-plane cef-exception
 Control Plane Cef-exception

  Service-policy input: CEF_EXCEPTION_RATE_LIMIT

    Class-map: class-default (match-any)
      3647 packets, 311514 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      police: rate 100 pps, burst 24 packets

        conformed 3647 packets; actions:
          transmit
        exceeded 0 packets; actions:
          drop
        conformed 0 pps, exceed 0 pps
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Management Plane Protection

Task

- Configure R3 with a local user **user1** with the password **cisco123** and a privilege level of 15.
- R3 should only listen for SSH and HTTPS management sessions, and authenticate these using the local user database.
- Configure Management Plane Protection so that SSH and HTTPS management requests can only come in on R3's link to the Test PC.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

```
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#ip domain-name ine.com
R3(config)#crypto key generate rsa modulus 1024
The name for the keys will be: R3.ine.com

% The key modulus size is 1024 bits
% Generating 1024 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 5 seconds)%SSH-5-ENABLED:SSH 1.99 has been enabled
R3:

username user1 privilege 15 password cisco123
!
```

https://t.me/learningnets

```
no ip http server
ip http authentication local
ip http secure-server
!
control-plane host
 management-interface FastEthernet0/0.11 allow https ssh
!
line vty 0 4
 login local
```

The Management Plane Protection (MPP) feature allows you to limit the interfaces on which the router listens for management sessions, such as Telnet, SSH, or HTTP. MPP works as part of Control Plane Protection (CPPr) under the **control-plane host** mode, as management sessions to the router are process switched.

Without MPP enabled, R3 accepts Telnet, SSH, or HTTPS sessions from any link:

```
R1#telnet 136.1.13.3
Trying 136.1.13.3 ... Open

User Access Verification

Username: user1
Password: R3#

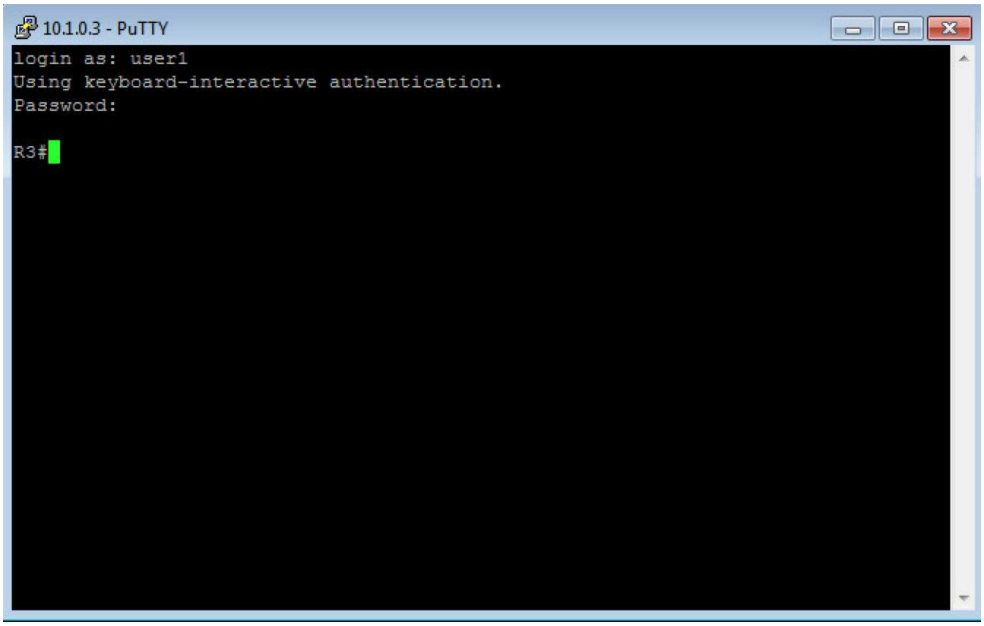
R2#telnet 136.1.23.3 443
Trying 136.1.23.3, 443 ... Open
```

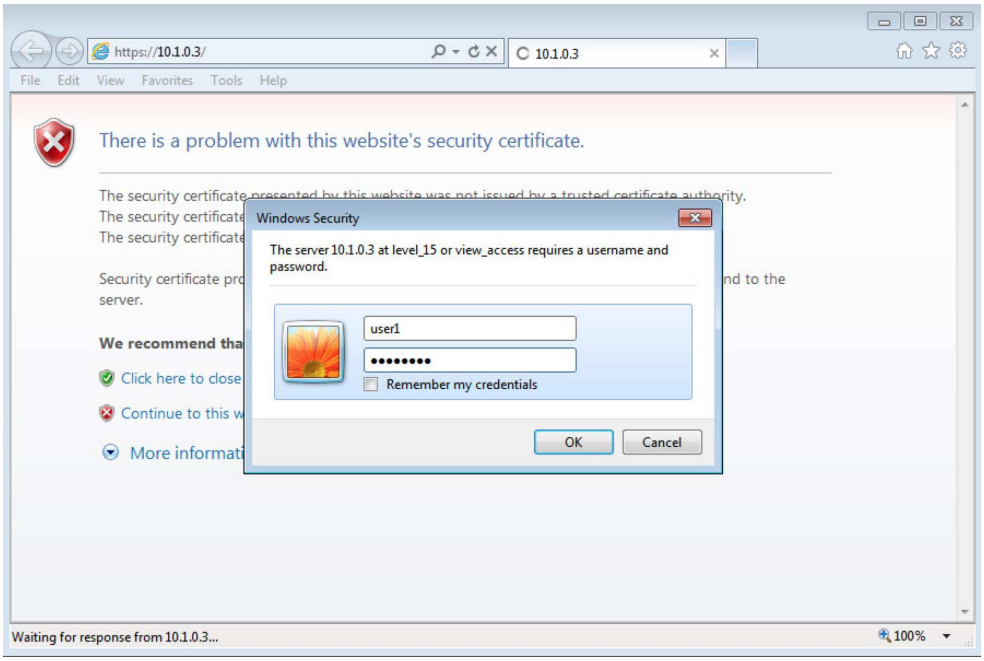
When MPP is enabled, these sessions are only allowed to come in from the link to VLAN11, and furthermore only SSH and HTTPs are allowed:

```
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#control-plane host
R3(config-cp-host)#management-interface FastEthernet0/0.11 allow https ssh
R3(config-cp-host)#end
R3#%CP-5-FEATURE: Management-Interface feature enabled on Control plane host path
R1#telnet 136.1.13.3
Trying 136.1.13.3 ... % Connection timed out; remote host not responding
R2#telnet 136.1.23.3 443
Trying 136.1.23.3, 443 ... % Connection timed out; remote host not responding
```

https://t.me/learningnets

From the above output we see that management sessions from R1 and R2 to R3 are now denied. Per the output below, however, management sessions are allowed in from the Test PC.





https://t.me/learningnets

```
R3#show management-interface
 Management interface FastEthernet0/0.11

        Protocol        Packets processed
           https               78
             ssh              222
R3#show control-plane host open-ports | i ESTA
      tcp*:443            10.1.0.100:64952
                HTTP CORE ESTABLIS      tcp*:22             10.1.0.100:64916
               SSH-Server ESTABLIS
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Device Hardening

Task

- Configure R3 to globally disable the following unnecessary services:
  - Finger
  - TCP and UDP Echo, Chargen, Discard, and Daytime
  - Bootstrap Protocol (BOOTP)
  - DHCP services
  - Packet Assembler/Disassembler (PAD)
  - HTTP server
  - Cisco Discovery Protocol (CDP)
  - TCP Client Identity Protocol Daemon (Identd)
  - Forwarding of packets with the Source Route field set
- Configure R3 to disable the following services on its Ethernet interfaces:
  - Maintenance Operation Protocol (MOP)
  - ICMP Redirects
  - ICMP Unreachables
  - ICMP Mask Replies
  - Proxy-ARP
  - Directed Broadcasts

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

https://t.me/learningnets

Configuration

R3:

```
no service finger
no ip finger
no service udp-small-servers
no service tcp-small-servers
no ip bootp server
no service dhcp
no service pad
no ip http server
no cdp run
no ip identd
no ip source-route
!
interface FastEthernet0/0
 no mop enabled
!
interface FastEthernet0/0.11
 no ip redirects
 no ip unreachables
 no ip mask-reply
 no ip proxy-arp
 no ip directed-broadcast
!
interface FastEthernet0/0.13
 no ip redirects
 no ip unreachables
 no ip mask-reply
 no ip proxy-arp
 no ip directed-broadcast
!
interface FastEthernet0/0.23
 no ip redirects
 no ip unreachables
 no ip mask-reply
 no ip proxy-arp
 no ip directed-broadcast
```

https://t.me/learningnets

Verification

Services that are unneeded but enabled by default on the router can leave it vulnerable to different types of malicious attacks. For example, with **ip directed-broadcast** on, the router can be an amplifier for either a Smurf or Fraggle attack; ICMP redirects, ICMP unreachables, ICMP mask replies, CDP, and Proxy ARP can be used to network mapping attacks, just to name a few.

For more information on these unneeded services, see the white paper Cisco Guide to Harden Cisco IOS Devices. Many of these services can also be found by using the **auto secure** feature in IOS, as seen below:

```
R3#auto secure
                --- AutoSecure Configuration ---

*** AutoSecure configuration enhances the security of
the router, but it will not make it absolutely resistant
to all security attacks ***

AutoSecure will modify the configuration of your device.
All configuration changes will be shown. For a detailed
explanation of how the configuration changes enhance security
and any possible side effects, please refer to Cisco.com for
Autosecure documentation.
At any prompt you may enter '?' for help.
Use ctrl-c to abort this session at any prompt.

Gathering information about the router for AutoSecure

Is this router connected to internet? [no]: yes
Enter the number of interfaces facing the internet [1]:

Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/0            unassigned      YES unset  up                    up
FastEthernet0/0.11         10.1.0.3        YES manual up                    up
FastEthernet0/0.13         136.1.13.3      YES manual up                    up
FastEthernet0/0.23         136.1.23.3      YES manual up                    up
FastEthernet0/1            unassigned      YES unset  administratively down down
Loopback0                  150.1.3.3       YES manual up                    up
Enter the interface name that is facing the internet: FastEthernet0/0.13

Securing Management plane services...

Disabling service finger
```

https://t.me/learningnets

```
Disabling service pad
Disabling udp & tcp small servers
Enabling service password encryption
Enabling service tcp-keepalives-in
Enabling service tcp-keepalives-out
Disabling the cdp protocol

Disabling the bootp server
Disabling the http server
Disabling the finger service
Disabling source routing
Disabling gratuitous arp

Here is a sample Security Banner to be shown
at every access to device. Modify it to suit your
enterprise requirements.

Authorized Access only
  This system is the property of So-&-So-Enterprise.
  UNAUTHORIZED ACCESS TO THIS DEVICE IS PROHIBITED.
  You must have explicit permission to access this
  device. All activities performed on this device
  are logged. Any violations of access policy will result
  in disciplinary action.

Enter the security banner {Put the banner between
k and k, where k is any character}:
"Authorized Access only
  This system is the property of So-&-So-Enterprise.
  UNAUTHORIZED ACCESS TO THIS DEVICE IS PROHIBITED.
  You must have explicit permission to access this
  device. All activities performed on this device
  are logged. Any violations of access policy will result
  in disciplinary action."
Enable secret is either not configured or
 is the same as the enable password
Enter the new enable secret:
Confirm the enable secret :
Enter the new enable password:
Confirm the enable password:
Configuring AAA local authentication
Configuring console, Aux and vty lines for
local authentication, exec-timeout, transport
Securing device against Login Attacks
Configure the following parameters
```

https://t.me/learningnets

```
Blocking Period when Login Attack detected: 10

Maximum Login failures with the device: 20

Maximum time period for crossing the failed login attempts: 30

Configure SSH server? [yes]:

Configuring interface specific AutoSecure services
Disabling the following ip services on all interfaces:

 no ip redirects
 no ip proxy-arp
 no ip unreachables
 no ip directed-broadcast
 no ip mask-reply
Disabling mop on Ethernet interfaces

Securing Forwarding plane services...

Enabling unicast rpf on all interfaces connected
to internet

Configure CBAC Firewall feature? [yes/no]: yes
This is the configuration generated:

no service finger
no service pad
no service udp-small-servers
no service tcp-small-servers
service password-encryption
service tcp-keepalives-in
service tcp-keepalives-out
no cdp run
no ip bootp server
no ip http server
no ip finger
no ip source-route
no ip gratuitous-arps
no ip identd
banner motd ^CAuthorized Access only
  This system is the property of So-&-So-Enterprise.
  UNAUTHORIZED ACCESS TO THIS DEVICE IS PROHIBITED.
  You must have explicit permission to access this
  device. All activities performed on this device
  are logged. Any violations of access policy will result
```

https://t.me/learningnets

```
  in disciplinary action.^C
security passwords min-length 6
security authentication failure rate 10 log
enable secret 5 $1$UHAJ$WlV9/FS.6wEuEIOmDc4Gk.
enable password 7 13061E010803557878707D
aaa new-model
aaa authentication login local_auth local
line console 0
 login authentication local_auth
 exec-timeout 5 0
 transport output telnet
line aux 0
 login authentication local_auth
 exec-timeout 10 0
 transport output telnet
line vty 0 4
 login authentication local_auth
 transport input telnet
line tty 1
 login authentication local_auth
 exec-timeout 15 0
login block-for 10 attempts 20 within 30
crypto key generate rsa general-keys modulus 1024
ip ssh time-out 60
ip ssh authentication-retries 2
line vty 0 4
 transport input ssh telnet
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
logging facility local2
logging trap debugging
service sequence-numbers
logging console critical
logging buffered
interface FastEthernet0/0
 no ip redirects
 no ip proxy-arp
 no ip unreachables
 no ip directed-broadcast
 no ip mask-reply
 no mop enabled
interface FastEthernet0/0.11
 no ip redirects
 no ip proxy-arp
 no ip unreachables
 no ip directed-broadcast
```

https://t.me/learningnets

```
 no ip mask-reply
interface FastEthernet0/0.13
 no ip redirects
 no ip proxy-arp
 no ip unreachables
 no ip directed-broadcast
 no ip mask-reply
interface FastEthernet0/0.23
 no ip redirects
 no ip proxy-arp
 no ip unreachables
 no ip directed-broadcast
 no ip mask-reply
interface FastEthernet0/1
 no ip redirects
 no ip proxy-arp
 no ip unreachables
 no ip directed-broadcast
 no ip mask-reply
 no mop enabled
access-list 100 permit udp any any eq bootpc
interface FastEthernet0/0.13
 ip verify unicast source reachable-via rx allow-default 100
ip inspect audit-trail
ip inspect dns-timeout 7
ip inspect tcp idle-time 14400
ip inspect udp idle-time 1800
ip inspect name autosec_inspect cuseeme timeout 3600
ip inspect name autosec_inspect ftp timeout 3600
ip inspect name autosec_inspect http timeout 3600
ip inspect name autosec_inspect rcmd timeout 3600
ip inspect name autosec_inspect realaudio timeout 3600
ip inspect name autosec_inspect smtp timeout 3600
ip inspect name autosec_inspect tftp timeout 30
ip inspect name autosec_inspect udp timeout 15
ip inspect name autosec_inspect tcp timeout 3600
ip access-list extended autosec_firewall_acl
 permit udp any any eq bootpc
 deny ip any any
interface FastEthernet0/0.13
 ip inspect autosec_inspect out
 ip access-group autosec_firewall_acl in
!
end

Apply this configuration to running-config? [yes]: no
```

https://t.me/learningnets

```
R3#
```

It is also possible to use the **auto secure** function for only a subset of features, depending on the required end result:

```
R3#auto secure ?

  firewall       AutoSecure Firewall
  forwarding     Secure Forwarding Plane
  full           Interactive full session of AutoSecure
  login          AutoSecure Login
  management     Secure Management Plane
  no-interact    Non-interactive session of AutoSecure
  ntp            AutoSecure NTP
  ssh            AutoSecure SSH
  tcp-intercept  AutoSecure TCP Intercept
  <cr>
```

The state of interface level services can be verified by using the command `show ip interface` , as follows:

```
R3#show ip interface fastEthernet 0/0.11 | i ICMP|Proxy|broadcast

  Directed broadcast forwarding is disabled
  Proxy ARP is disabled
  Local Proxy ARP is disabled
  ICMP redirects are never sent
  ICMP unreachables are never sent
  ICMP mask replies are never sent
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS CPU Protection Mechanisms

Task

- Configure RIPv2 routing between R1, R2, and R3.
- Configure R3 with the following filtering policy for traffic received from R2:
  - To protect against TTL expiry attacks, deny IP packets with a TTL value of less than 5.
  - Deny traffic that contains any IP options.
  - Dropped traffic should be logged to the syslog server located on the Test PC.
- Configure the following memory reservations on R3:
  - 2000 kilobytes of memory should be reserved for critical notifications.
  - Set a threshold of 20000 KB free processor memory before the router issues notifications.
  - Set a threshold of 20000 KB free I/O memory before the router issues notifications.
- Enable CPU threshold notifications on R3 as follows:
  - Trigger a notification when the total CPU utilization rises above 80 percent for a period of 5 seconds or longer.
  - Trigger a second notification when the total CPU utilization falls below 30 percent for a period of 5 seconds or longer.
  - Log these messages to the syslog server located on the Test PC.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.

https://t.me/learningnets

If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

https://t.me/learningnets

R1:

```
router rip
 version 2
 network 0.0.0.0
 no auto-summary
```

R2:

```
router rip
 version 2
 network 0.0.0.0
 no auto-summary
```

R3:

```
ip access-list extended TTL_AND_IP_OPTIONS_FILTER
 permit udp host 136.1.23.2 eq rip host 224.0.0.9 eq rip
 deny ip any any ttl lt 5 log
 deny ip any any option any-options log
 permit ip any any
!
interface FastEthernet0/0.23
 ip access-group TTL_AND_IP_OPTIONS_FILTER in
!
memory reserve critical 2000
memory free low-watermark processor 20000
memory free low-watermark IO 20000
!
process cpu threshold type total rising 80 interval 5 falling 30 interval 5
!
logging 10.1.0.100
logging trap informational
!
router rip
 version 2
 network 0.0.0.0
 no auto-summary
```

Verification

When the router receives packets with a TTL value of 0 or 1, they must be process switched. Because process switching uses the general-purpose CPU, as opposed to hardware acceleration, this can result in high CPU utilization. Therefore, another

https://t.me/learningnets

possible attack against the router is to flood it with lots of packets with a very low TTL. Filtering traffic with a low TTL (such as a TTL of 5 or lower) at the ingress of the network can help to protect the control plane of other downstream infrastructure devices from TTL expiry attacks.

> The local router's CPU cannot be protected against TTL expiry attacks by applying an ACL to the ingress interface. Instead, an ACL for TTL 0 or 1 packets should be applied through Control Plane Policing (CoPP). For more information, see the white paper ACL Support for Filtering on TTL Value.

Note that link local control plane traffic, such as our RIPv2 routing in this case, typically uses low TTL values. Therefore, we must account for this first in our ACL; otherwise, the routing control plane will be broken:

```
R3#
%SEC-6-IPACCESSLOGP: list TTL_AND_IP_OPTIONS_FILTER denied udp 136.1.23.2(0) -> 224.0.0.9(0), 11 packets

R3#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 2 subnets
R        150.1.1.1 [120/1] via 136.1.13.1, 00:00:09, FastEthernet0/0.13

R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#ip access-list extended TTL_AND_IP_OPTIONS_FILTER
R3(config-ext-nacl)# 1 permit udp host 136.1.23.2 eq rip host 224.0.0.9 eq rip
R3(config-ext-nacl)#end
R3#show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

https://t.me/learningnets

```
        o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
        + - replicated route, % - next hop override

Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 3 subnets
R        150.1.1.1 [120/1] via 136.1.13.1, 00:00:09, FastEthernet0/0.13
R        150.1.2.2 [120/1] via 136.1.23.2, 00:00:06, FastEthernet0/0.23
```

Like packets with a low TTL, packets with the IP Options field set can result in a DoS attack against the router in certain circumstances. Because the vast majority of legitimate IP flows do not need the Options field populated, a simple method to mitigate against these types of attacks is to filter packets with an ACL if they have any Options set.

> For more information on this, see the white paper Identifying and Mitigating Exploitation of the Crafted IP Option Vulnerability.

As a reactive measure to these types of attacks, the router can be configured to generate a syslog message when its CPU utilization reaches a certain threshold. However, if the router is already overloaded, it's possible that the syslog process will not be able to allocate memory to generate logs. This is where the `memory reserve` command is useful, to ensure that critical notifications can still be generated. Likewise, the router can be configured to generate logs when memory is low with the `memory free low-watermark` commands. To demonstrate this, we can flood the router with packets that cause process switching, and therefore cause high CPU utilization:

https://t.me/learningnets



When R3's CPU utilization reaches the threshold, a syslog message is generated as well for the memory:

```
R3#sh proc cpu | in utilization
CPU utilization for five seconds: 77%
/34%; one minute: 39%; five minutes: 26%
R3#sh proc cpu | in utilization
```

https://t.me/learningnets

```
CPU utilization for five seconds: 80%
/33%; one minute: 50%; five minutes: 30%



R3#
%SYS-4-FREEMEMLOW: Free Memory has dropped below low watermark Pool: I/O  Free: 6524080  Threshold: 20480000


R3#
%SYS-1-CPURISINGTHRESHOLD: Threshold: Total CPU Utilization(Total/Intr): 80%/32%, Top 3 processes(Pid/Util):  142/46
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS IPv4 Selective Packet Discard

Task

- Configure Selective Packet Discard on R1 as follows:
  - Enable SPD in aggressive mode.
  - Increase R1's input queue size on its link to VLAN 13 to twice the default.
  - Increase the amount of the memory headroom for IGP packets to 150 buffers.
  - Set the headroom for BGP packets to 120 packets.
  - Start dropping low-priority packets randomly when the input queue is 50% full.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

https://t.me/learningnets

```
R1:

spd extended-headroom 150
spd headroom 120
ip spd mode aggressive
ip spd queue threshold minimum 75 maximum 150
!
interface GigabitEthernet0/0
 hold-queue 150 in
```

Verification

Selective Packet Discard is the queue management technique for interface input queueing. The SPD commands are hidden in the IOS parser, but you can see them in the running configuration after you enter them:

```
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#spd ?
% Unrecognized command
R1(config)#ip spd ?
% Unrecognized command
R1(config)#do show running-config | include spd

spd extended-headroom 150
spd headroom 120
ip spd mode aggressive
ip spd queue threshold minimum 75 maximum 150
```

Every physical interface has an input FIFO queue. The router uses this queue to buffer packets going to the route processor. Usually these packets include control plane packets, such as Layer 2 keepalives (for example, HDLC/PPP keepalives), IGP packets (OSPF, ISIS, etc.), and BGP packets. The routing protocol packets are classified based on their default IP precedence of 6 or higher. In addition to control plane packets, the input queue holds other packets destined to the route processor, such as packets with an expired TTL, wrong header length, wrong checksum, or non-existent local UDP port numbers. The latter packets are malformed, in the sense that they require the router to generate an ICMP error message in response. The input queue holds packets that are to be process-switched, which is uncommon on modern CEF-based systems.

SPD input queueing is desirable for several reasons. The first is control plane

https://t.me/learningnets

security. It’s possible to block the router’s input queue with a high rate of malformed packets, which effectively blocks legitimate routing traffic. The result is a control plane DoS against the router.

The next reason is layer 2 keepalive, IGP, and BGP traffic separation. Large BGP tables generate a considerable number of updates. These updates could potentially block the input queue for some time, preventing the router from receiving keepalive packets or IGP updates. This may result in IGP adjacency flapping or layer 2 link loss reports.

The third reason is due to issues with process switching. If for some reason CEF is disabled, the **IP INPUT** process can result in regular IP traffic blocking the input queues, causing a loss of the control plane. SPD prevents this through input drops.

So how does SPD work? First, the input queue consists of two parts. One part is the regular hold queue, which is visible through the show interface command, and the other part is the priority queue, which stores routing updates and keepalives. The processor serves the priority queue first until it is empty, and then switches to the regular hold-queue. Additionally, the priority queue consists of two parts, the SPD Headroom and the SPD Extended Headroom. The Extended Headroom queue is emptied before the SPD Headroom in a priority manner. Specifically, input packets are queued as follows:

- Layer 2 keepalives and IGP packets go to the SPD Extended Headroom. If there is no space available in the SPD Extended Headroom, packets go to the SPD Headroom.
- As a last resort, if both the Extended Headroom and Headroom are full, these packets go to the regular Hold Queue.

BGP updates go directly to SPD Headroom. If the SPD Headroom is full, BGP packets hit the Hold Queue. All other IP packets (malformed or process-switched) go to the Hold Queue. The result is that L2 keepalive/IGP packets are serviced first, BGP next, and other packets last. Although the Hold Queue is FIFO, it uses the RED drop procedure. Two threshold variables (Minimum and Maximum) set for hold queue define the random drop behavior. If the current hold queue length is less than the Minimum Threshold, packets are never dropped. If the queue length grows beyond Minimum, but is less than Maximum, every new packet is randomly dropped with the probability proportional to queue depth:

Prob = (QueueDepth – MinimumThresh)/(MaximumThresh-MinimumThresh)

If the queue depth is above Maximum Threshold, SPD drops every new incoming packet. Note the important fact that SPD thresholds are global for all queues. SPD computes Min and Max thresholds based on the lowest hold-queue size in the system. Therefore, if you set the hold queue size lower on some interfaces, you will

https://t.me/learningnets

affect all other interface drop thresholds.

Finally, SPD has two modes of operation: normal and aggressive. They differ in their treatment of malformed packets (packets that require the router to generate ICMP responses). When SPD is set for normal mode (the default), it treats malformed packets as it would all regular IP packets: It places them in the hold queue, subject to random drop. However, in aggressive mode, the malformed packets are dropped as soon as the hold queue grows above the minimum threshold. Effectively, SPD Aggressive mode replaces the random drop for malformed packets with an unconditional drop.

SPD configuration can be verified as follows:

```
R1#show interface GigabitEthernet0/0
GigabitEthernet0/0 is up, line protocol is up
  Hardware is CN Gigabit Ethernet, address is 0007.7d3d.94e8 (bia 0007.7d3d.94e8)
  Internet address is 136.1.13.1/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full Duplex, 100Mbps, media type is RJ45
  output flow-control is unsupported, input flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters never Input queue: 0/150/0/0 (size/max/drops/flushes)
; Total output drops:
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
<snip>
R1#show ip spd
Current mode: normal. Queue min/max thresholds: 75/150, Headroom: 120, Extended Headroom: 150
IP normal queue: 0, priority queue: 0. SPD special drop mode: aggressively drop bad packets
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Transit Traffic Control with FPM

Task

- Configure Deep Packet Inspection on R3 to match a custom pattern attack known as the W32.Blaster Worm, using the following match criteria:
  - IP packet length exceeding 402 bytes
  - UDP port 69
  - Pattern match 0x20a29010 at 50 bytes from start of IP header, to match on 4 bytes
- You should not need to configure an ACL to complete this task.
- Assume that the attack is coming from the direction of R1.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

R3:

```
load protocol system:fpm/phdf/ip.phdf
load protocol system:fpm/phdf/udp.phdf
!
class-map type access-control match-all W32.BLASTER_CLASS
 match field UDP dest-port eq 0x45
 match field IP length gt 0x192
 match start l3-start offset 50 size 4 eq 0x20A29010
```

https://t.me/learningnets

```
!
class-map type stack match-all UDP
 match field IP protocol eq 0x11 next UDP
!
policy-map type access-control W32.BLASTER_POLICY
 class W32.BLASTER_CLASS
  drop
!
policy-map type access-control ACCESS_CONTROL_POLICY
 class UDP
  service-policy W32.BLASTER_POLICY
!
interface FastEthernet0/0.13
 service-policy type access-control input ACCESS_CONTROL_POLICY
```

Verification

Flexible Packet Matching (FPM) is a feature that allows for granular packet inspection in Cisco IOS routers. FPM is similar to Network Based Application Recognition (NBAR), that is, the `match protocol` command in QoS configuration, but the difference is that FPM can be used to match applications or traffic patterns that aren't yet categorized by NBAR. The main application of FPM is to mitigate against zero day attacks, meaning that you can write your own pattern matching before new NBAR or IPS signatures are introduced to address a new threat.

To accomplish this, FPM allows you to match any string, byte, or even bit at any position in the IP (or theoretically non-IP) packet. FPM configuration uses the Cisco Common Classification Policy Language (C3PL), which you may already be familiar with through other configurations such as the Zone Based Policy Firewall.

Before configuring FPM, you must first load the Protocol Header Data File (PHDF). This is necessary because without it, the router can't interpret where it is looking in a header. When the PHDF is loaded, you can proceed to configure the class-maps, policy-maps, and service policies as required.

> For more information on FPM, see Understanding Flexible Packet Matching on the INE Blog and the white paper Getting Started with Cisco IOS Flexible Packet Matching.

When the policy is complete, it can be verified similar to other C3PL based configurations, as seen below:

```
R3#show policy-map type access-control interface FastEthernet0/0.13
 FastEthernet0/0.13
```

https://t.me/learningnets

```
  Service-policy access-control input: ACCESS_CONTROL_POLICY

    Class-map: UDP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps
      Match: field IP protocol eq 0x11 next UDP

      Service-policy access-control : W32.BLASTER_POLICY

        Class-map: W32.BLASTER_CLASS (match-all)
          0 packets, 0 bytes
          5 minute offered rate 0 bps Match: start l3-start offset 50 size 4 eq 0x20A29010
Match: field UDP dest-port eq 0x45
Match: field IP length gt 0x192
drop



        Class-map: class-default (match-any)
          0 packets, 0 bytes
          5 minute offered rate 0 bps, drop rate 0 bps
          Match: any

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

The PHDF files define the respective protocol headers, which can be viewed once activated with command `show protocols phdf`：

```
R3#show protocols phdf UDP

Protocol ID: 2
Protocol name: UDP
Description: UDP-Protocol
Original file name: system:fpm/phdf/udp.phdf
Header length: 8
Constraint(s):
Total number of fields: 5
 Field id: 0, source-port, UDP-Source-Port
        Fixed offset. offset 0
        Constant length. Length: 16
 Field id: 1, dest-port, UDP-Destination-Port
        Fixed offset. offset 16
        Constant length. Length: 16
```

https://t.me/learningnets

```
Field id: 2, length, UDP-Length
        Fixed offset. offset 32
        Constant length. Length: 16
Field id: 3, checksum, UDP-Checksum
        Fixed offset. offset 48
        Constant length. Length: 16
Field id: 4, payload-start, UDP-Payload-Start

        Fixed offset. offset 64
        Constant length. Length: 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Congestion Management

Task

- On R3, create a policy to match ICMP traffic coming from VLAN 23 and mark it with a DSCP value of AF12.
- On R3, create a policy to match ICMP traffic coming from VLAN 11 and mark it with a DSCP value of AF23.
- Configure the following traffic policy outbound on R3's VLAN 13 interface:
  - AF12 Traffic should be given 25% of the interface bandwidth.
  - AF23 Traffic should be given 15% of the interface bandwidth.
  - Voice traffic with a DSCP value of EF should be prioritized and given 35% bandwidth.
  - Remaining traffic should be fair-queued.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Configuration

Class-maps are configured for both traffic marking and marked traffic matching. For each traffic class, required queuing is configured and bandwidth reserved for cases of congestion. Note that in case of congestion, priority traffic is policed at the upper limit (35%), whereas traffic from other classes is not and will exceed the allocated bandwidth by taking from the class-default.

https://t.me/learningnets

R3:

```
ip access-list extended ICMP
 permit icmp any any
!
class-map ICMP
 match access-group name ICMP
!
policy-map MARK_VLAN23
 class ICMP
  set dscp AF12
!
interface FastEthernet0/0.23
 service-policy input MARK_VLAN23
!
!
policy-map MARK_VLAN11
 class ICMP
  set dscp AF23
!
interface FastEthernet0/0.11
 service-policy input MARK_VLAN11
!
!
class-map MATCH_AF12
 match dscp AF12
!
class-map MATCH_AF23
 match dscp AF23
!
class-map MATCH_EF
 match dscp EF
!
policy-map QUEUE_CHILD
 class MATCH_AF12
  bandwidth percent 25
 class MATCH_AF23
  bandwidth percent 15
 class MATCH_EF
  priority percent 35
 class class-default
  fair-queue
!
policy-map QUEUE_PARENT
 class class-default
  shape average 100m
```

https://t.me/learningnets

```
   service-policy QUEUE_CHILD
!
interface FastEthernet0/0.13
 service-policy output QUEUE_PARENT
```

Verification

Because logical interfaces, such as subinterfaces, do not support a state of congestion (this is related to the main physical interface), configuring CBWFQ at the subinterface level is not allowed and the router will log an error message. To simulate congestion, instead of applying just queuing, we configure shaping, but because no upper limit has been specified, we configure the actual interfaec speed of 100 Mpbs.

```
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#interface FastEthernet0/0.13
R3(config-subif)#service-policy output QUEUE
CBWFQ : Not supported on subinterfaces
```

To learn more about applying QoS features at the subinterface level, read [Applying QoS Features to Ethernet Subinterfaces](#).

Verify the configure marking:

```
R3#show policy-map interface fastEthernet0/0.11
 FastEthernet0/0.11

  Service-policy input: MARK_VLAN11

    Class-map: ICMP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps Match: access-group name ICMP
QoS Set
dscp af23
          Packets marked 0

    Class-map: class-default (match-any)
      351 packets, 25502 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any

R3#show policy-map interface fastEthernet0/0.23
```

https://t.me/learningnets

```
 FastEthernet0/0.23

  Service-policy input: MARK_VLAN23

    Class-map: ICMP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps Match: access-group name ICMP
QoS Set
dscp af12

          Packets marked 0

    Class-map: class-default (match-any)
      9 packets, 1098 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

Verify the configured shaping and allocated bandwidth:

```
R3#show policy-map interface fastEthernet0/0.13
 FastEthernet0/0.13

  Service-policy output: QUEUE_PARENT

    Class-map: class-default (match-any)
      5 packets, 520 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 5/610
      shape (average) cir 100000000, bc 400000, be 400000 target shape rate 100000000

      Service-policy : QUEUE_CHILD

        queue stats for all priority classes:

          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0

        Class-map: MATCH_AF12 (match-all)
          0 packets, 0 bytes
          5 minute offered rate 0 bps, drop rate 0 bps
```

https://t.me/learningnets

```
Match:  dscp af12 (12)

          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0 bandwidth 25% (25000 kbps)

        Class-map: MATCH_AF23 (match-all)
          0 packets, 0 bytes
          5 minute offered rate 0 bps, drop rate 0 bps Match:  dscp af23 (22)
          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0 bandwidth 15% (15000 kbps)

        Class-map: MATCH_EF (match-all)
          0 packets, 0 bytes
          5 minute offered rate 0 bps, drop rate 0 bps Match:  dscp ef (46)
 Priority: 35% (35000 kbps), burst bytes 875000, b/w exceed drops: 0


        Class-map: class-default (match-any)
          5 packets, 520 bytes
          5 minute offered rate 0 bps, drop rate 0 bps
          Match: any
          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops/flowdrops) 0/0/0/0
          (pkts output/bytes output) 5/610 Fair-queue: per-flow queue limit 16
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS File System Security

Task

- Configure R1 to mitigate the risk of malicious attempts to erase the contents of flash by securely archiving copies of the running IOS image and configuration files.
- Configure R2 and R3 to prevent corrupted IOS images from being loaded by verifying the integrity of IOS images automatically.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

Most current versions of IOS running on routers support features to mitigate the risk of malicious attempts to erase the contents of persistent storage (NVRAM and FLASH) and features to prevent corrupted IOS images from being loaded. These features are known as Cisco IOS Resilient Configuration and Cisco IOS Image Verification, respectively. The IOS Resilient Configuration feature enables a router to securely archive copies of the running IOS image and configuration files. In this way, if the running files are tampered with or erased, you can restore them quickly using the secure copies and, as a result, minimize downtime. The IOS Image Verification feature allows you to automatically verify the integrity of IOS images. This was traditionally an optional step in the IOS version upgrade process. IOS Image Verification is now automated so that the integrity of any IOS image file downloaded is automatically verified. Although IOS Resilient Configuration can be enabled from a remote telnet/ssh session, console access is required to disable it.

https://t.me/learningnets

Another implicit behavior is that to upgrade the IOS image on the router, console access is required so that you can first disable the feature, perform the upgrade, and re-enable the feature so that the new secure image is created. You would use the following commands to activate these features:

- secure boot-config (IOS Resilient Configuration)
- secure boot-image (IOS Resilient Configuration)
- file verify auto (IOS Image Verification)

The `secure boot-config` command is supported only on routers configured with a PCMCIA Advanced Technology Attachment (ATA) disk. The archived configuration is hidden and cannot be viewed, copied, modified, or removed using EXEC mode commands (although it may be viewed in ROMMON mode). The archived configuration will even survive a disk format operation.

When the command `secure boot-image` is first enabled, the running IOS image (as displayed in the show version command output) is securely archived in the persistent storage. This command is also supported only on routers configured with a PCMCIA ATA disk. Again, the archived image is hidden and cannot be viewed, copied, modified, or removed with EXEC mode commands. The archived image will even survive a disk format operation.

To enable automatic image verification, use the `file verify auto` command in IOS global configuration mode. Image verification is disabled by default within IOS. With this command enabled, each IOS image that is copied or reloaded will be automatically verified. This includes computing a local MD5 hash of the image and comparing it to the MD5 hash embedded within the image. When this verification process is run, the CCO MD5 hash is also displayed, which you can manually compare against the MD5 digest posted on www.cisco.com from where the image was downloaded. If the MD5 hashes do not match, image verification fails and the image will not be loaded or copied. This command cannot be used while `secure boot-image` is in use because the image is no longer visible at **show flash:**. Additionally, you can run a manual verification from the CLI as shown below:

```
R1#verify /md5 flash:/c2900-universalk9-mz.SPA.152-3.T3.bin
.....................................................................................................................
<output omitted>
.....................................................................Done!
verify /md5 (flash0:/c2900-universalk9-mz.SPA.152-2.T.bin) = 941bbe5470359ec46fb265d787b695a3
```

https://t.me/learningnets

Configuration

```
R1:
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#secure boot-config
%IOS_RESILIENCE-5-CONFIG_RESIL_ACTIVE: Successfully secured config archive [flash0:.runcfg-20130909-181912.ar]
R1(config)#secure boot-image
%IOS_RESILIENCE-5-IMAGE_RESIL_ACTIVE: Successfully secured running image

R2:
!
file verify auto

R3:

!
file verify auto
```

Verification

Attempts to disable the IOS Resilient Configuration from a telnet session will be denied and an error message will be logged:

```
R1(config)#no secure boot-image
 %You must be logged on the console to apply this command
```

Verify the IOS Resilient Configuration:

https://t.me/learningnets

```
R1#show secure bootset

IOS resilience router id FCZ16436080

IOS image resilience version 15.2 activated at 18:19:21 UTC Mon Sep 9 2013
Secure archive flash0:c2900-universalk9-mz.SPA.152-3.T3.bin type is image (elf) []

  file size is 87305860 bytes, run size is 87484912 bytes
  Runnable image, entry point 0x81000000, run from ram

IOS configuration resilience version 15.2 activated at 18:19:12 UTC Mon Sep 9 2013
Secure archive flash0:.runcfg-20130909-181912.ar type is config

configuration archive size 1403 bytes
```

Verify the IOS Image Verification:

```
R3#write memory
Building configuration...

[OK]R3#reload
Verifying file integrity of flash:/c2800nm-adventerprisek9-mz.151-3.T2.bin.....................................
Embedded Hash   MD5 : 49BD41F0589BB0F202DADB5E7A602ED6
Computed Hash   MD5 : 49BD41F0589BB0F202DADB5E7A602ED6

CCO Hash        MD5 : 0BFF67580330DA84034BD189901D94FE
Signature Verified

Proceed with reload? [confirm]
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS RITE and RMON

Task

- Configure R3 to export telnet traffic from VLAN 23 to the TEST-PC. You should be able to view this data in Wireshark. The sample rate should be 1 in every 3 packets.
- Ensure that all log and debug messages sent from R3 to the syslog server at 10.1.0.100 include the date and time of the error or event.
- Using RMON, configure R1 to monitor its own CPU utilization and trigger a log message when the value exceeds 70% or falls below 40% within a minute.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. Note that the **IOS Resilient Configuration** requires manual removal of the commands.

Overview

Similar to Switched Port Analyzer (SPAN) ports available with multilayer Ethernet switches such as the Catalyst product family, software-based IOS routers also allow for packet capture and export to traffic analysis systems. This is often useful for classifying an attack and determining the required mitigation action when TCP/IP header information is not sufficient. This also reduces the need to deploy traffic analysis systems inline and on the router itself, such as enabling IOS Intrusion Prevention System (IPS) functions. Using IP Traffic Export, traffic can be selectively captured using classification ACLs and packet sampling, as well as traffic direction (ingress or bidirectional) on the capturing interface. This feature is generally enabled in response to an attack to facilitate attack classification and mitigation, because its

https://t.me/learningnets

operation can be processor intensive and could negatively impact overall router performance. Configuring RITE (Router IP Traffic Export) may actually have a greater performance impact than the attack itself.

When the packets are exported, the destination MAC address is modified to that of the analyzer. The MAC address and interface where the analyzer is reachable are the only required parameters in the profile. Wireshark interprets the stream as having lost segments because we are not exporting all packets. Also notice that the destination MAC address has been modified according to the configuration.

The Remote Monitoring (RMON) feature is used to monitor specified SNMP MIB variables and generate events (syslog messages or SNMP traps) when a threshold or change threshold is crossed. For example, you can configure a router to monitor its own CPU utilization and trigger an SNMP trap or syslog message when the value exceeds or goes below a defined threshold. RMON configuration consists of defining events and creating an alarm for a known MIB variable.

Configuration

Find the MAC address of TEST-PC, based on the assigned rack, this will differ:

```
R3#show ip arp 10.1.0.100
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.0.100              0  48f8.b32e.2423
  ARPA   FastEthernet0/0.11


R3:

ip access-list extended RITE
 permit tcp any any eq telnet
 permit tcp any eq telnet any
!
ip traffic-export profile RITE
 interface FastEthernet0/0.11
 bidirectional
 incoming access-list RITE
 outgoing access-list RITE
 mac-address 48f8.b32e.2423
 incoming sample one-in-every 3
 outgoing sample one-in-every 3
 exit
!
interface FastEthernet0/0.23
 ip traffic-export apply RITE
```

https://t.me/learningnets

```
!
!
service timestamps log datetime msec localtime
service timestamps debug datetime msec localtime
logging host 10.1.0.100
logging trap informational
!
!
rmon event 1 log description "CPU on R3 has exceeded 70% threshold"
rmon event 2 log description "CPU on R3 has is below 40% threshold"
rmon alarm 1 lsystem.57.0 60 absolute rising-threshold 70 1 falling-threshold 40 2
```

Verification

Open Wireshark on the TEST-PC (ensure to select the USB interface, the default option of **Capture all in promiscuous mode** can be unchecked) and initiate a telnet connection from R2 to R3:

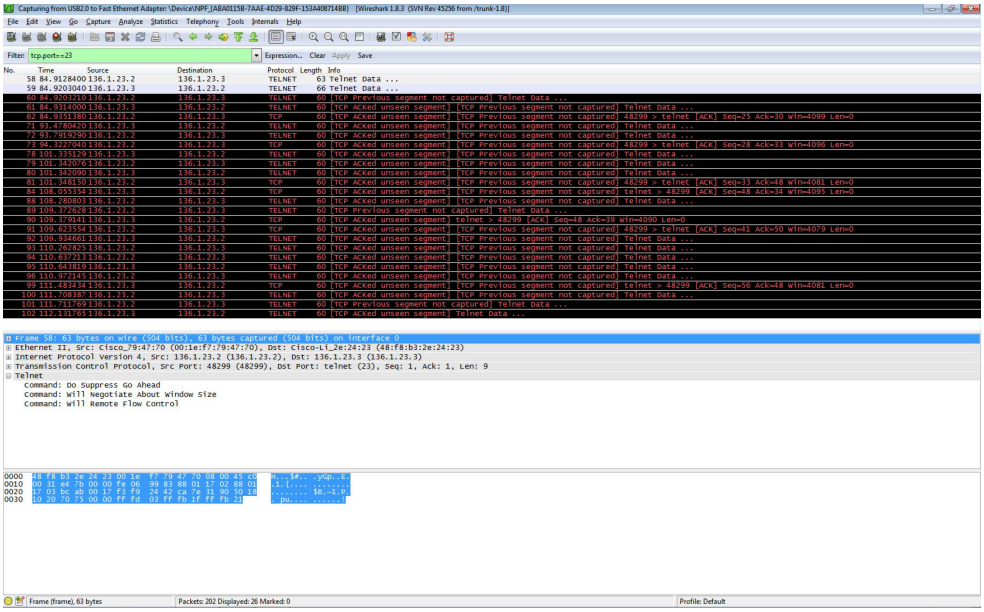
```
R2#telnet 136.1.23.3

Trying 136.1.23.3 ... Open

R3>enable
Password:
R3#
```

Note on the Wireshark capture that, because only one packet out of three is captured and sent to the analyzer, Wireshark does not see the complete flow and thus complains as such:

https://t.me/learningnets



Look at the statistics on R3. As you can see below, the counters are incrementing:

```
R3#show ip traffic-export

Router IP Traffic Export Parameters
Monitored Interface              FastEthernet0/0
        Export Interface                 FastEthernet0/0.11
        Destination MAC address 48f8.b32e.2423
        bi-directional traffic export is on
Output IP Traffic Export Information     Packets/Bytes Exported     11/465
        Packets Dropped            22
        Sampling Rate              one-in-every 3 packets
        Access List        RITE [named extended IP]
Input IP Traffic Export Information      Packets/Bytes Exported     15/621

        Packets Dropped            32
        Sampling Rate              one-in-every 3 packets
        Access List        RITE [named extended IP]
        Profile RITE is Active
```

Generate a UDP flood from TEST-PC to trigger the CPU rise over the configured threshold, for minimum one minute; note that it will take few minutes for the log message to be generated. Verify that log messges are also timestamped:

https://t.me/learningnets




```
R3#show proc cpu sorted | i CPU
CPU utilization for five seconds: 73%/31%; one minute: 64%
; five minutes: 26%
!
!R3#show proc cpu sorted | i CPU
CPU utilization for five seconds: 73%/30%; one minute: 68%
; five minutes: 31%
```

https://t.me/learningnets

```
!
!R3#show proc cpu sorted | i CPU
CPU utilization for five seconds: 74%/30%; one minute: 71%
; five minutes: 36%
!
!
R3#
*Sep  9 22:09:25.202: %RMON-5-RISINGTRAP: Rising threshold has been crossed because the value of lsystem.57.0 exceed
R3#
*Sep  9 22:11:25.210: %RMON-5-FALLINGTRAP: Falling threshold has been crossed because the value of lsystem.57.0 has
```

Stop the UDP flood generation from TEST-PC.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS BGP TTL Security Hack

Task

- Configure a BGP peering between R1 and R2 and advertise Loopback0 interfaces.
- R1 is member of AS 11, whereas R2 is member of AS 22.
- Both R1 and R2 are experiencing a large number of DoS attacks on TCP port 179 from spoofed IP addresses.
- Configure R1 and R2 so that the BGP session is formed only if the BGP peer is one hop away.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

The BGP TTL Security Hack (BTSH) is designed to protect the BGP (RFC1771) infrastructure from CPU-utilization based attacks. In Cisco IOS, the following is the default behavior:

https://t.me/learningnets

- Any IP packet with a TTL value greater than or equal to 1 is considered valid per RFC 791 and accepted by the router.
- iBGP sessions are initiated with a TTL value of 255, and eBGP sessions are initiated with a TTL value of 1.
- For an eBGP Multihop session to be even initiated and established, usually **ebgp-multihop"** is configured, basically modifying the default TTL value of 1.

General TTL Security Mechanism (GTSM) is defined in RFC 3682 and specifies a protection method against BGP session hijacking and resource exhaustion attacks. Generally, the BGP process listens on the TCP port 179 and accepts all TCP SYN packets destined to this port, unless they are filtered by an ACL for the source. For this reason, it is possible to generate a number of spoofed packets imitating a valid BGP session and inject false information (if the session is unauthenticated) or generate a TCP SYN-flooding attack. eBGP Multihop configuration does not offer any protection, because no inbound verifications are done before accepting the BGP packet. Typically, border network and router security policies will permit BGP protocol packets only from valid peers that are optionally using MD5 authentication. Authentication can help prevent a remote attacker from injecting spoofed BGP protocol packets into the session.

GTSM utilizes the simple fact that every router on the path to the BGP peer decrements the TTL field in IP packets by one. Based on this, it is possible to identify potentially spoofed packets by looking at their TTL field; the packets sent from attackers will have the TTL field below some configured threshold because these are from outside of your network. It is possible to define the maximum number of hop counts an eBGP session can traverse before being established. Notice that the usefulness of GTSM feature decreases as the diameter of eBGP Multihop session grows. GTSM is configured on a per-neighbor basis using the command **`neighbor <IP> ttl-security hops <hop-count>`** . For example, if configuration is `neighbor 1.1.1.1 ttl-security hops 5` , the following two actions are enforced:

- If the local router initiates the eBGP session (is the BGP client) toward 1.1.1.1, it will use a TTL vlue of 255.
- If the local router receives the eBGP session (is the BGP server) from 1.1.1.1, it will accept the session only if the TTL value is minimum 250, value coming from the maximum of 255 minus the configured hop-count value for GTSM.

This means that the feature is mutually exclusive with the eBGP MultiHop feature because eBGP Multihop sets the TTL to the configured value, whereas GTSM sets it to 255 unconditionally. More information can be found at Protecting Border Gateway Protocol for the Enterprise.

https://t.me/learningnets

Configuration

```
R1:
ip route 136.1.23.0 255.255.255.0 136.1.13.3
!
router bgp 11
 neighbor 136.1.23.2 remote-as 22
 neighbor 136.1.23.2 ttl-security hops 2
 network 150.1.1.1 mask 255.255.255.255

R2:

ip route 136.1.13.0 255.255.255.0 136.1.23.3
!
router bgp 22
 neighbor 136.1.13.1 remote-as 11
 neighbor 136.1.13.1 ttl-security hops 2
 network 150.1.2.2 mask 255.255.255.255
```

Verification

Verify that the BGP peering is UP and that configured GTSM settings are enforced:

```
R1#show ip bgp summary

BGP router identifier 150.1.1.1, local AS number 11
BGP table version is 3, main routing table version 3
2 network entries using 296 bytes of memory
2 path entries using 128 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 720 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
136.1.23.2      4           22       7       7        3    0    0 00:02:35        1

R1#show ip bgp neighbors | i TTL|hop
  External BGP neighbor may be up to 2 hops away.
```

https://t.me/learningnets

```
Connection is ECN Disabled, Mininum incoming TTL 253, Outgoing TTL 255



R2#show ip bgp neighbors | i TTL|hop
  External BGP neighbor may be up to 2 hops away.
Connection is ECN Disabled, Mininum incoming TTL 253, Outgoing TTL 255
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS IPv6 Selective Packet Discard

Task

- Configure R3 for SPD to match the following highlighted output:

```
R3#show ipv6 spd

Current mode: normalQueue min/max thresholds: 80/100, Headroom: 100, Extended Headroom: 20
IPv6 queues normal/priority/resolution 0/0/0
IPv6 queues high water normal/priority/resolution 3/1/1
SPD special drop mode: aggressively drop bad packets

SPD drop events: none
```

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

SPD implementation for IPv6 is very similar to the IPv4 implementation. More information can be found on IPv6 Selective Packet Discard.

Configuration

https://t.me/learningnets

```
R3:

!
ipv6 spd mode aggressive
ipv6 spd queue min-threshold 80
ipv6 spd queue max-threshold 100
spd extended-headroom 20
```

Verification

Verify that configured SPD values match the requested output:

```
R3#show ipv6 spd
Current mode: normalQueue min/max thresholds: 80/100, Headroom: 100, Extended Headroom: 20
IPv6 queues normal/priority/resolution 0/0/0
IPv6 queues high water normal/priority/resolution 3/1/1
SPD special drop mode: aggressively drop bad packets

SPD drop events: none
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

Section 1 Introduction

This section utilizes Firewall Services

Lab Topics

The first blueprint topic of the CCIE Security lab exam is System Hardening and Availability. At a high level, these topics include the following:

- Routing plane security features (e.g., protocol authentication, route filtering)
- Control Plane Policing
- Control Plane Protection and Management Plane Protection
- Broadcast control and switchport security
- Additional CPU protection mechanisms (e.g., options drop, logging interval)
- Disable unnecessary services
- Control device access (e.g., Telnet, HTTP, SSH, Privilege levels)
- Device services (e.g., SNMP, Syslog, NTP)
- Transit Traffic Control and Congestion Management

As you explore these topics in further detail, you can expand the areas in the following checklist. We recommend making it your goal to master each of these technologies.

Technology	Mastered?
Configuring Routing Protocol Authentication	
Route Filtering and Protocol-Specific Filters	

https://t.me/learningnets

Technology	Mastered?
ICMP Techniques to Reduce the Risk of ICMP-Related DoS Attacks (IP Unreachable, IP Redirect, IP Mask Reply, etc.)	
Selective Packet Discard (SPD)	
MQC and FPM Types of Service Policy on the CoPP Interface	
Broadcast Control on a Switch	
Catalyst Switch Port Security	
IPv6 Selective Packet Discard	
Cisco IOS Software-Based CPU Protection Mechanisms (Options Drop, Logging Interval, CPU Threshold)	
The Generalized TTL Security Mechanism Known as “BGP TTL Security Hack” (BTSH)	
Device Access Control (vty ACL, HTTP ACL, SSH Access, Privilege Levels)	
SNMP Security	
System Banners	
Secure Cisco IOS File Systems	
Understanding and Enabling Syslog	

https://t.me/learningnets

Technology	Mastered?
NTP with Authentication	
Role-Based CLI Views and Cisco Secure ACS Setup	
Service Authentication on Cisco IOS Software (FTP, Telnet, HTTP)	
Network Telemetry Identification and Classification of Security Events (IP Traffic Flow, NetFlow, SNMP, Syslog, RMON)	

The following technology labs are designed to cover these possible lab topics in such a way that you gain the technical expertise required to succeed in the CCIE Security Lab Exam. We strongly recommend that you focus on gaining in-depth knowledge of how each technology works.

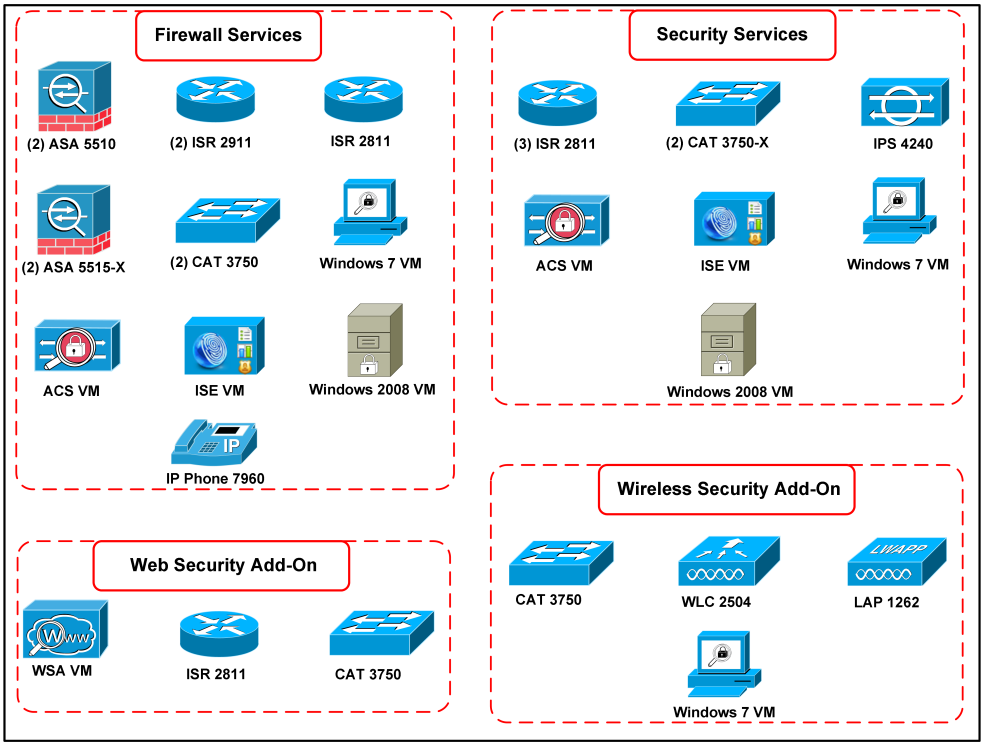
Configurations

Initial configuration files are provided at the top of this page. Each lab has an initial configuration that is assumed. The initial configuration in some cases is the configuration that the previous lab ended with. If there is only one configuration file, then the other devices did not change in the previous lab.

Section 1 General Lab Topology

The INE CCIE Security Racks are divided into modules, as seen in the figure below. This section, Section 1, **utilizes Firewall Services.**

https://t.me/learningnets



Use the Section 1 Topology Diagram for this section. Additional diagrams for each task are found in the **Task Information** area. The topology may change for various lab topics, so pay close attention to the diagrams provided.

Using Your Own Equipment

If you use your own rack to practice these labs, we recommend that you use the following hardware/software versions to obtain the same results and outputs, because these are inline with the CCIE Security Lab Equipment and Software v4.0 list provided by Cisco:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124
- ACS is running version 5.3

https://t.me/learningnets

- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS RIPv2 Authentication

Task

- Configure RIPv2 on R1, R2, and R3.
- Advertise the routers' Ethernet links and Loopback networks into RIP.
- Authenticate the RIPv2 exchange between R1 and R3 using the clear-text string **CLEARTEXT**.
- Authenticate the RIPv2 exchange between R2 and R3 using an MD5 hash of the string **MD5HASH**.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** or by clicking the **Resources** button.

Configuration

```
R1:
key chain TEXT
 key 1
  key-string CLEARTEXT
!
interface GigabitEthernet0/0
 ip rip authentication mode text
 ip rip authentication key-chain TEXT
!
router rip
 version 2
 network 136.1.0.0
 network 150.1.0.0
 no auto-summary
R2:
key chain MD5
 key 1
```

https://t.me/learningnets

```
  key-string MD5HASH
!
interface GigabitEthernet0/0
 ip rip authentication mode md5
 ip rip authentication key-chain MD5
!
router rip
 version 2
 network 136.1.0.0
 network 150.1.0.0
 no auto-summary
```

R3:

```
key chain TEXT
 key 1
  key-string CLEARTEXT
!
key chain MD5
 key 1
  key-string MD5HASH
!
interface FastEthernet0/0.13
 ip rip authentication mode text
 ip rip authentication key-chain TEXT
!
interface FastEthernet0/0.23
 ip rip authentication mode md5
 ip rip authentication key-chain MD5
!
router rip
 version 2
 network 136.1.0.0
 network 150.1.0.0
 no auto-summary
```

Verification

RIP authentication is configured in three distinct steps:

1. Define a global key chain.
2. Enable authentication mode (clear-text or MD5) on the RIP interfaces.
3. Apply the key chain to the interface.

https://t.me/learningnets

This configuration is nearly identical to both EIGRP and IS-IS authentication, with minor syntax differences at the interface level of how the authentication is enabled. Note that for MD5-based authentication, the key number in the key chain must match between the neighbors, because this value is used as a salt for the MD5 hash.

Verification that RIP authentication is working is very straightforward, whereas if the routes from the RIP neighbor are installed in the routing table, authentication was successful, as seen below:

```
R3#show ip route rip

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 2 subnets
R        150.1.1.1 [120/1] via 136.1.13.1, 00:00:08, FastEthernet0/0.13
R        150.1.2.2 [120/1] via 136.1.23.2, 00:00:13, FastEthernet0/0.23
```

A process-level debug with `debug ip rip` can also verify that authentication is working, as seen below.

```
R3#debug ip rip
RIP protocol debugging is on
R3#
RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0.23 (136.1.23.3)
RIP: build update entries
        136.1.13.0/24 via 0.0.0.0, metric 1, tag 0
        150.1.1.1/32 via 0.0.0.0, metric 2, tag 0RIP: received packet with MD5 authentication
RIP: received v2 update from 136.1.23.2 on FastEthernet0/0.23
     150.1.2.2/32 via 0.0.0.0 in 1 hopsRIP: received packet with text authentication CLEARTEXT

RIP: received v2 update from 136.1.13.1 on FastEthernet0/0.13
     150.1.1.1/32 via 0.0.0.0 in 1 hops
RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0.13 (136.1.13.3)
RIP: build update entries
        136.1.23.0/24 via 0.0.0.0, metric 1, tag 0
```

https://t.me/learningnets

```
        150.1.2.2/32 via 0.0.0.0, metric 2, tag 0
```

This debug output will likewise show if there is a failure in authentication, similar to the below output.

```
R3#debug ip rip
RIP protocol debugging is on RIP: ignored v2 packet from 136.1.13.1 (invalid authentication)

RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0.13 (136.1.13.3)
RIP: build update entries
        136.1.23.0/24 via 0.0.0.0, metric 1, tag 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Message Logging

Task

- Configure logging on R1:
  - Save debugging messages in the internal buffer's size of 16384 bytes.
  - Set debugging level for console and informational for VTY lines.
  - Log messages should not interrupt any command output.
  - Disable log messages of severity 7.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

Logging messages can be sent to multiple destinations simultaneously, such as console/VTY lines, internal buffer, remote syslog server. Some platforms also allow storing of logging messages in the local flash memory. For each destination, you can configure the message severity level, which controls which messages are logged. Logging at severity 7 (debugging) includes all messages 0 through 7. Logging at severity 4 includes all messages 0 through 4.

For syslog, the logging facility is also available, which controls the format of the log message. Using different facilities for different devices, for example local5 for all routers and local6 for all switches, can make sorting through log files on the syslog server itself more manageable. Synchronous logging on a terminal line makes the logging process wait for the terminal to finish printing a line before outputting its own

https://t.me/learningnets

message. Logging rate-limiting shrinks the number of messages (actually, message lines) sent per second, either to all destinations or to the console line specifically.

Because debug messages can negatively affect performance and result in system instability (or even worse, cause the router to crash/reload because of high CPU load), it is possible to disable logging messages of severity 7 using the command **`no logging buginf .`**

Configuration

```
R1:

logging on
logging buffered 16384 debugging
logging console debugging
logging rate-limit console all 10
logging monitor informational
no logging buginf
!
line console 0
 logging synchronous
!
line vty 0 14
 logging synchronous
```

Verification

Verify logging configuration:

```
R1#show logging
Syslog logging: enabled (0 messages dropped, 2 messages rate-limited, 0 flushes, 0 overruns, xml disabled, filtering

No Active Message Discriminator.



No Inactive Message Discriminator.

 Console logging: level debugging
, 48 messages logged, xml disabled,
                     filtering disabled Monitor logging: level informational
, 0 messages logged, xml disabled,
                     filtering disabled
```

https://t.me/learningnets

```
Buffer logging:  level debugging
, 9 messages logged, xml disabled,
                    filtering disabled
    Exception Logging: size (4096 bytes)
    Count and timestamp logging messages: disabled
    Persistent logging: disabled

No active filter modules.

    Trap logging: level informational, 50 message lines logged
        Logging Source-Interface:       VRF Name:

Log Buffer (16384 bytes):
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Banners

Task

- Display the following banners on R1:
  - All connecting users should see the message **This is the CCIE Security Lab Exam**.
  - Before giving the shell prompt, display the message **You are connected to [hostname] on line [line]**; replace *hostname* with the router's name and *line* with the line number on which the user is connected; do not show the message for console sessions.
  - Before issuing the login prompt, display the message **Do you know your username and password?**.
  - All reverse telnet sessions should see the message **What device do you want to connect on?**.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

For remote access and local console connections, the following types of banners can be displayed:

- MOTD banner, displayed first to everyone connecting to the router.
- Login banner, displayed before the authentication prompt is presented to the user (if

https://t.me/learningnets

login is disabled, the banner is not displayed), but after MOTD banner.

- Exec banner, displayed before the exec prompt is displayed to the user, so after the login banner.
- Incoming banner, displayed to users connecting via reverse telnet; it replaces the exec banner.

Special tokens can be used in the banner messages, such as $(line), $(line-desc), $(hostname), and $(domain), to create dynamic notifications. Some lines can be selectively exempted from viewing banner messages, such as the console in this example. Issuing `no exec-banner` will disable both MOTD and exec banners, whereas `no motd banner` will disable only the MOTD banner.

Configuration

R1:

```
banner motd *
This is the CCIE Security Lab Exam
*
!
banner login *
Do you know your username and password?
*
!
banner exec *
You are connected to $(hostname) on line  $(line)
*
!
banner incoming *
What device you want to connect on?
*
!
line console 0
 no exec-banner
```

Verification

https://t.me/learningnets

Because console and VTY lines are in the **no-login** mode, console by default and VTY lines from the loaded configurations, login banner will never be displayed. Verify that no banners are displayed at console; because the exec banner is removed, the MOTD banner is also:

```
R1#exit

R1>enable
Password: R1#
```

On a telnet session, the MOTD and exec banners should be displayed:

```
R1#telnet 150.1.1.1
Trying 150.1.1.1 ... Open
 This is the CCIE Security Lab Exam
 You are connected to R1 on line  388

R1>enable
Password:
R1#
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS Network Time Protocol

Task

- Configure NTP version 4 as follows:
  - R3 is the NTP server.
  - R1 synchronizes its clock with R3 via broadcast on 255.255.255.255.
  - R2 synchronizes its clock with R3 via multicast on 239.10.10.10.
  - All NTP messages require authentication.
  - R1 and R2 should trigger a message log when the clock is synchronized.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

NTP time distribution is based on a loopless tree topology. In the root of the tree resides the absolute time source (server), which is physically connected to an atomic clock, radio clock, or some other highly accurate time source. This server is described as being in stratum 1, meaning that it is one hop away from the source of the time. Cisco routers configured as NTP servers cannot work at stratum level 1.

Every next router that pulls time down from stratum N is automatically assigned to stratum N+1. That router in turn may become a time source (server) for clients in stratum N+2. A router in stratum N will not accept time updates from stratum N+1, enforcing a tree-level hierarchy of client-server relationships and preventing time synchronization loops. If a router has multiple servers, it will only select one for time

https://t.me/learningnets

synchronization (the one that is closer to NTP root and has the least offset from the currently running clock) and use others as backup. The preference may be adjusted using the `prefer` keyword when defining the NTP server, but only among candidates having other parameters equal (such as the same stratum).

Two routers in the same stratum may also form peer-to-peer, symmetric relationships. In this mode, both peers update each other's clocks. However, just one router in a pair may be configured for NTP peering (active peer); the other router will automatically convert into passive mode and will provide time updates to the dynamic peer.

In most cases, NTP is deployed in unicast mode, where the server is manually declared on clients. NTP is also capable of broadcasting or multicasting time updates on a shared interface, and clients may listen to NTP packets on the other side. This allows a single server to synchronize multiple hosts in unsolicited mode. Remember that NTP may take considerable time to synchronize, especially if router clocks are far out of sync. To speed up this process, adjust the clocks manually to values close to the server, to accelerate convergence (when the router considers all servers as insane). In the NTP show outputs of the NTP server, a reference of 127.127.1.1 means that local clock is used; it is not received from an upstream server.

NTP clients can be configured to authenticate the time source (NTP server) so that changes to the local clock is authenticated/authorized. This is achieved by using an HMAC MD5 signature, carrying also the signing key index number. The authentication key number and key string must be identically configured on both the server and the client; additionally, both server and client must have the key marked as trusted; otherwise, the server will not send any NTP packets with the respective key and clients will reject NTP messages. Only on the client side must NTP authentication be enabled.

The most important function that NTP version 4 adds is IPv6 support.

Configuration

```
R1:
ntp authenticate
ntp authentication-key 1 md5 cisco
ntp trusted-key 1
ntp logging
!
interface GigabitEthernet0/0
 ntp broadcast client
```

https://t.me/learningnets

```
R2:
ntp authenticate
ntp authentication-key 1 md5 cisco
ntp trusted-key 1
ntp logging
!
interface GigabitEthernet0/0
 ntp multicast client 239.10.10.10


R3:
ntp master 6
ntp authentication-key 1 md5 cisco
ntp trusted-key 1
ntp logging
!
interface FastEthernet0/0.13
 ntp broadcast version 4 key 1
!
interface FastEthernet0/0.23
 ntp multicast 239.10.10.10 version 4 key 1
```

Verification

Because of NTP logging, the following messages will be logged on R1 and R2, when their clocks are synchronized with R3:

```
%SYS-6-CLOCKUPDATE: System clock has been updated from 18:22:12 UTC Sun Dec 1 2013 to 19:30:36 UTC Sun Dec 1 2013, c
%SYS-6-CLOCKUPDATE: System clock has been updated from 18:21:55 UTC Sun Dec 1 2013 to 19:30:36 UTC Sun Dec 1 2013, c
```

Verify that R3 is synchronized with itself; otherwise, clients will not be able to synchronize:

https://t.me/learningnets

```
R3#show ntp status
Clock is synchronized, stratum 6, reference is 127.127.1.1

nominal freq is 250.0000 Hz, actual freq is 250.0015 Hz, precision is 2**24
reference time is D6460F18.43D0EAAF (19:40:08.264 UTC Sun Dec 1 2013)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.33 msec, peer dispersion is 0.24 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is -0.000005946 s/s
system poll interval is 16, last update was 6 sec ago.
```

Verify that R1 has its clock synchronized with R3 and that time source is authenticated:

https://t.me/learningnets

```
R1#show ntp status
Clock is synchronized, stratum 7, reference is 136.1.13.3

nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**20
reference time is D6460F5B.43FA5BC9 (19:41:15.265 UTC Sun Dec 1 2013)
clock offset is -0.1661 msec, root delay is 1.79 msec
root dispersion is 2.29 msec, peer dispersion is 0.92 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is -0.000000007 s/s
system poll interval is 64, last update was 62 sec ago.
!
!R1#show ntp associations


  address         ref clock       st   when   poll reach  delay  offset   disp
* 136.1.13.3      127.127.1.1      6      8     64   377  1.791  -0.158  0.940

 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
!
!R1#show ntp associations detail
136.1.13.3 dynamic, authenticated, our_master, sane, valid, stratum 6

ref ID 127.127.1.1    , time D6460F99.439D6A57 (19:42:17.264 UTC Sun Dec 1 2013)
our mode xcast client, peer mode xcast, our poll intvl 64, peer poll intvl 64

root delay 0.00 msec, root disp 0.27, reach 376, sync dist 2.42
delay 1.79 msec, offset -0.1585 msec, dispersion 0.94
precision 2**24, version 4
org time 00000000.00000000 (00:00:00.000 UTC Mon Jan 1 1900)
rec time D6460D1D.2C6AAFFF (19:31:41.173 UTC Sun Dec 1 2013)
xmt time D6460F9D.439A4FF5 (19:42:21.264 UTC Sun Dec 1 2013)
filtdelay =     1.79    1.79    1.79    1.79    1.79    1.79    1.79    1.79
filtoffset =   -0.15   -0.16   -0.15   -0.15   -0.15   -0.16   -0.19   -0.21
filterror =     0.00    0.99    1.93    2.91    3.87    4.84    5.80    6.75
minpoll = 6, maxpoll = 6
```

Verify that R2 has its clock synchronized with R3 and that time source is authenticated:

```
R2#show ntp status
Clock is synchronized, stratum 7, reference is 136.1.23.3
nominal freq is 250.0000 Hz, actual freq is 249.9998 Hz, precision is 2**21
reference time is D6461028.43B04584 (19:44:40.264 UTC Sun Dec 1 2013)
clock offset is -0.2392 msec, root delay is 1.87 msec
root dispersion is 2.15 msec, peer dispersion is 0.93 msec
```

https://t.me/learningnets

```
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000792 s/s
system poll interval is 64, last update was 33 sec ago.
!
!R2#show ntp associations

  address         ref clock       st   when   poll reach  delay  offset   disp
* 136.1.23.3      127.127.1.1      6     41     64   376  1.878  -0.239  0.933
 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
!
!R2#show ntp associations detail
136.1.23.3 dynamic, authenticated, our_master, sane, valid, stratum 6
ref ID 127.127.1.1    , time D646101B.4365D91E (19:44:27.263 UTC Sun Dec 1 2013)
our mode xcast client, peer mode xcast, our poll intvl 64, peer poll intvl 64

root delay 0.00 msec, root disp 0.44, reach 376, sync dist 3.12
delay 1.87 msec, offset -0.2392 msec, dispersion 0.93
precision 2**24, version 4
org time 00000000.00000000 (00:00:00.000 UTC Mon Jan 1 1900)
rec time D6460D1D.5F852C08 (19:31:41.373 UTC Sun Dec 1 2013)
xmt time D6461028.436309AA (19:44:40.263 UTC Sun Dec 1 2013)
filtdelay =     1.87    1.87    1.87    1.87    1.87    1.87    1.87    1.87
filtoffset =   -0.23   -0.24   -0.23   -0.25   -0.28   -0.24   -0.27   -0.28
filterror =     0.00    0.96    1.93    2.91    3.87    4.86    5.85    6.84
minpoll = 6, maxpoll = 6
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS NTP Access Control

Task

- Configure R2 and R3 to synchronize time with each other; no other device can update their clocks.
- Configure R1 as NTP client of R3.
- R3 requires strict control of its clients; allow only R1 to update its clock.
- All NTP messages require authentication and are sourced from Loopback0 interfaces.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

NTP access-control divides NTP messages in two categories, control messages and update/request messages. Control messages are those needed to extract specific management information, such as the peer status or set a management parameter. NTP control messages are not needed for proper time synchronization. NTP update/request messages are those messages needed for time synchronization.

NTP access-control defines four levels:

- Peer: permits NTP updates/requests to the host as well as control queries; this is the only access type that allows the host to be synchronized by others.
- serve: permits NTP requests, but rejects NTP updates (does not change local clock);

https://t.me/learningnets

control queries are also permitted.

- Serve-only: permits NTP requests only; it rejects attempts to synchronize the local system and does not accept control queries.
- Query-only: accepts NTP control queries; no responses to NTP requests are sent, and no local system time synchronization with a remote system is allowed.

When these levels are associated with access-lists on a router, incoming NTP message source IP address is matched in the order listed above (peer, serve, serve-only, and query-only). The first match determines the type of access. If some access-groups are configured but not all, all other types of access from any other sources are implicitly denied.

Configuration

```
R1:
ntp authenticate
ntp authentication-key 1 md5 R1
ntp trusted-key 1
ntp source Loopback0
!
ntp server 150.1.3.3 key 1
!
ip route 150.1.3.3 255.255.255.255 136.1.13.3


R2:
ntp master 6
ntp authenticate
ntp authentication-key 2 md5 R2
ntp trusted-key 2
ntp source Loopback0
!
ntp peer 150.1.3.3 key 2
!
access-list 2 permit host 150.1.3.3
ntp access-group ipv4 peer 2
!
ip route 150.1.3.3 255.255.255.255 136.1.23.3


R3:
ntp master 6
ntp authenticate
```

https://t.me/learningnets

```
ntp authentication-key 1 md5 R1
ntp authentication-key 2 md5 R2
ntp trusted-key 1
ntp trusted-key 2
ntp source Loopback0
!
ntp peer 150.1.2.2 key 2
!
access-list 2 permit host 150.1.2.2
ntp access-group peer 2
!
access-list 1 permit host 150.1.1.1
ntp access-group serve-only 1
!
ip route 150.1.1.1 255.255.255.255 136.1.13.1
ip route 150.1.2.2 255.255.255.255 136.1.23.2
```

Verification

Verify that R1 synchronized time with R3 and authenticated the time source:

https://t.me/learningnets

```
R1#show ntp status
Clock is synchronized, stratum 7, reference is 150.1.3.3

nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**20
reference time is D64624B7.2C606C9A (21:12:23.173 UTC Sun Dec 1 2013)
clock offset is 0.2767 msec, root delay is 1.53 msec
root dispersion is 188.33 msec, peer dispersion is 187.54 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000020 s/s
system poll interval is 128, last update was 5 sec ago.
!
!R1#show ntp associations detail
150.1.3.3 configured, authenticated, our_master, sane, valid, stratum 6

ref ID 127.127.1.1    , time D64624A8.3B4BF51A (21:12:08.231 UTC Sun Dec 1 2013)
our mode client, peer mode server, our poll intvl 128, peer poll intvl 128

root delay 0.00 msec, root disp 0.44, reach 1, sync dist 189.42
delay 1.53 msec, offset 0.2767 msec, dispersion 187.54
precision 2**24, version 4
org time 00000000.00000000 (00:00:00.000 UTC Mon Jan 1 1900)
rec time D64624B7.2C402DA3 (21:12:23.172 UTC Sun Dec 1 2013)
xmt time D64624B7.2C402DA3 (21:12:23.172 UTC Sun Dec 1 2013)
filtdelay =     1.53    1.66    1.55    1.59    1.59    1.58    0.00    0.00
filtoffset =    0.27    0.27    0.28    0.26    0.29    0.28    0.00    0.00
filterror =     0.00    0.03    0.06    0.09    0.12    0.15 16000.0 16000.0
minpoll = 6, maxpoll = 10
```

Verify NTP settings on R2 and R3:

```
R2#show ntp status
Clock is synchronized, stratum 6, reference is 127.127.1.1
nominal freq is 250.0000 Hz, actual freq is 249.9998 Hz, precision is 2**21
reference time is D6462F78.609EC2AD (21:58:16.377 UTC Sun Dec 1 2013)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.24 msec, peer dispersion is 0.23 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000774 s/s
system poll interval is 16, last update was 1 sec ago.
!
!R2#show ntp associations

  address         ref clock       st   when   poll reach  delay  offset   disp
*~127.127.1.1     .LOCL.           5     12     16   377  0.000   0.000  0.232
 ~150.1.3.3       127.127.1.1      6     53     64   376  1.912   0.101  2.635
```

https://t.me/learningnets

```
 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
!
!R3#show ntp status
Clock is synchronized, stratum 6, reference is 127.127.1.1
nominal freq is 250.0000 Hz, actual freq is 250.0015 Hz, precision is 2**24
reference time is D6462F8D.37055A38 (21:58:37.214 UTC Sun Dec 1 2013)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.27 msec, peer dispersion is 0.22 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is -0.000005946 s/s
system poll interval is 16, last update was 3 sec ago.
!
!R3#show ntp associations

  address         ref clock       st   when   poll reach  delay  offset   disp
*~127.127.1.1     .LOCL.           5     11     16   377  0.000   0.000  0.225
 ~150.1.2.2       127.127.1.1      6     30     64   376  1.930  -0.106  2.370

 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS SNMPv2

Task

- Configure R1 with an snmp community value of **CCIELAB**, allowing changes to MIB database from hosts in 150.1.0.0/16 subnet.
- Ensure that R1 can be reloaded through SNMP and that interface index survives a reload.
- Configure SNMP on R1, with location set to **Romania Bucharest** and contact set to **CCIE Trainer**.
- Configuration is allowed to be retrieved via FTP/TFTP servers in 150.1.0.0/16 subnet.
- Notifications about link states should be sent:
  - v2c as informs to host 136.1.13.30.
  - v2c as traps to host 136.1.13.40.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

By default, when you configure the SNMP agent with a community string, SNMPv2c is used, which uses only the community string for authentication. Issuing the global command `snmp-server community <value>` is the minimal configuration needed to enable the devices to be polled via SNMP. The read-write string allows the management station to make changes to the managed device, as opposed to the read-only string. Because the community is visible in clear-text in transport and read-

https://t.me/learningnets

write exposes the managed devices to certain risks, it may be desirable to restrict from which sources can it be managed; this is achieved by attaching an ACL to the community definition.

SNMP interface index (IfIndex) values are not preserved between reloads by default. This may confuse some network management systems (NMS), because interfaces are identified by the IfIndex value in the SNMP Management Information Base (MIB). To fix this, enable index persistency using the command `snmp-server ifindex persist` . To allow the NMS to reload the managed device, read-write access must be allowed, and the `snmp-server system-shutdown` feature must be enabled. The file-server list specifies the acceptable addresses to download/upload the router's configuration when instructed via SNMP, as well as protocols. The old command format that allowed transfers only via TFTP is `snmp-server tftp-server-list <acl> .`

SNMPv2c allows sending notifications as informs, which differ from traps in that they require acknowledgement from the NMS (network management server). Informs are kept in router local queue until they are acknowledged or timeout has expired. Informs make SNMP reliable even though the transport protocol is still UDP. The traps that are to be sent can be configured either globally or on a per-host basis. The host settings override the global settings, but some traps can only be enabled globally, such as link up and down notifications. If the command `snmp-server enable traps` is entered, all traps will be sent to all configured destinations, unless they are overridden on a per-host basis. Note that it is required to configure both the `snmp-server enable traps` and the `snmp-server host <IP_address>` commands to send notifications to a particular host. The same host may receive informs and traps in parallel, although this configuration is uncommon.

Configuration

https://t.me/learningnets

```
R1:


snmp-server community CCIELAB RW 1
snmp-server location Romania Bucharest
snmp-server contact CCIE Trainer
snmp-server ifindex persist
snmp-server system-shutdown
!
access-list 1 permit 155.1.0.0 0.0.255.255
snmp-server file-transfer access-group 1 protocol ftp tftp
!
snmp-server enable traps snmp linkup linkdown
snmp-server host 136.1.13.30 inform version 2c cisco
snmp-server host 136.1.13.40 traps version 2c cisco
```

Verification

Verify SNMP configuration:

```
R1#show snmp
Chassis: FCZ16436080Contact: CCIE Trainer
Location: Romania Bucharest
0 SNMP packets input
    0 Bad SNMP version errors
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    0 Number of requested variables
    0 Number of altered variables
    0 Get-request PDUs
    0 Get-next PDUs
    0 Set-request PDUs
    0 Input queue packet drops (Maximum queue size 1000)
0 SNMP packets output
    0 Too big errors (Maximum packet size 1500)
    0 No such name errors
    0 Bad values errors
    0 General errors
    0 Response PDUs
    0 Trap PDUs
SNMP Dispatcher:
   queue 0/75 (current/max), 0 dropped
SNMP Engine:
```

https://t.me/learningnets

```
   queue 0/1000 (current/max), 0 dropped
SNMP logging: enabled
Logging to 136.1.13.40.162, 0/10, 0 sent, 0 dropped.

SNMP Manager-role output packets
    0 Get-request PDUs
    0 Get-next PDUs
    0 Get-bulk PDUs
    0 Set-request PDUs
    0 Inform-request PDUs
    0 Timeouts
    0 Drops
SNMP Manager-role input packets
    0 Inform request PDUs
    0 Trap PDUs
    0 Response PDUs
    0 Responses with errors
SNMP informs: enabled
    Informs in flight 0/25 (current/max) Logging to 136.1.13.30.162

        0 sent, 0 in-flight, 0 retries, 0 failed, 0 dropped
```

Verify SNMP communities and hosts:

```
R1#show snmp community

Community name: ILMI
Community Index: cisco0
Community SecurityName: ILMI
storage-type: read-only  active

Community name: CCIELAB
Community Index: cisco5
Community SecurityName: CCIELAB storage-type: nonvolatile        active access-list: 1
!
!R1#show snmp host
Notification host: 136.1.13.30  udp-port: 162   type: inform
user: cisco     security model: v2c
Notification host: 136.1.13.40  udp-port: 162   type: trap

user: cisco     security model: v2c
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS SNMPv3

Task

- Configure two SNMP views on R2:
  - View **AUTH_VIEW** should include branch **system**.
  - View **SECURE_VIEW** should include branch **iso**.
- Configure an SNMP group named **AUTH_GROUP** with read access to view **AUTH_VIEW** and a security model of **auth**.
  - Assign a user named **AUTH_USER** to the **AUTH_GROUP**; use sha for authentication.
  - Restrict access to this group for hosts in subnet 150.1.0.0/16.
- Configure an SNMP group named **SECURE_GROUP** with red/write access to view **SECURE_VIEW** and a security model of **priv**.
  - Assign a user named **SECURE_USER** to the **SECURE_GROUP**; use md5 for authentication and 3des for encryption.
  - Enable SNMP traps to host 136.1.23.30 using the user **SECURE_USER**.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

https://t.me/learningnets

SNMPv3 introduces a new security model that replaces the old community-based authentication system and provides for communication privacy by means of encryption. The new concepts for SNMPv3 are the user, group, and security level.

A group defines the access rights that a set of users has. This access policy controls which SNMP objects (MIBs) can be accessed for reading and writing, or which SNMP objects can generate notifications to the members of a group. The policy is defined by associating a read, write, or notify view with the group. By using a notify view, a group determines the list of notifications its users can receive. The group also defines the security model (SNMP version) and the security level (authentication and/or encryption) for its users.

If a group is defined without a read view, all objects are available to be read (implicit permit). Contrary to that, if a write or notify view is not defined, no write access is granted, and no objects can send notifications to members of the group (implicit deny). The notify view is usually not configured manually and is auto-generated by the `snmp-server host` command when users in a group are bound to a notification target host.

The security models are defined as SNMPv1, SNMPv2, SNMPv3, and the security levels are defined as:

- noAuthNoPriv: `noauth` keyword in the IOS means no authentication and no encryption.
- AuthNoPriv: `auth` keyword in the IOS means authentication but no encryption.
- AuthPriv: `priv` keyword in IOS means both authentication and encryption.

SNMPv3 can implement any of the three security levels listed above. SNMPv1 and SNMPv2 only support noAuthNoPriv. If SNMPv3 uses noAuthNoPriv, the username serves as a replacement for the community string. All users sharing a group use the same security model, but the specific model settings (password and encryption key) are set per user. Note that SNMPv3 does not send passwords in clear-text, but instead uses MD5 or SHA1 hash-based authentication. For encryption, statically configured keys are used along with a single-DES (56-bit) symmetric cipher. This means that the same key should be configured on NMS for the particular user.

> Note that SNMPv3 users do not appear in the running configuration for security reasons.

Configuration

```
R2:
```

https://t.me/learningnets

```
access-list 1 permit 155.1.0.0 0.0.255.255
!
snmp-server view AUTH_VIEW system included
snmp-server view SECURE_VIEW iso included
!
snmp-server group AUTH_GROUP v3 auth read AUTH_VIEW access 1
snmp-server user AUTH_USER AUTH_GROUP v3 auth sha cisco
!
snmp-server group SECURE_GROUP v3 priv read SECURE_VIEW write SECURE_VIEW
snmp-server user SECURE_USER SECURE_GROUP v3 auth md5 cisco priv 3des cisco
!
snmp-server enable traps
snmp-server host 136.1.23.30 traps version 3 priv SECURE_USER
```

Verification

Verify the groups and users:

```
R2#show snmp group
groupname: ILMI                              security model:v1
contextname: <no context specified>          storage-type: permanent
readview : *ilmi                             writeview: *ilmi
notifyview: <no notifyview specified>
row status: active

groupname: ILMI                              security model:v2c
contextname: <no context specified>          storage-type: permanent
readview : *ilmi                             writeview: *ilmi
notifyview: <no notifyview specified>
row status: active
groupname: AUTH_GROUP                         security model:v3 auth
contextname: <no context specified>          storage-type: nonvolatile
readview : AUTH_VIEW                         writeview: <no writeview specified>
notifyview: <no notifyview specified>        row status: active      access-list: 1
groupname: SECURE_GROUP                       security model:v3 priv
contextname: <no context specified>          storage-type: nonvolatile
readview : SECURE_VIEW                       writeview: SECURE_VIEW
notifyview: *tv.FFFFFFFF.FFFFFFFF.FFFFFFFF.F
row status: active
!
!R2#show snmp user
 User name: AUTH_USER
Engine ID: 800000090300000000000000
storage-type: nonvolatile        active Authentication Protocol: SHA
Privacy Protocol: None
Group-name: AUTH_GROUP
```

https://t.me/learningnets

```
 User name: SECURE_USER

Engine ID: 800000090300000000000000
storage-type: nonvolatile        active Authentication Protocol: MD5
 Privacy Protocol: 3DES
 Group-name: SECURE_GROUP
```

Verify SNMP trap configuration:

```
R2#show snmp host
 Notification host: 136.1.23.30  udp-port: 162   type: trap
 user: SECURE_USER        security model: v3 priv
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS RTBH Destination Based

Task

- Configure RIPv2 on R1, R2, and R3 so that you achieve full connectivity between all IPv4 subnets.
- Several DoS attacks have been generated from behind R1, targeting **TEST-PC-A** at 10.1.0.100:
  - Configure BGP on R2 and R3 using AS 23.
  - Configure RTBH to silently stop the attack; consider R2 the trigger router and R3 the edge router.
  - Static routes are allowed.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

Imagine that your network is under attack and you are able to somehow identify the source of the attack, the destinations of the attack, or both. A simple but not very scalable solution would be to configure access-list filtering for the identified sources/destinations and apply it inbound on all your ingress points in the network. A more efficient method is to use the BGP RTBH (Remotely Triggered Black Hole) functionality, which was specifically designed for such cases and even available as RFC5635. How this works is very simple:

- All your edge routers (connecting your network to other networks) will run iBGP with

https://t.me/learningnets

one or two special routers within your network called trigger routers.
- When the source/destination of the attack is identified, you configure the trigger routers with a route to Null0 (recycle bin) for it.
- Trigger routers will send an iBGP update for the identified sources/destinations to the edge routers, setting a specific next-hop value.
- All edge/trigger routers have a pre-defined static route to Null0 for the specific next-hop value.

Based on the above, the end result is that edge routers will be updated dynamically (on-demand) by the trigger routers, and all received traffic for the respective destinations or from respective sources will be routed to Null0, effectively stopping the attack at the edge of the network. This example is for a destination-based scenario.

Although the main scope of the trigger router is to set the specific next-hop value (which you choose, but it needs to be unroutable and nonexistent within your network), note there are several other parameters configured in the route-map used for redistribution; their scope is to ensure that when edge routers run the BGP best-path selection algorithm, the route advertised by trigger routers is selected as best.

Normally, when a router drops a packet it will send back an ICMP message to the originator, signaling the reason for packet being dropped; in our case, the reason would be routing it to Null0. To silently drop packets (and save CPU resources on edge routers, because ICMP messages as any other router-generated packets are not CEF switched, but process switched), disable sending of ICMP unreachables for the Null0 interface.

Configuration

```
R1:
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0

R2:
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
!
```

https://t.me/learningnets

```
router bgp 23
 neighbor 136.1.23.3 remote-as 23
 neighbor 136.1.23.3 send-community
 redistribute static route-map STATIC_TO_BGP
!
route-map STATIC_TO_BGP permit 10
 match tag 23
 set local-preference 200
 set origin igp
 set community no-export
 set ip next-hop 192.0.2.1
!
ip route 192.0.2.1 255.255.255.255 Null0
ip route 10.1.0.100 255.255.255.255 Null0 tag 23
```

R3:

```
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
 network 10.0.0.0
!
router bgp 23
 neighbor 136.1.23.2 remote-as 23
!
ip route 192.0.2.1 255.255.255.255 Null0
!
interface Null0
 no ip unreachables
```

Verification

Verify that R3 received the BGP update from R2, and check the attributes attached to the prefix:

```
R3#show ip bgp 10.1.0.100
BGP routing table entry for 10.1.0.100/32, version 2
Paths: (1 available, best #1, table default, not advertised to EBGP peer)
  Not advertised to any peer
  Local
    192.0.2.1 from 136.1.23.2 (150.1.2.2) Origin IGP, metric 0, localpref 200, valid, internal, best
```

https://t.me/learningnets

```
Community: no-export
```

Verify through route-recursive lookups that packets for 10.1.0.100 destination will be routed to Null0 by R3:

```
R3#show ip route 10.1.0.100
Routing entry for 10.1.0.100/32
  Known via "bgp 23", distance 200, metric 0, type internal
  Last update from 192.0.2.1 00:03:05 ago
  Routing Descriptor Blocks: * 192.0.2.1, from 136.1.23.2, 00:03:05 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none
!
!R3#show ip route 192.0.2.1
Routing entry for 192.0.2.1/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks: * directly connected, via Null0
      Route metric is 0, traffic share count is 1
!
!R3#show ip cef 10.1.0.100
10.1.0.100/32 nexthop 192.0.2.1 Null0
```

Verify that packets toward 10.1.0.100 are silently dropped by R3 and no ICMP unreachables are being sent back:

```
R1#ping 10.1.0.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.0.3, timeout is 2 seconds: !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R1#ping 10.1.0.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.0.100, timeout is 2 seconds: .....
Success rate is 0 percent (0/5)
!
!R3#show interfaces null 0 | section 5 minute
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
```

https://t.me/learningnets

```
5 packets output, 590 bytes, 0 underruns

     0 output errors, 0 collisions, 0 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 1: System Hardening and Availability

IOS RTBH Source Based

Task

- Configure RIPv2 on R1, R2, and R3 so that you achieve full connectivity between all IPv4 subnets.
- Several DoS attacks have been generated from behind R1, sourced from 170.170.170.170:
  - Configure BGP on R2 and R3 using AS 23.
  - Configure RTBH to stop the attack; consider R2 the trigger router and R3 the edge router.
  - Static routes are allowed.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 1 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

This is mostly the same scenario as destination-based RTBH, but this time the source of the attack is known. We need to somehow drop the attack based on the source of traffic. This is achieved by enabling uRPF in loose mode on the interface through which traffic is entering your network and configuring the trigger BGP router to send an update with next-hop pointing to Null0 for the respective source. uRPF in loose mode will drop the traffic only if a route lookup for the source points to Null0, which is our case.

https://t.me/learningnets

Configuration

R1:

```
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
```

R2:

```
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
!
router bgp 23
 neighbor 150.1.3.3 remote-as 23
 neighbor 150.1.3.3 update-source Loopback0
 neighbor 150.1.3.3 send-community
 redistribute static route-map STATIC_TO_BGP
!
route-map STATIC_TO_BGP permit 10
 match tag 23
 set local-preference 200
 set origin igp
 set community no-export
 set ip next-hop 192.0.2.1
!
ip route 192.0.2.1 255.255.255.255 Null0
ip route 170.170.170.170 255.255.255.255 Null0 tag 23
```

R3:

```
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
 network 10.0.0.0
!
router bgp 23
```

https://t.me/learningnets

```
 neighbor 150.1.2.2 remote-as 23
 neighbor 150.1.2.2 update-source Loopback0
!
ip route 192.0.2.1 255.255.255.255 Null0
!
interface FastEthernet0/0.13
 ip verify unicast source reachable-via any
```

Verification

Verify that R3 received the BGP update from R2, and check the attributes attached to the prefix:

```
R3#show ip bgp 170.170.170.170
BGP routing table entry for 170.170.170.170/32, version 4
Paths: (1 available, best #1, table default, not advertised to EBGP peer)
  Not advertised to any peer
  Local
    192.0.2.1 from 150.1.2.2 (150.1.2.2) Origin IGP, metric 0, localpref 200, valid, internal, best
Community: no-export
```

Verify through route-recursive lookups that packets for 170.170.170.170 destination will be routed to Null0 by R3:

https://t.me/learningnets

```
R3#show ip route 170.170.170.170

Routing entry for 170.170.170.170/32
  Known via "bgp 23", distance 200, metric 0, type internal
  Last update from 192.0.2.1 00:02:22 ago
  Routing Descriptor Blocks: * 192.0.2.1, from 150.1.2.2, 00:02:22 ago

      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none
!
!R3#show ip route 192.0.2.1

Routing entry for 192.0.2.1/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks: * directly connected, via Null0

      Route metric is 0, traffic share count is 1
!
!R3#show ip cef 170.170.170.170

170.170.170.170/32 nexthop 192.0.2.1 Null0
```

Verify that packets sourced from 170.170.170.170 are dropped by R3 because of uRPF deployed in loose mode; ICMP unreachables are being sent back but, it will never reach the source because ICMP packets destined to 170.170.170.170 will be routed to Null0:

```
R1:
interface Loopback1
 ip address 170.170.170.170 255.255.255.255

R1#ping 150.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#ping 150.1.3.3 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 170.170.170.170 .....
```

https://t.me/learningnets

```
Success rate is 0 percent (0/5)
!
!R3#show ip interface fastEthernet0/0.13 | section IP verify
  IP verify source reachable-via ANY 5 verification drops
   0 suppressed verification drops
   0 verification drop-rate
!
!R3#show ip traffic | section Drop
IP statistics:
  Drop:  0 encapsulation failed, 0 unresolved, 0 no adjacency 0 no route, 5 unicast RPF, 0 forced drop

         0 options denied
  Drop:  0 packets with source IP address zero
  Drop:  0 packets with internal loop back IP address
         0 physical broadcast
  Drop due to input queue full: 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Dynamic Trunking Protocol

Task

- Ensure that SW1's active interfaces that are STP enabled for VLAN 13 and VLAN 23 do not negotiate trunking, regardless of the configured port administrative mode.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.

Overview

By default, all non-modular Cisco Catalyst switches running IOS/IOS XE have all ports enabled as layer 2 ports with Dynamic Trunking Protocol (DTP) enabled. The only exception is for switches that have a dedicated out-of-band management port that is a layer 3 port and cannot be converted to layer 2, such as Catalyst 3750-X. DTP is the magic protocol that makes two switches negotiate the interconnecting links as trunk, as well as the trunking protocol (802.1q or ISL, with ISL having priority over 802.1q), without any required configurations. This might not always be the case, because depending on the switch model, there are two possible DTP default port states:

- Dynamic Desirable (DTP Active): the port actively sends DTP messages, so it initiates trunk formation.
- Dynamic Auto (DTP Passive): the port waits for DTP messages from the other end in order to respond and negotiate the trunk formation.

If you connect two switches that outside of the box have ports in Dynamic Auto mode (as is the case for Catalyst 3750-X), no trunk would be formed because there is no switch to initialize the DTP negotiation.

In common trunk port configurations, at a minimum you specify the trunking protocol

https://t.me/learningnets

and administratively set the port as trunk, with the following interface-level commands:

```
switchport trunk encapsulation dot1q
switchport mode trunk
```

In common access port configuration, at a minimum you specify the VLAN membership and administratively set the port as access, with the following interface-level commands:

```
 switchport access vlan 13
 switchport mode access
```

Ports administratively configured as trunks still have DTP enabled, whereas ports administratively configured as access have DTP disabled. The logic is simple:

- Even if you configure a port as static trunk, you still want DTP enabled because the other end of the link might not yet be configured as static trunk, and you don't want to break it.
- After you have configured a port as static access, you do not want it to be trunk, so there is no need to leave DTP enabled.

The inteface-level command to manually disable DTP is `switchport nonegotiate` . You might want to use this command on access ports just to make it visible in the configuration, whereas on trunk ports it is mandatory to disable DTP. The command that implicitly disables DTP on access ports is `switchport mode access` and not `switchport access vlan <vlan_number>` .

Configuration

SW1:

```
interface range FastEthernet1/0/1 - 2
 switchport nonegotiate
!
interface FastEthernet1/0/20
 switchport nonegotiate
```

https://t.me/learningnets

Verification

Determine which interfaces run STP in VLAN 13 and VLAN 23:

https://t.me/learningnets

```
SW1#show spanning-tree vlan 13


VLAN0013
  Spanning tree enabled protocol ieee
  Root ID    Priority    16397
             Address     001b.5408.f380
             Cost        19
             Port        22 (FastEthernet1/0/20)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32781  (priority 32768 sys-id-ext 13)
             Address     001b.d58b.2180
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Fa1/0/1             Desg FWD 19        128.3    P2p Edge
Fa1/0/20            Root FWD 19        128.22   P2p

!
!SW1#show spanning-tree vlan 23


VLAN0023
  Spanning tree enabled protocol ieee
  Root ID    Priority    16407
             Address     001b.5408.f380
             Cost        19
             Port        22 (FastEthernet1/0/20)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32791  (priority 32768 sys-id-ext 23)
             Address     001b.d58b.2180
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
```

https://t.me/learningnets

```
------------------ ---- --- --------- -------- --------------------------------
Fa1/0/2            Desg FWD 19        128.4    P2p Edge
Fa1/0/20           Root FWD 19        128.22   P2p
```

Verify that DTP is disabled for STP active ports in VLAN 13 and VLAN 23. On access ports, it is visible by the `nonegotiate` keyword and the fact that hello messages are stopped. If the remote end still runs DTP, as is our case for trunk ports because DTP is enabled on SW2, you'll see the dropped packets counter increasing, as each DTP message received inbound is dropped:

```
SW1#show dtp interface fastEthernet1/0/1
DTP information for FastEthernet1/0/1: TOS/TAS/TNS:                              ACCESS/OFF/ACCESS
 TOT/TAT/TNT:                                NATIVE/NEGOTIATE/NATIVE
   Neighbor address 1:                        000000000000
   Neighbor address 2:                        000000000000
Hello timer expiration (sec/state):        never/STOPPED
   Access timer expiration (sec/state):       never/STOPPED
   Negotiation timer expiration (sec/state): never/STOPPED
   Multidrop timer expiration (sec/state):    never/STOPPED
   FSM state:                                 S1:OFF
   # times multi & trunk                      0
   Enabled:                                   no
   In STP:                                    no

   Statistics
   ----------0 packets received (0 good)
 0 packets dropped
        0 nonegotiate, 0 bad version, 0 domain mismatches,
        0 bad TLVs, 0 bad TAS, 0 bad TAT, 0 bad TOT, 0 other
   0 packets output (0 good)
        0 native, 0 software encap isl, 0 isl hardware native
   0 output errors
   0 trunk timeouts
   2 link ups, last link up on Mon Mar 01 1993, 00:15:19
   3 link downs, last link down on Mon Mar 01 1993, 00:16:30
!
!SW1#show dtp interface fastEthernet1/0/20
DTP information for FastEthernet1/0/20:
TOS/TAS/TNS:                               TRUNK/NONEGOTIATE/TRUNK
 TOT/TAT/TNT:                               802.1Q/802.1Q/802.1Q
   Neighbor address 1:                        001B5408F396
   Neighbor address 2:                        000000000000
Hello timer expiration (sec/state):        never/STOPPED
```

https://t.me/learningnets

```
  Access timer expiration (sec/state):      never/STOPPED
  Negotiation timer expiration (sec/state): never/STOPPED
  Multidrop timer expiration (sec/state):   never/STOPPED
  FSM state:                                S6:TRUNK
  # times multi & trunk                     0
  Enabled:                                  yes
  In STP:                                   no

  Statistics
  ----------79 packets received (72 good)
7 packets dropped

      7 nonegotiate, 0 bad version, 0 domain mismatches,
      0 bad TLVs, 0 bad TAS, 0 bad TAT, 0 bad TOT, 0 other
  72 packets output (72 good)
      72 native, 0 software encap isl, 0 isl hardware native
  0 output errors
  0 trunk timeouts
  3 link ups, last link up on Mon Mar 01 1993, 00:21:48
  3 link downs, last link down on Mon Mar 01 1993, 00:17:01
```

A commmonly used method to identify DTP state for interfaces is to view layer 2 port state information:

```
SW1#show interfaces fastethernet1/0/1 switchport
Name: Fa1/0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: nativeNegotiation of Trunking: Off

Access Mode VLAN: 13 (R1-R3)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: none
Operational private-vlan: none
```

https://t.me/learningnets

```
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
Appliance trust: none
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Port Security

Task

- Configure SW1 to guard against MAC address flooding attacks originated at R1, R2, and R3.
- Limit the number of MAC addresses learned simultaneously on a port to one per VLAN, and configure VLAN 100 as the voice VLAN on access ports:
  - In case of violation from R1, simply drop the traffic.
  - In case of violation from R2, drop offending packets and generate log records of the violation.
  - In case of violation from R3 on VLAN 133, apply the err-disabled state without affecting traffic on VLAN 143 and ensure automatic recovery after 3 minutes.
- Age the learned secure entries after 10 minutes of inactivity.
- Retain the MAC addresses learned on the port connected to R2 in the switch configuration.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Port security is a layer 2 feature that enforces a limit on the number of MAC addresses allowed per port. The two main purposes are to prevent unauthorized

https://t.me/learningnets

connections (from unauthorized/unknown MAC addresses) on a port and to thwart MAC-address flooding attacks. A MAC address flooding attack consists of sending a barrage of packets with different source MAC addresses, forcing the switch to overpopulate its MAC address table. This attack may result in switch performance degradation and privacy violation. The latter may occur in cases when the switch starts behaving like a hub, flooding frames out all ports and all VLANs because the MAC address table overflows, exceeding the maximum number of MAC address that the switch can learn.

Note that port security works only on ports configured as static access or static trunk. The port security feature does not work on dynamic ports on which operational port state is negotiated through DTP. Port security logic is as follows:

- On a port configured for port security, the switch keeps a table of secure MAC address entries. The total number of entries allowed in this table is configurable using the interface-level command `switchport port-security maximum-address <number> .`
- On trunk ports, the above command specifies the maximum number of MAC addresses for all VLANs, the aggregate limit. Note that the switch treats the same MAC address on different VLANs as two different MAC addresses. For trunk ports, you may additionally specify the maximum number of MAC addresses per VLAN using the interface-level command `switchport port-security maximum <number> vlan <vlan-number> .`
- If the port is an access-port configured with both data and voice VLANs, you can use commands to impose restrictions on any of the two VLANs with the interface-level command `switchport port-security maximum <number> vlan [access|voice] .`

When a switch has reached the maximum allowed number of MAC addresses on the port level or VLAN level, and a frame with a new source MAC address arrives on the port, the switch may take any of the following actions:

- **Shutdown:** The port actually enters in an **err-disabled** state and all frames received on the port are discarded.
- **Shutdown VLAN:** The VLAN enters in an **err-disabled** state but only for the port where the violation occurred, and all frames received on that port for the respective VLAN are discarded. In this case, a syslog message is also generated.
- **Protect:** All frames are silently discarded on the VLAN where the violation occurred for the respective port. Protect mode is not recommended for trunk ports because as soon as any VLAN on a trunk reaches its MAC address limit, the port stops learning MAC addresses on any other VLAN. The worst thing about this mode is that the switch does not notify you with a logging message.
- **Restrict:** All frames are discarded for the respective port, but a syslog message and

https://t.me/learningnets

SNMP trap are generated. You must additionally configure the SNMP hosts to send the actual traps.

The default port security violation action, unless otherwise configured, is **shutdown**. The switch does not allow the same MAC address to appear on more than one secure port at the same time. Thus, if a switch has learned a MAC address on a secure port, it will not allow the same address to appear on other layer 2 ports until the secure entry has expired. The switch ages out secure MAC address entries using a configurable timeout. You can set the timeout and its functional mode per port using following two commands:

```
switchport port-security aging timeout <timeout>
switchport port-security aging type {absolute|inactivity}
```

Together, these commands define the aging behavior. Absolute aging instructs the switch to age out each MAC address entry when the timeout period has elapsed, so it is unconditional. Inactivity aging instructs the switch to age out each MAC address entry only if it has been inactive for an interval equal to the timeout period, so it is conditional.

If the port security feature has shutdown a port, the port can be restored to an operational state using the automatic error-disable recovery procedure, or manually by issuing a `shutdown` command followed by a `no shutdown` command on the port. There are multiple possible reasons that can trigger a port to enter the err-disabled state, so we must specify both the cause for which the port entered in the err-disabled state, and the interval for keeping it in this state. The interval is a global value, which affects the switch behavior for all possible err-disabled causes:

```
errdisable recovery cause <cause>
errdisable recovery interval <seconds>
```

In addition to setting up dynamic learning of secure MAC addresses, you may configure static secure MAC address entries using the interface-level command `switchport port-security mac-address <mac-address>` . The static entries also count against the maximum number of allowed MAC addresses on an interface. You may configure a port to age static secure MAC address entries as well by using the interface-level command `switchport port-security aging static` . This may be useful when you need to set up guaranteed access for a specific MAC address for some amount of time.

The last port-security feature is known as *sticky learning*. It allows you to transform dynamically learned MAC addresses into static secure MAC addresses. When a

https://t.me/learningnets

switch learns new MAC addresses on a port in sticky mode, it generates a configuration line for the corresponding MAC address as a secure static entry. This line appears in the running configuration, so you need to save it to make the static entry truly permanent; otherwise, if the switch reloads the command is lost. To enable the feature, use the interface-level command `switchport port-security mac-address sticky` . Intead of manual configuration saving, a kron policy or EEM script can be used to automatically save the configuration periodically or triggered by an event.

Last but most important, for the feature to be enabled it is required to use the interface-level command `switchport port-security .`

Configuration

SW1:

```
interface FastEthernet1/0/1
 switchport voice vlan 100
 switchport port-security
 switchport port-security maximum 2
 switchport port-security maximum 1 vlan voice
 switchport port-security maximum 1 vlan access
 switchport port-security violation protect
 switchport port-security aging time 10
 switchport port-security aging type inactivity
!
!
interface FastEthernet1/0/2
 switchport voice vlan 100
 switchport port-security
 switchport port-security maximum 2
 switchport port-security maximum 1 vlan voice
 switchport port-security maximum 1 vlan access
 switchport port-security mac-address sticky
 switchport port-security violation restrict
 switchport port-security aging time 10
 switchport port-security aging type inactivity
!
!
interface FastEthernet1/0/3
 switchport port-security
 switchport port-security maximum 2
 switchport port-security maximum 1 vlan 133
 switchport port-security maximum 1 vlan 143
```

https://t.me/learningnets

```
 switchport port-security violation shutdown vlan
 switchport port-security aging time 10
 switchport port-security aging type inactivity
!
!
errdisable recovery cause psecure-violation
errdisable recovery interval 180
```

Verification

Verify that ports are configured as requested, as well as the err-disabled automatic port recovery function:

```
SW1#show port-security interface fastEthernet1/0/1
Port Security              : Enabled
Port Status                : Secure-upViolation Mode             : Protect
Aging Time                 : 10 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : d48c.b5da.d548:13
Security Violation Count   : 0
!
!SW1#show port-security interface fastEthernet1/0/2
Port Security              : Enabled
Port Status                : Secure-upViolation Mode             : Restrict
Aging Time                 : 10 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : d867.d9e0.bbc0:23
Security Violation Count   : 0
!
!SW1#show port-security interface fastEthernet1/0/3
Port Security              : Enabled
Port Status                : Secure-upViolation Mode             : Shutdown Vlan
Aging Time                 : 10 mins
Aging Type                 : Inactivity
```

https://t.me/learningnets

```
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 2
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 001e.f779.4770:133
Security Violation Count   : 0
!
!SW1#show errdisable detect | i secure
psecure-violation            Enabled      port/vlan
!
!SW1#show errdisable recovery | i secure|interval
psecure-violation            Enabled
Timer interval: 180 seconds
```

Verify that the MAC address limit is at the VLAN level:

```
SW1#show port-security interface fastEthernet1/0/1 vlan

Default maximum: not set, using 6144
VLAN  Maximum    Current
   13          1          1
  100          1          0SW1#show port-security interface fastEthernet1/0/2 vlan
Default maximum: not set, using 6144
VLAN  Maximum    Current
   23          1          1
  100          1          0SW1#show port-security interface fastEthernet1/0/3 vlan

Default maximum: not set, using 6144
VLAN  Maximum    Current
  133          1          1
  143          1          1
```

Verify that interface FastEthernet1/0/2 uses sticky learning:

```
SW1#show port-security address
          Secure Mac Address Table
-----------------------------------------------------------------------
Vlan    Mac Address       Type                       Ports   Remaining Age
                                                                (mins)
----    -----------       ----                       -----   -------------
  13    d48c.b5da.d548    SecureDynamic              Fa1/0/1       10 (I)
23    d867.d9e0.bbc0    SecureSticky               Fa1/0/2         -
```

https://t.me/learningnets

```
 133    001e.f779.4770    SecureDynamic              Fa1/0/3       10 (I)
 143    001e.f779.4770    SecureDynamic              Fa1/0/3       10 (I)
----------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)     : 1
Max Addresses limit in System (excluding one mac per port) : 6144
!
!SW1#show running-config interface fastEthernet1/0/2
Building configuration...

Current configuration : 628 bytes
!
interface FastEthernet1/0/2
 description *** TO R2 Gi0/0 ***
 switchport access vlan 23
 switchport mode access
 switchport nonegotiate
 switchport voice vlan 100
 switchport port-security maximum 2
 switchport port-security maximum 1 vlan access
 switchport port-security maximum 1 vlan voice
 switchport port-security
 switchport port-security aging time 10
 switchport port-security violation restrict
 switchport port-security aging type inactivity
 switchport port-security mac-address sticky
 switchport port-security mac-address sticky d867.d9e0.bbc0 vlan access

 logging event spanning-tree
 spanning-tree portfast
end
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Storm Control

Task

- Configure SW1 to limit unicast traffic received from R1 to 100 pps.
- Configure SW1 to limit broadcast traffic received from R2 to 10 Mbps.
- Configure SW1 to limit unicast traffic received from R3 to 1 Mbps using a percentage of the interface bandwidth.
- If traffic limits are exceeded, ensure that a syslog messages is generated and an SNMP trap is sent to 136.1.13.33 using the v2c community string of **cisco**.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Storm control is used to limit the amount of unicast, multicast, or broadcast traffic received inbound on a layer 2 port. The most common application of this feature is to prevent broadcast storms, but it can also be used to police individual ports so that they do not exceed a desired rate. Depending on the version of IOS, the storm-control command may take units in interface bandwidth percentage, packets per second, or bits per second. Make sure to use the question mark when implementing this command so that the units entered achieve the desired result. The command to enable the feature at the interface level is `storm-control [unicast|multicast|broadcast [level [bps|pps] <rising-threshold> <falling-threshold> .`

https://t.me/learningnets

When a violation of the configured thresholds occurs, the switch can take one or more of the following actions:

- **Syslog:** This is the default and non-configurable action when the configured rising threshold is reached; the switch will silently drop exceeding traffic until it goes bellow falling threshold, and a syslog message is generated similar to the following:
  `%STORMCONTROL-3-FILTERED: A Unicast storm detected on Fa1/0/1. A packet filter action has been applied on the interface .`
- **Trap:** A syslog message is generated and also an SNMP trap (considering that SNMP is configured appropriately) when the rising threshold is reached, in addition to traffic being silently dropped until it goes below falling threshold.
- **Shutdown:** A syslog message is generated and the port enters the **err-disabled** state. You can recover the port from the err-disabled state by using the same options as described in the port security task: manually with a **`shutdown/no shutdown`** sequence of commands on the port, or automatically byusing the err-disabled recovery feature.

Note that the syslog message is generated at all times, and you can have all options simultaneously enabled on the port. For example, you can configure both **trap** and **shutdown** actions at the same time. To configure thresholds violation actions, in addition to to the default of generating a syslog message, the interface-level command of **`storm-comtrol action [trap|shutdwon]`** is used.

Configuration

SW1:

```
interface FastEthernet1/0/1
 storm-control unicast level pps 100
 storm-control action trap
!
!
interface FastEthernet1/0/2
 storm-control broadcast level bps 10m
 storm-control action trap
!
!
interface FastEthernet1/0/3
 storm-control unicast level 1
 storm-control action trap
!
!
```

https://t.me/learningnets

```
snmp-server enable traps storm-control trap-rate 60
snmp-server host 136.1.13.33 traps version 2c cisco
```

Verification

First, verify the configuration of storm control and SNMP; note that there is no specific command to verify the storm control violation actions, other than variations of `show running-config`:

```
SW1#show storm-control unicast

Interface  Filter State   Upper        Lower        Current
---------  -------------  -----------  -----------  ----------
Fa1/0/1       Forwarding          100 pps      100 pps         0 pps
Fa1/0/3       Forwarding        1.00%        1.00%        0.00%

!
!SW1#show storm-control broadcast

Interface  Filter State   Upper        Lower        Current
---------  -------------  -----------  -----------  ----------
Fa1/0/2       Forwarding          10m bps      10m bps         0 bps
!
!SW1#show snmp host

Notification host: 136.1.13.33  udp-port: 162   type: trap
user: cisco     security model: v2c
```

Second, issue `debug snmp packets` on SW1 and generate traffic from R1 to exceed the configured threshold:

https://t.me/learningnets

```
SW1#debug snmp packets

SNMP packet debugging is on
!
!R1#ping 136.1.13.3 repeat 200

Type escape sequence to abort.
Sending 200, 100-byte ICMP Echos to 136.1.13.3, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (200/200), round-trip min/avg/max = 1/1/4 ms
```

On SW1, you should see both the syslog message and SNMP traps being generated:

```
%STORM_CONTROL-3-FILTERED: A Unicast storm detected on Fa1/0/1. A packet filter action has been applied on the inter
!
!
SNMP: Queuing packet to 136.1.13.33
SNMP: V2 Trap, reqid 40, errstat 0, erridx 0
 sysUpTime.0 = 90471274 snmpTrapOID.0 = ciscoPortStormControlMIB.0.2

 cpscStatus.10001.3 = 3
SNMP: Packet sent via UDP to 136.1.13.33
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Port Blocking

Task

- Configure port Fa1/0/13 and Fa1/0/15 of SW1 as static access ports in VLAN 13.
- Ensure that direct traffic between Fa1/0/13 and Fa1/0/15 is disallowed without using any type of ACL filtering.
- Traffic between any other two hosts within VLAN 13 attached to SW1 should be allowed.
- Additionally, configure SW1 to protect R1 from unknown unicast flooding attacks.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

Protected ports are used to prevent traffic from being directly exchanged at layer 2 between two or more ports that are in the same VLAN. Traffic received inbound on a protected port cannot be sent outbound on another protected port, but traffic received inbound on a protected port can be sent outbound on a non-protected port and vice-versa. This feature is a much smaller subset of the Private VLAN feature, so it is called Private VLAN Edge, and cannot span between multiple physical switches. Note that the configuration affects only data plane traffic, which is hardware switched; for example, control-plane data such as PIM packets are allowed between protected ports because these are CPU/process switched. It is enabled at the interface level by using the command `switchport protected .`

When the switch receives a frame with a destination MAC address for which it does not have an entry in its database, it forwards that frame on all ports where the VLAN on which the frame was received is allowed, except for the port on which it was received. This might be seen as a security issue because hosts may receive traffic that is not destined to them, or as a possible attack because a high volume of such traffic may overwhelm all hosts in that VLAN. You can configure switches so that unicast traffic received with a yet-unknown MAC address is not forwarded out on specific ports. Although you can perform the same for multicast traffic, it only affects layer 2 multicast (sender and receiver are in the same VLAN), because layer 3 multicast is controlled by PIM and IGMP (a layer 3 protocol which runs directly over IP) on the VLAN where multicast receivers exist. It is enabled at the interface level by using the command `switchport block {unicast|multicast} .`

Configuration

SW1:

```
interface range FastEthernet1/0/13 , FastEthernet1/0/15
 switchport mode access
 switchport access vlan 13
 switchport protected
 no shutdown
!
interface FastEthernet1/0/1
 switchport block unicast
```

Verification

https://t.me/learningnets

Verify that both ports are configured as protected and R1 is guarded against unknown unicast flooding attacks:

```
SW1#show interfaces fastEthernet1/0/13 switchport | i Protected
Protected: true
!
!SW1#show interfaces fastEthernet1/0/15 switchport | i Protected
Protected: true
!
!SW1#show interfaces fastEthernet1/0/1 switchport | i block
Unknown unicast blocked: enabled

Unknown multicast blocked: disabled
```

On these ports, Ethernet0/1 interfaces of ASA1 and ASA2 are connected; configure these with IP addresses from VLAN 13 and verify that traffic is disallowed, while traffic to other hosts in VLAN 13 is allowed:

```
ASA1:
interface Ethernet0/1
 nameif VLAN13
 ip address 136.1.13.11 255.255.255.0
 no shutdown
!
!ASA2:

interface Ethernet0/1
 nameif VLAN13
 ip address 136.1.13.12 255.255.255.0
 no shutdown
```

Verify that connectivity within VLAN 13 works as expected:

```
ASA1#ping 136.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.1, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
!
!ASA1# ping 136.1.13.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
!
```

https://t.me/learningnets

```
!ASA1# ping 136.1.13.12

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.12, timeout is 2 seconds:?????

Success rate is 0 percent (0/5)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Private VLAN

Task

- Considering that SW2's ports Fa1/0/1, Fa1/0/2, Fa1/0/12, and Fa1/0/14 have been allocated to VLAN 200, configure the following intra-VLAN restrictions:
  - Fa1/0/1 can communicate with Fa1/0/2 and Fa1/0/14.
  - Fa1/0/12 can communicate only with Fa1/0/14.
- If additional VLANs are required, use VLANs 201 and 202.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Without any special configurations, a VLAN represents a broadcast domain where all hosts reside in the same subnet and can communicate directly (at layer 2) to each other without any restrictions or the need for a layer 3 (default gateway) device. For inter-VLAN communication, you need a layer 3 device to connect multiple broadcast domains and forward packets between.

What if you wanted to restrict direct (layer 2) communication between hosts within the same VLAN based on some rules, and you needed this configuration to span across multiple switches? This is why Private VLANs (PVLANs) exist. PVLANs allows you to logically divide one broadcast domain (one primary VLAN) into multiple broadcast sub-domains (sub-VLANs or secondary VLANs) while keeping all

https://t.me/learningnets

hosts addressed within the same IP subnet. When this is configured, all traffic within the VLAN is allowed/restricted based on PVLAN rules. Malicious hosts/users could bypass these rules by forcing communication through a layer 3 device within the same subnet (which usually is the default gateway); however, for this to function the layer 3 device must have Local Proxy ARP enabled. Local Proxy ARP can be enabled at the interface level using the command `ip local-proxy-arp`, and is disabled by default. By configuring appropriate IP access-list rules at the default gateway level, you're sure to disallow communication within the same subnet through a layer 3 device.

Common implementations of PVLANs are on switches managed by ISPs, where multiple customers belong to the same VLAN, but some rules regarding direct communication between customers are required. Another example is a DMZ network, in which you have multiple services (like SMTP, FTP, WWW) within the same VLAN and IP subnet, but you want to configure some restrictions so that SMTP servers can communicate with FTP servers but not with WWW servers, for example.

In PVLAN configuration, the following new port type concepts are introduced:

- **Promiscuous port:** connects to a shared resource (such as the default gateway) that is accessible by all hosts within allowed secondary VLANs.
- **Isolated port:** connects to a secondary isolated VLAN type, and is allowed to communicate only with promiscuous ports.
- **Community port:** connects to a secondary community VLAN type, and is allowed to communicate with promiscuous ports or ports within the same secondary community VLAN.

In PVLAN configuration, the following new VLAN concepts are introduced:

- **Primary VLAN:** This is the VLAN that all hosts initially belonged to. It can have only one isolated secondary VLAN and multiple community secondary VLANs associated, and it is used to forward frames downstream from **promiscuous** ports to **isolated** or **community** ports.
- **Secondary isolated VLAN:** This is a VLAN associated with only one primary VLAN and is used to forward frames upstream from **isolated** ports to **promiscuous** ports.
- **Secondary Community VLAN:** This is a VLAN associated with only one primary VLAN and is used to forward frames upstream from **community** ports to **promiscuous** ports, or between hosts within the same community VLAN.

From the configuration perspective, use the following steps:

https://t.me/learningnets

1. If your switch runs VTP version 1 or VTP version 2, configure it in VTP transparent mode using the global command `vtp mode transparent` .
2. The secondary VLANs are created and defined as either community or isolated using the VLAN-level command `private-vlan [isolated|community]` .
3. The primary VLAN is defined and the secondary VLANs are associated with the primary VLAN using the VLAN-level command `private-vlan primary` followed by the command `private-vlan association <secondary-vlan>` .
4. Configure layer 2 interfaces using the command `switchport mode private-vlan promiscuous` or `switchport mode private-vlan host` . As you might guess, the promiscuous option defines that the port role is promiscuous, and the host option defines that the port role as either community or isolated. The interface role is specified as isolated or community and attached to the required secondary VLAN using the interface-level command `switchport private-vlan host-association <primary-vlan> <secondary-vlan>` .
5. Configure which secondary VLANs can communicate with a promiscuous port by using the promiscuous interface-level command `switchport private-vlan mapping <primary-vlan> <secondary-vlan>` ; basically, this command functions similar to `switchport trunk allowed vlan <vlan-number>` because you can optionally restrict which secondary VLANs can communicate with a specific promiscuous port.

Note that if your private VLAN configuration spans multiple switches, you must manually configure all settings on all switches. VTP version 1 and VTP version 2 miss the necessary TLV values to propagate PVLAN information, and because of this you must set your VTP in transparent mode when PVLAN is configured. VTP version 3 allows for PVLAN configuration advertisements, so you need to do the VLAN related configuration only on the switch designated as primary in your VTP domain; all other switches running in VTP client/server mode will dynamically learn the PVLAN configuration (switches running in VTP transparent or off mode will not learn it). So if VTP version 3 is being used, only interface-related configuration need be manually performed on all switches.

Based on task requirements, we must configure Fa1/0/1 and Fa1/0/2 as community ports, Fa1/0/14 as a promiscuous port, and Fa1/0/12 as an isolated port.

Configuration

SW2:

```
vtp mode transparent
!
vlan 201
```

https://t.me/learningnets

```
 private-vlan isolated
!
vlan 202
 private-vlan community
!
vlan 200
 private-vlan primary
 private-vlan association 201,202
!
!
interface range FastEthernet1/0/1 , FastEthernet1/0/2
 switchport mode private-vlan host
 switchport private-vlan host-association 200 202
 no shutdown
!
interface FastEthernet1/0/12
 switchport mode private-vlan host
 switchport private-vlan host-association 200 201
 no shutdown
!
interface FastEthernet1/0/14
 switchport mode private-vlan promiscuous
 switchport private-vlan mapping 200 201,202
 no shutdown
```

Verification

Verify PVLAN configuration:

```
SW2#show vlan private-vlan


Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
200     201       isolated          Fa1/0/12, Fa1/0/14
200     202       community         Fa1/0/1, Fa1/0/2, Fa1/0/14
!
!SW2#show interfaces fastEthernet1/0/1 switchport

Name: Fa1/0/1
Switchport: Enabled
Administrative Mode: private-vlan host
Operational Mode: private-vlan host
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
```

https://t.me/learningnets

```
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: noneAdministrative private-vlan host-association: 200 (VLAN0200) 202 (VLAN0202)

Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: noneOperational private-vlan:
 200 (VLAN0200) 202 (VLAN0202)

Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
Appliance trust: none
!
!SW2#show interfaces fastEthernet1/0/14 switchport

Name: Fa1/0/14
Switchport: Enabled
Administrative Mode: private-vlan promiscuous
Operational Mode: private-vlan promiscuous
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: 200 (VLAN0200) 201 (VLAN0201) 202 (VLAN0202)

Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
```

https://t.me/learningnets

```
Administrative private-vlan trunk mappings: noneOperational private-vlan:
  200 (VLAN0200) 201 (VLAN0201) 202 (VLAN0202)

Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
Appliance trust: none
```

Assign IP addresses on devices connected to SW2's ports using IP subnet 136.1.200.0/24:

```
R1:
interface GigabitEthernet0/1
 ip address 136.1.200.1 255.255.255.0
 no shutdown
!
!R2:
interface GigabitEthernet0/1
 ip address 136.1.200.2 255.255.255.0
 no shutdown
!
!ASA1:
interface Ethernet0/0
 nameif VLAN200
 ip address 136.1.200.11 255.255.255.0
 no shutdown
!
!ASA2:

interface Ethernet0/0
 nameif VLAN200
 ip address 136.1.200.12 255.255.255.0
 no shutdown
```

Verify that traffic within VLAN 200 is restricted as requested:

```
R1#ping 136.1.200.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.200.2, timeout is 2 seconds:!!!!!
```

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R1#ping 136.1.200.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.200.11, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R1#ping 136.1.200.12
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.200.12, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!ASA1# ping 136.1.200.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.200.2, timeout is 2 seconds:?????
Success rate is 0 percent (0/5)
!
!ASA1# ping 136.1.200.12
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.200.12, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Private VLAN Propagation

Task

- Ensure that VLAN database configurations related to PVLAN from the previous task are dynamically propagated from SW2 to SW1.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> For this task to be functional, you must have completed the previous PVLAN task and keept the resulting configuration files; or, simply copy and paste the configurations from previous PVLAN task.

Overview

For VTP to announce PVLAN configuration changes to the VLAN database, VTP version 3 must be used. VTP version 3 comes with multiple VLAN database security improvements, the most significant being that only the VTP switch designated as **primary** can update the VLAN database within one VTP domain, regardless of the configuration revision number value. The switch designated as **primary** must run in server mode. Note that the command `vtp primary` is configured from `#` mode. VTP version 3 is modular, in that it supports advertisements propagation for:

- **VLAN** database configuration
- **MST** configuration
- **Unknown**, reserved for future use

For each of the above modules, a switch can run in the following modes:

- Server
- Client

https://t.me/learningnets

- Transparent
- Off

The difference between transparent and off is that in off mode, advertisements are not relayed over trunk ports but silently dropped. VTP can now be disabled at the interface level with the command `no vtp`, even though globally it remains enabled.

Configuration

```
SW1:
!
vtp version 3


SW2:

!
vtp version 3
vtp mode server
!
end
vtp primary
```

Verification

Check the VTP status and PVLAN configuration propagation to SW1:

https://t.me/learningnets

```
SW1#show vlan private-vlan


Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
200     201       isolated
          200     202       community

!
!SW2#show vtp devices

Retrieving information from the VTP domain. Waiting for 5 seconds.

VTP Feature  Conf Revision Primary Server Device ID      Device Description
------------ ---- -------- -------------- -------------- ----------------------
VLAN         No   1        001c.0e58.4380 0024.1406.4b80 SW1
```

Verify that SW2 is elected as primay VTP server:

```
SW2# show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 3
VTP Domain Name                 : SECTION2
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : 001c.0e58.4380

Feature VLAN:
--------------VTP Operating Mode                : Primary Server

Number of existing VLANs          : 15
Number of existing extended VLANs : 0
Configuration Revision            : 1
Primary ID                        : 001c.0e58.4380
Primary Description               : Rack1SW2
MD5 digest                        : 0x8F 0xF0 0x47 0x66 0xDA 0xEC 0x1D 0x63
                                    0x74 0x41 0xB3 0x51 0xA9 0x08 0x5F 0xD2



Feature MST:
--------------
VTP Operating Mode                : Transparent
```

https://t.me/learningnets

```
Feature UNKNOWN:
--------------
VTP Operating Mode                : Transparent
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch Port ACL

Task

- Configure SW1 so that R1 cannot telnet to R3; all other traffic should not be affected.
- Your configuration should not be applied at VLAN level.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Because an IP access-list (also known as a router ACL or RACL) that filters IPv4 traffic can be applied inbound on a layer 2 port, the feature is called port ACL (PACL). You cannot apply a PACL in the outbound direction or more than one PACL in the inbound direction. PACL is used in most 802.1x deployments to restrict inbound traffic before and after 802.1x authentication.

Configuration

SW1:

```
ip access-list extended DENY_TELNET
 deny tcp host 136.1.13.1 host 136.1.13.3 eq 23
 permit ip any any
!
```

https://t.me/learningnets

```
interface FastEthernet1/0/1
 ip access-group DENY_TELNET in
```

Verification

Verify that only telnet traffic is filtered and access-list is correctly applied:

```
R1#ping 136.1.13.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#telnet 136.1.13.3
Trying 136.1.13.3 ... % Connection timed out; remote host not responding
!
!SW1#show ip interface fastEthernet1/0/1
FastEthernet1/0/1 is up, line protocol is up Inbound  access list is DENY_TELNET
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch MAC ACL

Task

- You have been alerted that high volumes of non-IP traffic transit VLAN 13. Configure SW1 to drop all such traffic received on its Fa1/0/1 interface, except for ARP.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

MAC access-lists are used to match and filter non-IP traffic. A single MAC access-list can be applied inbound on a layer 2 port, to filter traffic such as ARP, STP, and IPX. Note that although the name is MAC access-lists, it is not actually used to filter traffic based on Ethernet MAC addresses; for this you must use the port-security feature. Recall that layer 2 OSI (Data Link Layer) has two sub-layers: logical link control (LLC) and media access control (MAC); this access-list operates at the MAC sublayer, hence the name.

From the configuration perspective, you first create the access-list using the global command `mac access-list extended <acl-name>`. Then simply create ACL rules just like in IP access-lists, but with the following available parameters:

- **deny|permit**: the action to take in case of frame match
- **source**: where you specify the source MAC address of the frame; the `any` keyword

https://t.me/learningnets

can be used instead

- **destination**: where you specify the destination MAC address of the frame; the `any` keyword can be used instead
- **protocol**: where you specify the non-IP protocol to match based on its Ethertype value.

We are not required to match on a specific source or destination MAC address, but we must match the ARP protocol exactly, thus the protocol value of **0x806** and the mask of **0x0** for an exact match. The ethertype mask functions the same way as a wildcard mask from IP access-lists, where a bit of zero means match and a bit of one means don't care.

Configuration

```
SW1:


!
mac access-list extended PERMIT_ARP
 permit any any 0x806 0x0
!
interface FastEthernet1/0/1
 mac access-group PERMIT_ARP in
```

Verification

Ensure that ARP traffic is allowed:

```
R1#debug arp
R1#clear ip arp 136.1.13.3
!
IP ARP: sent req src 136.1.13.1 d48c.b5da.d548,
                 dst 136.1.13.3 001e.f779.4771 GigabitEthernet0/0
IP ARP: rcvd rep src 136.1.13.3 001e.f779.4771, dst 136.1.13.1 GigabitEthernet0/0
IP ARP: creating entry for IP address: 136.1.13.3, hw: 001e.f779.4771
!R1#ping 136.1.13.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.3, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

https://t.me/learningnets

Verify that MAC access-list is applied:

```
SW1#show mac access-group interface fastEthernet1/0/1
Interface FastEthernet1/0/1: Inbound access-list is PERMIT_ARP

   Outbound access-list is not set
```

Verify that ARP traffic is blocked by modifying the MAC ACL:

```
SW1:
!
mac access-list extended PERMIT_ARP
 no permit any any 0x806 0x0
 permit any any 0x805 0x0
!
!R1#clear ip arp 136.1.13.3
!IP ARP: sent req src 136.1.13.1 d48c.b5da.d548,
 dst 136.1.13.3 001e.f779.4771 GigabitEthernet0/0
!R1#ping 136.1.13.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.3, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)
```

After testing, configure the ACL back:

```
SW1:

!
mac access-list extended PERMIT_ARP
 no permit any any 0x805 0x0
 permit any any 0x806 0x0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Switch VLAN ACL

Task

- Configure SW2 so that no telnet traffic is allowed across VLAN 23.
- Configure SW2 so that only the following layer 2 protocols are allowed within VLAN 23:
  - CDP, VTP, UDLD
  - STP, ARP

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

VLAN ACL (VACL) is a filtering method applied at the VLAN level; it can be used to filter both IP and non-IP traffic, for both intra-VLAN traffic (bridged) or inter-VLAN traffic (routed). The VACL has no direction, it just inspects all IP and/or non-IP packets crossing the VLAN against the configured rules.

https://t.me/learningnets

To activate the function for IP or non-IP traffic, at least one rule for each type of traffic needs to be configured, otherwise traffic is allowed and not inspected by the VACL. For example, if you configure VACL rules for only non-IP traffic, all IP traffic is allowed without even being inspected by the VACL. However, after a rule is configured for IP traffic, all IP traffic traversing the VLAN is inspected and appropriate action is taken as configured.

To enable VACL, follow these steps:

1. Configure a MAC ACL and/or a IP ACL to match on the non-IP/IP traffic.
2. Configure VLAN access-map rules for IP and/or non-IP traffic.
3. Apply the VLAN access-map to one or multiple VLANs.

VLAN access-maps functionality is very similar to route-maps:

- You use sequence numbers to configure multiple entries within the same access-map, and the map is processed top down until a match occurs.
- Within each rule, you identify traffic using `match` commands, but you can only match on IP or MAC access-lists.
- If no `match` command is present within a rule, it matches on all IP and non-IP traffic.
- There are only two possible actions to be configured for each entry: `drop` or `forward`, where `drop` is similar to `deny` from access-lists and `forward` is similar to `permit`.
- When a rule is configured for IP traffic, if a packet doesn't match any of the configured rules, it will be implicitly denied.
- When a rule is configured for non-IP traffic, if a packet doesn't match any of the configured rules, it will be implicitly denied.

If all possible filters are enabled for a packet, the order of operation is as follows:

1. Port MAC ACL (applied inbound on a layer 2 port).
2. Port IP ACL (applied inbound on a layer 2 port).
3. VLAN ACL (applied at the VLAN level).
4. Router Inbound ACL (applied inbound on a layer 3 port).
5. Router Outbound ACL (applied inbound on a layer 3 port).

When you perform Layer 2 filtering, you usually do that based on EtherTypes or LSAP (Link Service Access Point) values. LSAP value consists of two parts (one byte each): SSAP (Source SAP) and DSAP (Destination SAP). You can apply Layer 2 filtering based on MAC addresses, but this is not commonly used; most of the time you'll just match on the non-IP protocols.

IEEE STP BPDUs use 802.2 LLC encapsulation with SSAP/DSSP values of **0x42** or LSAP value of **0x4242**. Cisco switches and routers run IEEE STP on access-ports

https://t.me/learningnets

and also on the native VLAN of an 802.1q trunk. STP packets are sent across ISL trunks using the same LSAP value of **0x42**. Cisco's implementation does not modify the STP packets sent over ISL trunks, it just adds the ISL encapsulation.

PVST+ SSTP BPDUs use 802.2 SNAP encapsulation with an LSAP value of **0xAAAA** and SNAP PID (Protocol ID) value of **0x010B**. You can match this Protocol ID value as the EtherType number in MAC ACLs on the Catalyst Switches. Cisco switches send PVST+ BPDUs only across 802.1q trunks, in addition to IEEE STP BPDUs. The PVST+ BPDUs appear on native as well as non-native VLANs of a trunk port.

The following protocols use 802.2 SNAP encapsulation with the SNAP Protocol ID values listed below:

- VTP: 0x2003
- CDP: 0x2000
- DTP: 0x2004
- UDLD: 0x0111

All SNAP-encapsulated packets can be matched using an LSAP value of **0xAAAA**. The above mentioned packet types have no VLAN tag header and use the default VLAN 1 or the native VLAN of the 802.1q trunk.

Configuration

SW2:

```
mac access-list extended PERMIT_LAYER2
 permit any any lsap 0x4242 0x0
 permit any any 0x010B 0x0
 permit any any 0x806 0x0
 permit any any lsap 0xAAAA 0x0
!
ip access-list extended DENY_TELNET
 permit tcp any any eq 23
!
ip access-list extended PERMIT_IP
 permit ip any any
!
vlan access-map VLAN23 10
 match mac address PERMIT_LAYER2
 action forward
vlan access-map VLAN23 20
 match ip address DENY_TELNET
 action drop
```

https://t.me/learningnets

```
vlan access-map VLAN23 30
 match ip address PERMIT_IP
 action forward
!
vlan filter VLAN23 vlan-list 23
```

Verification

Verify that telnet is not allowed within or across VLAN 23, but other IP traffic such as ICMP is allowed:

```
R2#ping 136.1.23.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R2#ping 136.1.13.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.13.1, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R2#telnet 136.1.23.3
Trying 136.1.23.3 ... % Connection timed out; remote host not responding
!
!R2#telnet 136.1.13.1
Trying 136.1.13.1 ... % Connection timed out; remote host not responding
```

Verify that non-IP protocols such as CDP and STP are functional within VLAN 23:

```
SW2#show cdp neighbors  fastEthernet1/0/3 | b Device
Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
R3.ine.com       Fas 1/0/3         150             R S I  2811      Fas 0/1
!
!SW2#show spanning-tree vlan23 root port
VLAN0023         This bridge is root
!
!SW1#show spanning-tree vlan 23 root port

VLAN0023         FastEthernet1/0/20
```

Verify VACL configuration:

https://t.me/learningnets

```
SW2#show vlan filter access-map VLAN23

VLAN Map VLAN23 is filtering VLANs: 23
!
!SW2#show vlan access-map VLAN23
Vlan access-map "VLAN23"  10
  Match clauses:
    mac address: PERMIT_LAYER2
  Action: forward
Vlan access-map "VLAN23"  20
  Match clauses:
    ip  address: DENY_TELNET
  Action: drop
Vlan access-map "VLAN23"  30
  Match clauses:
    ip  address: PERMIT_IP
  Action: forward
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing STP Attacks Using BPDU Guard

Task

- Configure SW2 so that if it receives a BPDU on its Fa1/0/1 - Fa1/0/3 ports, the port goes into err-disabled state.
- The switch should attempt to recover the port automatically after five minutes.
- No configuration changes are allowed on Fa1/0/1 or Fa1/0/2 ports.
- Match the following output:

```
SW2#show spanning-tree vlan 13 interface fastEthernet1/0/3 detail | i guard
Bpdu guard is enabled
```

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The STP BPDU Guard feature is used to enforce access-layer security on the termination of the STP domain. When an interface running BPDU Guard receives a BPDU (STP packet), the interface is transitioned into err-disable state. This ensures that unauthorized switches cannot be plugged in to the network, for example to perform a layer 2 man-in-the-middle (MiTM) attack. If configured, the err-disable recovery function can be used to bring the interface out of err-disable state

https://t.me/learningnets

automatically after a configured interval.

- BPDU Guard can be enabled at the interface level using the command `spanning-tree bpduguard enable`, in which case it's applied unconditionally to the port.
- BPDU Guard can be enabled globally using the command `spanning-tree portfast bpduguard default`, in which case it's applied only to portfast-enabled ports, also called Edge Ports. Note that if you enable portfast (access or trunk) on a port, if it receives a BPDU it gets out of the portfast-enabled state and starts negotiating its STP state, so this feature is conditional.

You can determine whether BPDU Guard is active on an interface by using the command `show spanning-tree interface <interface-number> detail`. If BPDU Guard was enabled at global level, you would see an output such as **Bpdu guard is enabled by default**, whereas if it was enabled at port level, you would see an output such as **Bpdu guard is enabled**. Given the task requirements, we need to enable it at port level for Fa1/0/3 and at global level for Fa1/0/1 - Fa1/0/2.

It's important to remember that BPDU Guard has the same functionality, regardless of how it's enabled.

Configuration

SW2:

```
interface FastEthernet1/0/3
 spanning-tree bpduguard enable
!
!
spanning-tree portfast default
spanning-tree portfast bpduguard default
!
errdisable recovery cause bpduguard
errdisable recovery interval 300
```

Verification

Verify automatic port recovery and that BPDU is enabled as requested:

```
SW2#show spanning-tree vlan 13 interface fastEthernet0/3 detail | i guard
 Bpdu guard is enabled
!
!SW2#show spanning-tree interface fastEthernet0/1 detail | i guard
```

https://t.me/learningnets

```
 Bpdu guard is enabled by default

!
!SW2#show spanning-tree interface fastEthernet0/2 detail | i guard
 Bpdu guard is enabled by default
!
!SW2#show errdisable recovery | i guard|interval
bpduguard                       Enabled
l2ptguard                       DisabledTimer interval: 300 seconds
!
!SW2#show errdisable detect | i guard
bpduguard                       Enabled      port

l2ptguard                       Enabled      port
```

You can also verify that BPDU Guard or PortFast is enabled at global level using the following command:

```
SW2#show spanning-tree summary | i Default
       Portfast Default              is enabled
PortFast BPDU Guard Default  is enabled

Portfast BPDU Filter Default is disabled
Loopguard Default            is disabled
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing STP Attacks Using BPDU Filter

Task

- Configure SW1 so that if a traffic capture is performed on its ports connecting to R1 or R2, no STP topology information will be exposed.
- If a BPDU is received inbound on Fa1/0/1 or Fa1/0/2, only Fa1/0/2 should start negotiating its STP port state.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The BPDU Filter feature, like the BPDU Guard feature, is used to terminate the STP domain, and can also be enabled at interface level or global level. BPDU Filter has different behavior, depending on how it's enabled:

- When configured at the interface level, BPDU Filter unconditionally drops all inbound BPDUs and does not send BPDUs out the interface.
- When configured at global level, it sends out exactly 11 BPDUs on all portfast-enabled interfaces, and if no BPDU is received within this interval it stops sending BPDUs; however, if a BPDU is received inbound on the interface at any time after, it is accepted and STP port state negotiation starts, so BPDU Filter becomes inactive.

Unlike BPDU Guard, the interface does not go into err-disable when a violation

https://t.me/learningnets

occurs. Non-STP traffic will continue to be forwarded in both directions without any restrictions. BPDU Filter can be enabled at the port level using the command `spanning-tree bpdufilter enable` . BPDU Filter can be enabled globally using the command `spanning-tree portfast bpdufilter default` and it gets applied only to portfast-enabled ports, also called Edge Ports. Note that if you enable portfast (access or trunk) on a port, if it receives a BPDU, it gets out of the portfast-enabled state, so this feature is conditional.

If both BPDU Guard and BPDU Filter are enabled at the same time on an interface, the order of operation dictates that BPDU Filter be applied first, so the port will never reach the err-disabled state if a BPDU is received because the inbound BPDU is filtered. The conclusion is that enabling both features at the same time on the same port does not accomplish the desired results. BPDU Filter could be used in ISP topologies where the ISP and customer have a layer 2 connection but don't want to share a common STP domain, so BPDU Filter is applied at port level to restrict any STP communication between ISP and its clients.

Based on the task requirements, we must enable BPDU Filter on Fa1/0/1 at interface level and on Fa1/0/2 at global level. Note that portfast is enabled on Fa1/0/2 from the initial configuration.

Configuration

SW1:

```
interface FastEthernet1/0/1
 spanning-tree bpdufilter enable
!
spanning-tree portfast bpdufilter default
```

Verification

Verify that BPDU Filter is enabled on both ports, and notice the output difference between the two ports:

https://t.me/learningnets

```
SW1#show spanning-tree vlan 13 interface fastEthernet1/0/1 detail | i filter
 Bpdu filter is enabled


!
!SW1#show spanning-tree vlan 23 interface fastEthernet1/0/2 detail | i filter
 Bpdu filter is enabled by default
```

You can also verify whether the BPDU Filter is enabled at global level by using the following command:

```
SW1#show spanning-tree summary | i Filter
 Portfast BPDU Filter Default is enabled
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing STP Attacks Using Root Guard

Task

- Configure SW2 so that its trunk link to SW1 is disabled if SW1 tries to become the STP root bridge for any VLAN.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The scope of Root Guard is to guard/protect the root bridge role in your STP topology from being taken away by unauthorized devices/switches, basically ensuring root bridge placement in the network. Based on your layer 2 network design, you'll almost always manually configure which switch is the root bridge for which VLANs. Based on the STP root bridge selection algorithm (lowest switch priority followed by lowest MAC address) and the fact that the MAC address is not user configurable, even if you configure your switch with the lowest possible priority of zero, you cannot guarantee its root bridge role because a different switch with a lower MAC address can be injected in the network and take it over.

Being related to STP, Root Guard can only be enabled on layer 2 ports. In theory, when enabled on a port, if superior STP packets are received, the port enters into the **Root Inconsistent** state; this state is different than the **err-disabled** state because the port remains in the **UP** state and it is just logically blocked at the STP level, thus restricting any traffic on that port (remains in a STP state similar to **listening**

https://t.me/learningnets

). For this reason, the port can be automatically recovered/unblocked from the **Root Inconsistent** state as soon as the superior BPDUs are no longer received, without any special configurations like in the case of **err-disabled** state. A superior BPDU is defined as carrying information about the root bridge ID that is better than the one currently stored, or indicates a better cost to the root bridge than what is currently installed. Therefore, this feature is used to prevent a potentially rogue device from announcing itself as the new root bridge and possibly implementing a layer 2 man-in-the-middle attack.

In practice, when enabled on a port, there are two possible situations, depending on the STP port state:

- The port is in the designated state and Root Guard ensures that it remains designated; otherwise it is transitioned into **Root Inconsistent** state.
- The port is in the non-designated state (Root or Alternate) and Root Guard immediately transitions it into **Root Inconsistent** state because Root Guard cannot actively modify STP configuration to ensure that the port becomes designated.

In other words, you can say that Root Guard keeps the port enabled only if it's in the designated state, because on any given segment there is only one designated port, the one that has the best STP cost toward the root bridge, and therefore the one that issues the superior BPDU for that segment. Based on this logic, you'll enable it only on key designated ports to protect the root bridge. Enabling Root Guard on non-designated ports will block your upstream ports toward the root bridge, effectively isolating the switch; remember that non-designated ports receive superior BPDUs and do not send any BPDUs out, unless STP bridge assurance is configured.

Root Guard is enabled only at the interface level by using the command `spanning-tree guard root` , and therefore is activated on all VLANs for which STP is active on the respective interface. Even though it is enabled at the port level, it takes actions at the VLAN level. For example, if it receives superior BPDUs for VLAN 100, it will transition the port into **Root-inconsistent** only for VLAN 100, traffic for all other VLANs will not be affected. Root Guard and Loop Guard cannot be configured simultaneously on the same port because of their inherent functionality (Root Guard ensures that the port on which it is enabled is selected as designated, whereas Loop Guard ensures that the port on which it is enabled is not selected as designated).

Based on the initial configurations, SW2 was configured as the root bridge for all VLANs with a non-default priority value of 16384.

Configuration

https://t.me/learningnets

```
SW2:

!
interface FastEthernet1/0/20
 spanning-tree guard root
```

Verification

Verify that Root Guard is enabled for all VLANs:

```
SW2#show spanning-tree interface fastEthernet1/0/20 detail | i VLAN|Root
 Port 23 (FastEthernet0/21) of VLAN0001 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0013 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0023 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0100 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0114 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0133 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0143 is designated forwarding  Root guard is enabled on the port
 Port 23 (FastEthernet0/21) of VLAN0213 is designated forwarding  Root guard is enabled on the port
```

To verify that ports transition into root-inconsistent state if a superior BPDU is received, configure SW1 with a better STP bridge priority than 16384:

```
SW1:

!
spanning-tree vlan 1-4094 priority 8192
```

On SW2, you will see messages logged similar to the following:

```
%SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port FastEthernet1/0/20 on VLAN0001.
%SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port FastEthernet1/0/20 on VLAN0023.
```

Verify that interfaces are in the root inconsistent state:

```
SW2# show spanning-tree interface fastEthernet 1/0/20 detail | i Inconsistent
 Port 22 (FastEthernet1/0/20) of VLAN0001 is broken  (Root Inconsistent)
 Port 22 (FastEthernet1/0/20) of VLAN0013 is broken  (Root Inconsistent)
 Port 22 (FastEthernet1/0/20) of VLAN0023 is broken  (Root Inconsistent)
 Port 22 (FastEthernet1/0/20) of VLAN0100 is broken  (Root Inconsistent)
```

https://t.me/learningnets

```
 Port 22 (FastEthernet1/0/20) of VLAN0114 is broken  (Root Inconsistent)
 Port 22 (FastEthernet1/0/20) of VLAN0133 is broken  (Root Inconsistent)
 Port 22 (FastEthernet1/0/20) of VLAN0143 is broken  (Root Inconsistent)
 Port 22 (FastEthernet1/0/20) of VLAN0213 is broken  (Root Inconsistent)
!
!SW2#show spanning-tree interface fastEthernet1/0/20 inconsistency

VLAN0001            Root Inconsistent
VLAN0013            Root Inconsistent
VLAN0023            Root Inconsistent
VLAN0100            Root Inconsistent
VLAN0114            Root Inconsistent
VLAN0133            Root Inconsistent
VLAN0143            Root Inconsistent
VLAN0213            Root Inconsistent
```

Remove the test configuration to re-establish network connectivity:

```
SW1:


!
no spanning-tree vlan 1-4094 priority 8192
```

On SW2, you will see messages logged similar to the following, confirming that interfaces are active again, removed from the root-inconsistent state:

```
%SPANTREE-2-ROOTGUARD_BLOCK: Root guard unblocking
 port FastEthernet1/0/20 on VLAN0001.%SPANTREE-2-ROOTGUARD_BLOCK: Root guard unblocking
 port FastEthernet1/0/20 on VLAN0023.
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing STP Loops Using Loop Guard

Task

- Based on the initial configurations, SW2 should be the STP root bridge for all VLANs.
- If SW2 fails to send BPDUs out on its Fa1/0/20 port, ensure that no loops are formed if there is inter-switch link redundancy.
- Use the minimum commands necessary to accomplish this task.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Without any special features like bridge assurance configured, on any given segment only the designated port sends BPDUs out, and there is only one designated port per segment. The segment health and convergence from STP point of view relies on these BPDUs being sent and received. If a non-designated port stops receiving BPDUs from the designated port, the following happens:

- If STP/PVST is being used: Every BPDU contains two fields, **Max_Age** and **Message_Age**. The **Message_Age** field is incremented every time a BPDU traverses a switch (so it might be compared to the hop count). When a switch stores the BPDU received on a port, it counts the time in seconds, starting from **Message_Age** and up to **Max_Age**. If during this interval no further BPDUs are received, the current BPDU is wiped out and the port is eventually declared

https://t.me/learningnets

designated. It first enters the listening state, and after two times the **forward_delay** timer, it reaches the forwarding state.
- If RSTP/PVRSTP/MST is being used: The switch waits three times the **hello-interval** for a BPDU to be received, after which it transitions to listening state, and after two times the **forward_delay** time it reaches the forwarding state, basically becoming designated.

Although the built-in STP functionality ensures port converges in the case of indirect link failure, it may lead loops in specific circumstance. SW2 is the root bridge in our scenario and has only one port connected to SW1 (Fa1/0/20), so Fa1/0/20 is the Root Port. Let's imagine that there are three inter-switch links, so that Fa1/0/22-Fa1/0/23 are Alternate Ports in blocking state. If SW2 is alive but fails to send BPDUs out on Fa1/0/22-Fa1/0/23 ports, eventually SW1 will transition these two ports to designated forwarding state and a loop will be formed between SW1 and SW2.

Loop Guard was designed to stop a loop from happening in such cases. Its scope is to prevent accidental loops generated by the inability of designated ports to generate BPDUs (software issues or sustained high CPU) or by missing BPDUs because of unidirectional links. If STP decides to transition a Loop Guard enabled port from non-designated to designated state because of missing BPDUs, the port transitions to **Loop Inconsistent** state, which is very similar to **Root Inconsistent** state: The port remains in the **UP** state and it is just logically blocked at the STP level, thus restricting any traffic on that port (remains in a STP state similar to **listening**).

The following must be considered for Loop Guard:

- Although configured, Loop Guard is not active on an interface unless it is considered point-to-point by STP; for a link to be considered point-to-point, it must be active in full duplex mode.
- If it's configured on a STP designated port, even though commands show it as being active, it doesn't actually monitor the STP port state.
- If it's configured on a STP non-designated port (alternate or root), it monitors the STP port state and prevents it from becoming designated.

You should be aware of the following functionality, given a switch with one Root Port and one Alternate Port (secondary/backup path toward the root bridge) and Loop Guard enabled on both ports:

- If the Root Port is lost (interface link is down), it allows the Alternate Port to become Root Port, because Root Port is also a non-designated port.

https://t.me/learningnets

- If the Root Port is functional and the switch stops receiving BPDUs on the Alternate Port, it will not transition it to Designated Forwarding; instead the port transitions into **Loop Inconsistent** state.
- If at this moment the Root Port goes down, the switch does not recover the **Loop Inconsistent** port unless a BPDU is received, so the switch remains isolated.
- The port is automatically recovered from the **Loop Inconsistent** state only after it receives a BPDU and starts negotiating its STP port state.

Note that Root Guard and Loop Guard cannot be enabled at the same time on any given interface because they serve opposite purposes:

- Root Guard should be used on designated ports and ensures that these do not transition to non-designated.
- Loop Guard should be used on non-designated ports and ensures that these do not transition to designated.

Loop Guard can be activated at the global level with the command `spanning-tree loop guard default`, and the switch enables it on all ports, or at the interface level with the command `spanning-tree guard loop`. If enabled at the global level and you want to disable it on a specific port, you can use the interface-level command `spanning-tree guard none`.

Configuration

```
SW1:

!
spanning-tree loopguard default
```

Verification

Verify that Loop Guard is enabled at the global level:

https://t.me/learningnets

```
SW1#show spanning-tree summary | i Loop
Loopguard Default            is enabled

!
!SW1#show spanning-tree vlan 13 interface fastEthernet1/0/20 detail | i Loop

   Loop guard is enabled by default on the port
```

Simulate that SW2 stops sending BPDUs out on its Fa1/0/20 port by enabling BPDU Filter at the interface level:

```
SW2:

!
interface FastEthernet1/0/20
 spanning-tree bpdufilter enable
```

Verify that Fa1/0/20 on SW1 enters the Loop Inconsistent state, SW1 should issue a log similar to the following:

```
%SPANTREE-2-LOOPGUARD_BLOCK: Loop guard blocking port FastEthernet1/0/20 on VLAN0213.
!
!SW1#show spanning-tree inconsistentports

Name                 Interface               Inconsistency
-------------------- ----------------------- ------------------
VLAN0001             FastEthernet0/22        Loop Inconsistent
VLAN0013             FastEthernet0/22        Loop Inconsistent
VLAN0023             FastEthernet0/22        Loop Inconsistent
VLAN0100             FastEthernet0/22        Loop Inconsistent
VLAN0114             FastEthernet0/22        Loop Inconsistent
VLAN0133             FastEthernet0/22        Loop Inconsistent
VLAN0143             FastEthernet0/22        Loop Inconsistent
VLAN0213             FastEthernet0/22        Loop Inconsistent
Number of inconsistent ports (segments) in the system : 11
```

Remove the test configuration from SW2:

```
SW2:

!
interface FastEthernet1/0/20
```

https://t.me/learningnets

```
no spanning-tree bpdufilter enable
```

After the port is recovered, SW1 should log a message similar to the following:

```
%SPANTREE-2-LOOPGUARD_UNBLOCK: Loop guard unblocking port FastEthernet1/0/20 on VLAN0202.
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing DHCP Attacks Using DHCP Snooping

Task

- Configure R2 as a DHCP server and ASA1's management interface as a DHCP client in VLAN23.
- Ensure that SW1 protects VLAN 23 from rogue DHCP servers attached on any of its current or future ports in VLAN 23.
- SW1 should store the respective database in its flash memory using the file name **VLAN23_DHCP.txt**, and consecutive write operations should not be faster than 30 seconds.
- If ASA1 exceeds 15 DHCP messages per second, the port should be disabled and automatically recovered in 5 minutes.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

DHCP Snooping is a security feature that inspects DHCP packets initiated/destined from/to hosts of an attached VLAN, thus is a switch functionality. With SNopping configured, the switch enforces DHCP packet flow integrity and protects the network from DHCP-based attacks, regardless of its role in the DHCP process: it can be the DHCP Server, DHCP Layer 3 Relay Agent, or have no role in the DHCP end-to-end flow (act like a Layer 2 switch for the VLAN). Before we start detailing the DHCP snooping functionality, let's review the possible device roles from the DHCP process:

- **DHCP Client** requests an IP address as well as other configuration parameters using a **DHCPDISCOVER** packet. When an answer is received from the DHCP server, the client will accept the offer using a **DHCPREQUEST** message. When the client no longer needs DHCP obtained configuration, it instructs the server to make it available for other clients using a **DHCPRELEASE** packet.
- **DHCP Relay** acts as a man-in-the-middle (MiTM) agent between the DHCP clients and the DHCP server and is required when the DHCP server and its clients are not directly attached to the same VLAN. The primary function of a DHCP Relay is to forward messages from directly attached clients to the remote DHCP server. When a DHCP relay receives a DHCP broadcast packet from a connected client, it encapsulates it in a unicast packet towards the DHCP server. In Cisco IOS, you can specify the IP address of the server using the interface-level command `ip helper-address <server-ip-address>` . This enables the DHCP Relay Agent function and is configured on the layer 3 interface facing the DHCP clients. The DHCP Relay Agent changes the **giaddr** field in DHCP packets from the default of zero (an IP address of 0.0.0.0) to the IP address of the interface the original packet was received on. Because in this case the DHCP server is not in the same VLAN with the client, it uses the IP address from **giaddr** field to identify the correct DHCP pool and respond with correct IP addressing information in its **DHCPOFFER** packet.
- **DHCP Server** responds to discover and request messages from DHCP clients with **DHCPOFFER** and **DHCPACK/DHCPNAK** messages. The offer contains the IP address and other settings for the client, whereas the ACK/NACK acknowledge or reject a request message. A server may also explicitly request a client to release an IP address through a **DHCPRELEASEREQUERY** message.

https://t.me/learningnets

DHCP messages contain the MAC address of the client (client hardware address) explicitly stored in one of DHCP fields. This MAC address identifies a client to the server if the Client ID field is set in the DHCP packet. DHCP servers allocate IP addresses based on the client hardware address found in the DHCP payload, not based on the source MAC address of the Ethernet frame (this would not be possible anyway if the server is not in the same VLAN with the clients, because the source Ethernet MAC address is lost).

The primary goal of using DHCP Snooping is to enforce DHCP security. When you enable DHCP Snooping on a switch, it starts treating every port as connected to a DHCP client, also called **untrusted** port. Thus, the switch applies DHCP message filtering on all ports and only accepts messages expected from directly connected DHCP clients (DHCPREQUEST, DHCPDISCOVER, DHCPRELEASE). The switch allows any type of DHCP messages on trusted ports. You may explicitly configure a port as trusted with the interface-level command `ip dhcp snooping trust`. Usually you need this command on the ports connected to DHCP servers, or uplink ports to upstream switches where the DHCP server resides. To be more specific, on trusted ports, DHCP Snooping does not perform any security checks or inspections.

Regardless of whether the port is trusted or not, you can use the interface-level command `ip dhcp snooping limit rate <pps>` to control the number of DHCP packets per second received on a particular port. If the port exceeds the threshold set, the switch will transition the port into an `error-disabled` state. The port can be manually enabled out of this state or automatically through error-disable recovery.

If a DHCP message contains a non-default value for the **giaddr** field, meaning it has already been relayed, it must be received on a trusted port. Untrusted ports simply discard such messages because untrusted ports allow only directly connected DHCP clients. This check can be disabled with the global command `no ip dhcp snooping verify no-relay-agent-address`. This is very helpful when you have DHCP Snooping enabled at the access switch level and you don't want to configure on the distribution switch level the downstream port as trusted, in order to accept the DHCP message that had the **giaddr** field modified by DHCP Snooping.

In addition to filtering messages based on their types, DHCP snooping also populates a special DHCP snooping table based on messages exchanged across untrusted ports. This table contains the IP addresses that were allocated to clients by the DHCP server and other information such as client MAC address, client port, VLAN number, and lease duration. The switch consults this table when it receives DHCP messages such as **DHCPRELEASE** to make sure that it's valid and not spoofed from a malicious host. The switch also ensures that client MAC addresses from DHCP packet payload match the source MAC address from the Ethernet frame of the host sending the message. This feature could be disabled with the global command `no ip dhcp snooping verify mac-address`. Note that messages received on

https://t.me/learningnets

trusted ports do not create any DHCP Snooping binding entries, because packets received on trusted ports are not inspected.

You may instruct the switch to save the DHCP Snooping database contents in flash memory, NVRAM, or on external server through TFTP (it allows the switch to preserve DHCP binding mappings across reloads). The global command to specify the external storage is `ip dhcp snooping database <path>`. You may also want to control the interval between database updates by using the command `ip dhcp snooping database write-delay <seconds>`.

To start the DHCP Snooping process, you must enable it globally and also activate it per-VLAN using the commands `ip dhcp snooping` and `ip dhcp snooping vlan <vlan-number>`. Note that by default, when DHCP Snooping is enabled, the switch also adds DHCP Information Option (or Option 82) to all packets received on untrusted ports, the purpose of this option being to identify the device and port that the DHCP client is connected to; these values can be used by the DHCP server to allocate certain IP addresses and configurations to the requesting clients. For DHCP Snooping to enforce DHCP security, you don't need Option 82 to be inserted. One possible issue that may arise here is if the switch inserts the option but leaves the **giaddr** field at its default value of zero. In such case, a DHCP server may assume that the option has been formatted incorrectly, because a DHCP relay is supposed to set the **giaddr** field to its own IP address. An IOS DHCP server would reject by default such DHCP messages, so the following options to prevent this are available:

- Configure the IOS DHCP server to accept DHCP messages with a zero address in the **giaddr** field by using the global command `ip dhcp relay information trust-all` or the interface-level command `ip dhcp relay information trusted`.
- Configure DHCP Snooping not to insert Option 82. This is accomplished by using the global command `no ip dhcp-snooping information option`.
- Configure the port towards the DHCP client as trusted, because DHCP Snooping does not insert any Information Option for packets received on trusted ports.

DHCP Snooping protects against DHCP attacks in the following ways:

- Prevents rogue DHCP servers from functioning and thus allocating rogue IP address to DHCP clients (through the use of untrusted ports).
- Prevents rogue DHCP clients from depleting the DHCP pool and thus leasing all IP addresses (through the validation between Ethernet MAC address and DHCP Client Hardware address).
- Prevents DHCP client impersonation, which results in IP connectivity issues for a client (through maintenance of DHCP Snooping database):
  - A rogue DHCP client may try to decline a lease received from the DHCP

https://t.me/learningnets

server in the name of the actual DHCP client.

- A rogue DHCP client may try to release a leased IP address from the DHCP server in the name of the actual DHCP client.

Configuration

```
SW1:
ip dhcp snooping
ip dhcp snooping vlan 23
ip dhcp snooping database flash:/VLAN23_DHCP.txt
ip dhcp snooping database write-delay 30
!
interface FastEthernet1/0/12
 ip dhcp snooping limit rate 15
!
interface FastEthernet1/0/2
 ip dhcp snooping trust
!
errdisable recovery cause dhcp-rate-limit
errdisable recovery interval 300
!
interface FastEthernet1/0/12
 switchport mode access
 switchport access vlan 23
 spanning-tree portfast
 no shutdown

ASA1:
interface Management0/0
 nameif VLAN23
 ip address dhcp
 no shutdown

R2:

!
ip dhcp relay information trust-all
!
ip dhcp pool VLAN23
 network 136.1.23.0/24
```

https://t.me/learningnets

Verification

Verify that R2 allocates IP addresses to ASA1:

```
ASA1# show ip address management0/0
System IP Address:
Interface                  Name                   IP address      Subnet mask     Method
Management0/0              VLAN23                 136.1.23.1      255.255.255.0   DHCP
Current IP Address:
Interface                  Name                   IP address      Subnet mask     Method
Management0/0              VLAN23                 136.1.23.1      255.255.255.0   DHCP
!
!R2#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.23.1          0063.6973.636f.2d30.    Apr 22 2013 05:40 PM    Automatic
                    3031.622e.3063.6532.
                    2e34.6531.612d.564c.
                    414e.3233.2d52.6163.
                    6b31.4153.4131.00
!
!SW1#more flash:VLAN23_DHCP.txt
2ba0f221 TYPE DHCP-SNOOPING
VERSION 1
BEGIN 136.1.23.1 23 001b.0ce2.4e1a 2BA2436D Fa0/12                              8019ba9c

END
```

Verify the DHCP Snooping configuration:

```
SW1#show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
23
DHCP snooping is operational on following VLANs:
23
DHCP snooping is configured on the following L3 Interfaces:

Insertion of option 82 is enabled
   circuit-id default format: vlan-mod-port
   remote-id: 0024.1406.4b80 (MAC)
Option 82 on untrusted port is not allowed
```

https://t.me/learningnets

```
Verification of hwaddr field is enabled
Verification of giaddr field is enabled
DHCP snooping trust/rate is configured on the following Interfaces:

Interface                  Trusted    Allow option    Rate limit (pps)
-----------------------    -------    ------------    ----------------
FastEthernet1/0/2            yes         yes             unlimited
  Custom circuit-ids:FastEthernet1/0/12           no          no                15

  Custom circuit-ids:
!
!SW1#show ip dhcp snooping database
Agent URL : flash:/VLAN23_DHCP.txt
Write delay Timer : 30 seconds

Abort Timer : 300 seconds

Agent Running : No
Delay Timer Expiry : Not Running
Abort Timer Expiry : Not Running

Last Succeded Time : 20:27:13 UTC Fri Mar 12 1993
Last Failed Time : None
Last Failed Reason : No failure recorded.

Total Attempts       :        2   Startup Failures :        0
Successful Transfers :        2   Failed Transfers :        0
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        2   Failed Writes    :        0
Media Failures       :        0
```

Verify err-disable port recovery:

```
SW1#show errdisable recovery | i dhcp|interval
dhcp-rate-limit              Enabled
Timer interval: 300 seconds
!
!SW1#show errdisable detect | i dhcp
dhcp-rate-limit              Enabled      port
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing DHCP Attacks Using Port-Security

Task

- Assuming that the previous DHCP Snooping task was completed, ensure that a rogue device connected to SW1's Fa1/0/12 interface cannot starve the DHCP pool.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> For this task to be functional, you must have completed the previous DHCP Snooping task and keep the resulted configuration files or simply copy/paste the configurations from the previous DHCP Snooping task.

Overview

The previous task stated that by default, DHCP Snooping checks that the source MAC address from the Ethernet frame is the same as the Client Hardware address present in the DHCP payload. This ensures that a rogue client cannot use the same Ethernet MAC address with different DHCP payload MAC addresses and thus request multiple IP addresses from the DHCP server, potentially leasing it all. However, what would stop a rogue client from creating random DHCP request packets, where Ethernet MAC address equals DHCP payload MAC address, so basically spoofing its MAC address? For best protection, DHCP Snooping should be coupled with port-security, thus limiting the maximum number of Ethernet MAC addresses that can be learned on a port.

Because we have only one device connected to SW1's Fa1/0/12 interface, we'll leave the default limit of maximum one MAC address.

https://t.me/learningnets

Configuration

```
SW1:

!
interface FastEthernet1/0/12
 switchport port-security
```

Verification

Verify that port-security is enabled on Fa1/0/12 port:

```
SW1#show port-security interface fastEthernet1/0/12
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0Last Source Address:Vlan    : 001b.0ce2.4e1a:23

Security Violation Count   : 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing ARP Attacks Using Dynamic ARP Inspection

Task

- Configure SW1 to protect VLAN 23 from ARP poisoning attacks.
- Ensure that IP connectivity inside VLAN 23 is maintained between R2 and R3.
- Based on your configuration, ensure that all possible logging is enabled, but store no more than 10 log entries in the buffer.
- Configure SW1 to perform all sanity checks for ARP packets.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Dynamic ARP Inspection (DAI) is a security feature that fixes some well-known weaknesses in the ARP protocol. Generally, ARP operation on an Ethernet segment allows any host to spoof a MAC address for any IP address on the segment. These attacks, commonly known as **Man-in-the-Middle (MITM)** attacks, cannot be prevented by using only port-security, access-lists, or other well-known security features. DAI is used to prevent ARP poisoning attacks:

- It is enabled per VLAN using the global command `ip arp inspection vlan <vlan-id>` .
- By default, all ports are untrusted; to configure a port as trusted, use the interface-level command `ip arp inspection trust` .

https://t.me/learningnets

- DAI only inspects ARP packets from untrusted ports.
- ARP packets are validated against the DHCP Snooping database information or against statically configured ARP entries from ARP access-lists.

When the switch receives an ARP packet on an untrusted port, it compares the IP-to-MAC address binding with entries from the DHCP Snooping database or ARP access-lists. If there is no match, the ARP packet is dropped. Note that implementing DAI may break some services, such as Proxy ARP. To resolve these issues, ARP inspection allows you to configure some ports as trusted, because all ARP messages from trusted ports bypass the ARP inspection engine. Most of the time, for example, access layer switch uplinks toward the distribution or core layer are configured as trusted.

By default, DAI inspects all ARP packets from untrusted ports to ensure valid IP-to-MAC address table bindings, but additional inspections can be enabled:

- It can check that target MAC addresses inside the ARP packet payload match the destination MAC address in Ethernet frames for ARP responses. It can be enabled with the global command `ip arp inspection validate dst-mac .`
- It can check that sender MAC addresses inside the ARP packet payload match the source MAC address in Ethernet frames for ARP requests and responses. It can be enabled with the global command `ip arp inspection validate src-mac .`
- It can check the ARP packet payload for invalid or unexpected IP addresses, ensuring that no host tries to bind MAC addresses to IP addresses, such as 0.0.0.0, 255.255.255.255, or multicast addresses. It can be enabled with the global command `ip arp inspection validate ip .`

DAI validates ARP packets against the DHCP Snooping database. If any hosts on the segment are not using DHCP for address allocation, DHCP is not used, or DHCP Snooping is not enabled, you must configure ARP access-lists to allow ARP traffic:

- The access-list is created using the global command `arp access-list <acl-name> .`
- Access-list entries are configured using the command `permit ip host <sender-IP-address> mac host <sender-MAC-address> [log] .`
- The access-list is applied per VLAN using the command `ip arp inspection filter <acl-name> vlan <vlan-id> [static] .`

DAI first checks the ARP access-list to determine whether an ARP packet is legitimate; if there are no matches on **permit** entries for the given IP and MAC address pair and there is no explicit **deny ip any mac any** statement at the end of the access-list, DAI also checks the DHCP Snooping database. If there are no

https://t.me/learningnets

matches, the ARP packet is dropped. However, if there is an explicit deny statement at the end of the access-list or the access-list has been configured with the **static** keyword, ARP inspection does not consult the DHCP Snooping database.

DAI also allows for ARP message rate-limiting. This is enabled by default on untrusted ports with a maximum of 15 pps, and disabled on trusted ports. Using the interface-level command `ip arp inspection limit rate {<pps> [burst interval <seconds>]}` , you can modify the default thresholds. When the port exceeds the configured rate, it will go into **error-disabled** state. The feature limits the aggregate rate on trunks or EtherChannels, so you may need to adjust it for real-world situations.

From the logging perspective, DAI is very granular, with multiple options available:

- You may specify the **log** keyword for entries in the ARP access-list, but logs are not generated by default unless you enable it with the global command `ip arp inspection vlan <vlan-id> logging acl-match matchlog .`
- You may enable logging when entries in the DHCP Snooping database are matched by using the global command `ip arp inspection vlan <vlan-id> logging dhcp-bindings {all|permit} .`
- You may enable logging for suspect ARP packets (such as those having a source IP of 0.0.0.0) with the global command `ip arp inspection vlan <vlan-id> logging arp-probe` .
- The switch accumulates all DAI-related logging in a dedicated internal buffer, and its size can be modified by using the global command `ip arp inspection log-buffer entries <number> .`

Configuration

Note that the MAC addresses of R2 and R3 will be different than those used below, based on your rack number.

SW1:

```
arp access-list DAI_VLAN23
 permit ip host 136.1.23.2 mac host d867.d9e0.bbc0 log
 permit ip host 136.1.23.3 mac host 001e.f779.4771 log
!
ip arp inspection filter DAI_VLAN23 vlan 23
!
ip arp inspection vlan 23
ip arp inspection vlan 23 logging acl-match matchlog
ip arp inspection vlan 23 logging dhcp-bindings all
```

https://t.me/learningnets

```
ip arp inspection vlan 23 logging arp-probe
ip arp inspection log-buffer entries 10
ip arp inspection validate src-mac dst-mac ip
```

Verification

Verify DAI configuration:

```
SW1#show ip arp inspection vlan 23
 Source Mac Validation      : Enabled
Destination Mac Validation : Enabled
IP Address Validation      : Enabled

 Vlan     Configuration    Operation   ACL Match          Static ACL
 ----     -------------    ---------   ---------          ----------
   23     Enabled          Active      DAI_VLAN23         No

 Vlan     ACL Logging      DHCP Logging      Probe Logging
 ----     -----------      ------------      -------------
   23     Acl-Match        All               Permit
!
!SW1#show ip arp inspection log
Total Log Buffer Size : 10
Syslog rate : 5 entries per 1 seconds.

No entries in log buffer.
!
!SW1#show ip arp inspection interfaces fastEthernet1/0/2

 Interface        Trust State     Rate (pps)    Burst Interval
 ---------------  -----------     ----------    -------------- Fa1/0/2Untrusted
               15                 1
!
!SW1#show ip arp inspection interfaces fastEthernet1/0/20

 Interface        Trust State     Rate (pps)    Burst Interval
 ---------------  -----------     ----------    -------------- Fa1/0/20Untrusted
               15                 1
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing Attacks Using IP Source Guard

Task

- Configure SW1's Fa1/0/9 and Fa1/0/10 in VLAN 23:
  - assume connected hosts will receive IP addresses via DHCP from R2.
  - ensure these hosts cannot spoof their layer2/layer3 addresses after receiving IP address form DHCP server.
- Match the following output to maintain IP connectivity between R2 and R3:

```
SW1#show  ip verify source interface fastEthernet1/0/2

Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Fa1/0/2    ip-mac       inactive-trust-port
```

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IP Source Guard is a security feature intended to prevent packet spoofing based on MAC or IP addresses (MiTM attacks). When you enable IP Source Guard on a layer 2 port, the switch applies a layer 3 filter to this port, only accepting the packets with

https://t.me/learningnets

source IP addresses matching DHCP Snooping bindings created for the port, thus enabling DHCP Snooping is a prerequisite. If a host does not use DHCP to obtain IP addressing, all its traffic will be dropped by the swich unless you create a static entry using the global command `ip source binding <MAC-address> vlan <vlan-id> <IP-address> interface <interface-id> .`

IP Source Guard relieves the necessity of applying IP ingress filtering on individual layer 2 ports to prevent IP address spoofing. Combined with Dynamic ARP Inspection and DHCP Snooping, it allows for a much more secure environment than default layer 2 deployments.

You enable IP Source Guard per interface by using the command `ip verify source [port-security]` . The **port-security** option requires port-security to be enabled on the switch port as well. With this option, the switch filters packets based on both the source IP and MAC addresses, and the secure MAC address is taken from the DHCP snooping database or a static mapping entry.

Note that on ports configured as trusted from the DHCP Snooping point of view, configuring IP Source Guard has no effect; it acts as though disabled. This is logical because IP Source Guard relies on the DHCP Snooping database, and because snooping is disabled on that port, DHCP packets are not inspected and no entries can be added in the DHCP Snooping database. For this reason, we don't need create manual IP-to-MAC-to-PORT binding for R2's IP in VLAN 23.

Configuration

https://t.me/learningnets

```
SW1:

interface FastEthernet1/0/2
 ip verify source port-security
 switchport port-security
 ip dhcp snooping trust
!
interface range FastEthernet1/0/9 - 10
 switchport mode access
 switchport access vlan 23
 ip verify source port-security
 switchport port-security
 no shutdown
!
!
ip dhcp snooping
ip dhcp snooping vlan 23
```

Verification

Verify that IP Source Guard is not active without enabling DHCP Snooping:

```
SW1#show ip verify source interface FastEthernet1/0/2
Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----Fa1/0/2    ip-mac
inactive-no-snooping-vlan
```

Enable DHCP snooping. Without configuring Fa1/0/2 as trusted from DHCP Snooping perspective, all inbound traffic will be dropped because no bindings are found in the DHCP Snooping database:

```
SW1#show ip verify source interface FastEthernet1/0/2
Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Fa1/0/2    ip-mac       active deny-all deny-all
           23
!
!SW1#show ip dhcp snooping binding
MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
Total number of bindings: 0
```

https://t.me/learningnets

```
!
!R2#ping 136.1.23.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.3, timeout is 2 seconds:....

Success rate is 0 percent (0/5)
```

After interface Fa1/0/2 is configured as DHCP Snooping trusted, IP Source Guard will appear as inactive, and thus intra-VLAN 23 traffic is not restricted:

```
SW1#show ip verify source interface FastEthernet1/0/2
Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Fa1/0/2       ip-mac        inactive-trust-port
!
!R2#ping 136.1.23.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R3#ping 136.1.23.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Because ports Fa1//0/9 and Fa1/0/10 are not active (no hosts connected), the filtering mode shows up as **inactive-no-snooping-vlan**, however DHCP snooping is active on the VLAN:

```
SW1#show ip verify source

Interface  Filter-type  Filter-mode  IP-address       Mac-address        Vlan
---------  -----------  -----------  ---------------  -----------------  ----
Fa1/0/2    ip-mac       inactive-trust-portFa1/0/9    ip-mac       inactive-no-snooping-vlan
Fa1/0/10   ip-mac       inactive-no-snooping-vlan
!
!SW1#show ip dhcp snooping
 Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
23
DHCP snooping is operational on following VLANs:
23
DHCP snooping is configured on the following L3 Interfaces:
```

https://t.me/learningnets

```
Insertion of option 82 is enabled
   circuit-id default format: vlan-mod-port
   remote-id: 001b.d58b.2180 (MAC)
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Verification of giaddr field is enabled
DHCP snooping trust/rate is configured on the following Interfaces:

Interface                  Trusted    Allow option    Rate limit (pps)
-----------------------    -------    ------------    ----------------
FastEthernet1/0/2          yes        yes             unlimited

  Custom circuit-ids:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing VLAN Hopping Attacks

Task

- Configure SW2's Fa1/0/13 interface in VLAN 133.
- Ensure that frames received on this port cannot initiate a VLAN hopping attack, considering the switch does not accept 802.1q frames on access ports.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

A VLAN hopping attack happens when a frame initiated from one VLAN is able to reach another VLAN without passing through a layer 3 device. There are two possible scenarios in which this attack could occur. Imagine two switches connected via a trunk link with a native VLAN of 10:

- If a rogue host connected to one switch can successfully negotiate trunking via DTP, it can send double tagged frames to the switch: the inner 802.1q header representing VLAN 20 and the outer 802.1q header representing VLAN 10:
  - If the switch decides to forward the frame out on its trunk port to the second switch, because of the native VLAN 10, it will remove the outer 802.1q header and transmit the frame with the inner 802.1q header.
  - The second switch receives the frame with an 802.1q header for VLAN 20,

https://t.me/learningnets

and thus forwards the frame in VLAN20.

- If a rogue host cannot negotiate the trunk port with the switch but the switch accepts tagged frames on access ports, it can send double tagged frames to the switch: the inner 802.1q header representing VLAN 20 and the outer 802.1q header representing VLAN 10.
  - Given the native VLAN 10 between switches, this will lead again to the same attack.

Note that if a switch accepts double or triple tagged frames on trunk ports, switching is always done based on the outer 802.1q header, just like when using Q-in-Q.
There are several solutions for mitigating such attacks:

- Disable DTP on all ports by using the interface-level command `switchport nonegotiate` or `switchport mode access .`
- Ensure that the native VLAN used between switches is a dedicated VLAN, not used on any access ports.
- Enable 802.1q tagging for all VLANs, which basically disables the native VLAN functionality by using the global command `vlan dot1q tag native .`

Given the task requirements, we need to protect ourselves from the first scenario, in which the host can negotiate a trunk.

Configuration

SW2:

```
interface FastEthernet1/0/13
 switchport mode access
 switchport access vlan 133
 no shutdown
```

Verification

Verify that DTP is disabled on the port:

https://t.me/learningnets

```
SW2#show interfaces fastEthernet1/0/13 switchport | i Trunking

Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: nativeNegotiation of Trunking: Off

Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Implementing RFC 1918 Anti-Spoofing Filtering

Task

- You have been notified that your network is under attack from the Internet. Configure R3 so that it drops and logs each packet sourced from RFC 1918 IPv4 address space, inbound on its connection to R2.
- Ensure that by viewing a log message, you can identify the access-list that was matched, based on a router-generated tag.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

RFC 1918 specifies that a range from each IPv4 unicast class of addresses is reserved for private/internal usage and should not be used to represent Internet/publicly available resources. These are: 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. Because these addresses are not publicly routable on the Internet, there is no valid reason for traffic to be received on an Internet-facing interface that is sourced from a host in these ranges. Therefore, best security practice dictates that these addresses should be denied entrance to a network on the Internet-facing connection. Because we need to configure filtering based on source address, both standard and extended access-lists could be used; however, because of the added functionality, we'll use extended.

https://t.me/learningnets

The **log** keyword is required to generate log messages for dropped packets, but we need to log a message each and every time an access-list entry is matched. By default, if the **log** option is configured, messages are generated for only the first matching packet; after this, messages are cumulated for a five-minute interval and a single message is logged with the total number of matched packets during that interval. To generate a log message for each match, we need to modify the default threshold using the global command **`ip access-list log-update threshold <message-threshold> .`**

IOS can generate and attach an MD5 hash to each access-list entry, which survives on a router reload. This hash value will be appended to each generated log message and can help to easily filter messages on a syslog server based on a specific access-list entry.

Configuration

```
R3:

ip access-list extended RFC1918
 deny ip 10.0.0.0 0.255.255.255 any log
 deny ip 172.16.0.0 0.15.255.255 any log
 deny ip 192.168.0.0 0.0.255.255 any log
 permit ip any any
!
ip access-list log-update threshold 1
ip access-list logging hash-generation
!
interface FastEthernet0/1.23
 ip access-group RFC1918 in
```

Verification

First we must verify that the access-list is applied correctly inbound on R3's interface:

```
R3#show ip interface FastEthernet0/1.23 | i RFC1918
 Inbound  access list is RFC1918
!
!R3#show ip access-lists RFC1918

Extended IP access list RFC1918
    10 deny ip 10.0.0.0 0.255.255.255 any log (hash = 0xE2F226F6)
    20 deny ip 172.16.0.0 0.15.255.255 any log (hash = 0xC65E057E)
    30 deny ip 192.168.0.0 0.0.255.255 any log (hash = 0x1015B05A)
```

https://t.me/learningnets

```
    40 permit ip any any
```

We must also ensure that each match of an access-list entry is logged and the MD5 hash is also appended to the log. For this, we generate packets from the Internet, sourced from RFC1918 IPv4 address ranges:

```
R2#ping 136.1.23.3



Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
!
!R2#ping 136.1.23.3 source loopback10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.3, timeout is 2 seconds:
Packet sent with a source address of 10.10.10.2 U.U.U
Success rate is 0 percent (0/5)
!
!R3#show ip access-lists RFC1918
Extended IP access list RFC1918
10 deny ip 10.0.0.0 0.255.255.255 any log (5 matches) (hash = 0xE2F226F6)
    20 deny ip 172.16.0.0 0.15.255.255 any log (hash = 0xC65E057E)
    30 deny ip 192.168.0.0 0.0.255.255 any log (hash = 0x1015B05A)
    40 permit ip any any (5 matches)
!
!R3#show logging | i RFC1918
%SEC-6-IPACCESSLOGDP: list RFC1918 denied icmp 10.10.10.2 -> 136.1.23.3 (0/0), 1 packet  [0xE2F226F6]
%SEC-6-IPACCESSLOGDP: list RFC1918 denied icmp 10.10.10.2 -> 136.1.23.3 (0/0), 1 packet  [0xE2F226F6
]%SEC-6-IPACCESSLOGDP: list RFC1918 denied icmp 10.10.10.2 -> 136.1.23.3 (0/0), 1 packet  [0xE2F226F6
]%SEC-6-IPACCESSLOGDP: list RFC1918 denied icmp 10.10.10.2 -> 136.1.23.3 (0/0), 1 packet  [0xE2F226F6
]%SEC-6-IPACCESSLOGDP: list RFC1918 denied icmp 10.10.10.2 -> 136.1.23.3 (0/0), 1 packet  [0xE2F226F6
]
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Implementing RFC 2827 Anti-Spoofing Filtering

Task

- As you own the network of 136.1.0.0/8, ensure bidirectional RFC2827 filtering on R3's connection to ASA4.
- Each security policy violation should be logged.
- Configure the router so that when viewing a log message, you can identify the name and line of the access-list from which it was generated.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

RFC 2827 defines a filtering mechanism based on your internally assigned IPv4 address space, to ensure a basic form of security. Traffic leaving your network should have source IPv4 addresses only from your address space, and traffic entering your network should have any source IPv4 addresses except the ones you own. Usually these filters are applied at the edge of the network, so that it doesn't affect your internal network traffic but still drops inbound spoofed traffic as soon as possible.

Syslog messages can be tagged using router random-generated MD5 hash values, or user-defined cookies. The advantage of cookies is that you can specify a more meaningful tag value, such as matching the access-list name and the access-list

https://t.me/learningnets

entry to help you on troubleshooting scenarios. The user-defined cookie is configured as a tag value at the end pf each access-list entry.

Note that user-defined cookies take precedence over router MD5 hash-generated values; if both options are configured, there will be no MD5 hash-generated value for access-lists with cookies.

Configuration

R3:

```
ip access-list extended RFC2827_IN
 deny ip 136.1.0.0 0.255.255.255 any log RFC2827_IN_LINE1
 permit ip any any
!
ip access-list extended RFC2827_OUT
 permit ip 136.1.0.0 0.255.255.255 any
 deny ip any any log RFC2827_OUT_LINE2
!
interface FastEthernet0/0.143
 ip access-group RFC2827_IN in
 ip access-group RFC2827_OUT out
!
ip access-list log-update threshold 1
```

Verification

Verify that the access-list is applied correctly inbound/outbound on R3's interface:

https://t.me/learningnets

```
R3#show ip interface fastEthernet0/0.143 | i RFC
    Outgoing access list is RFC2827_OUT
 Inbound  access list is RFC2827_IN

!
!R3#show ip access-lists RFC2827_IN

Extended IP access list RFC2827_IN
    10 deny ip 136.1.0.0 0.255.255.255 any log (tag = RFC2827_IN_LINE1)
    20 permit ip any any (661 matches)
!
!R3#show ip access-lists RFC2827_OUT

Extended IP access list RFC2827_OUT
    10 permit ip 136.1.0.0 0.255.255.255 any
    20 deny ip any any log (tag = RFC2827_OUT_LINE2)
```

Ensure that nonconforming packets are logged and messages are tagged using the configured cookie. For this, we generate packets from the SW1, sourced from the internal 136.0.0.0/8 IPv4 address range:

```
SW1#ping 172.1.143.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.1.143.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
!
!SW1#ping 172.1.143.3 source loopback1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.1.143.3, timeout is 2 seconds:
Packet sent with a source address of 136.7.7.7 .....
Success rate is 0 percent (0/5)
!
!R3#show ip access-lists RFC2827_IN
Extended IP access list RFC2827_IN
10 deny ip 136.0.0.0 0.255.255.255 any log (5 matches) (tag = RFC2827_IN_LINE1)
    20 permit ip any any (14 matches)
!
!R3#show logging | i RFC2827_IN
%SEC-6-IPACCESSLOGDP: list RFC2827_IN denied icmp 136.7.7.7 -> 172.1.143.3 (0/0), 1 packet  [RFC2827_IN_LINE1]
%SEC-6-IPACCESSLOGDP: list RFC2827_IN denied icmp 136.7.7.7 -> 172.1.143.3 (0/0), 1 packet
[RFC2827_IN_LINE1]
```

https://t.me/learningnets

```
%SEC-6-IPACCESSLOGDP: list RFC2827_IN denied icmp 136.7.7.7 -> 172.1.143.3 (0/0), 1 packet
[RFC2827_IN_LINE1]
%SEC-6-IPACCESSLOGDP: list RFC2827_IN denied icmp 136.7.7.7 -> 172.1.143.3 (0/0), 1 packet
[RFC2827_IN_LINE1]
%SEC-6-IPACCESSLOGDP: list RFC2827_IN denied icmp 136.7.7.7 -> 172.1.143.3 (0/0), 1 packet
[RFC2827_IN_LINE1]
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Implementing RFC 3330 Anti-Spoofing Filtering

Task

- Configure inbound filtering in R2's connection to R3, in accordance with RFC 3330.
- Configure R2 for detailed logging information of each policy violation.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IANA is responsible for allocating IPv4 address space on the Internet. By definition, **bogon networks** are IPv4 address ranges that should not be seen on the Internet, because they are either reserved or unallocated, or have a special use. Because these ranges are not Internet routable, inbound filtering on the Internet edge helps mitigate attacks based on IPv4 address spoofing. www.team-cymru.com maintains an updated list of **bogon networks**, and it also provides a free option of BGP peering to get the updated list in real time.

However, RFC 3330 defined the static list of bogon addresses back in 2002. Although RFC 5735 has made RFC 3330 obsolete, it still may show up on the exam. The IPv4 address range defined by RFC 3330 is follows (as you can see, it includes RFC 1918):

```
0.0.0.0/8
10.0.0.0/8
14.0.0.0/8
```

https://t.me/learningnets

```
24.0.0.0/8
39.0.0.0/8
127.0.0.0/8
128.0.0.0/16
169.254.0.0/16
172.16.0.0/12
191.255.0.0/16
192.0.0.0/24
192.0.2.0/24
192.88.99.0/24
192.168.0.0/16
198.18.0.0/15
223.255.255.0/24
224.0.0.0/4
240.0.0.0/4
```

The IPv4 address range defined by RFC 5735 is as follows (as you can, see it includes RFC 1918):

```
0.0.0.0/8
10.0.0.0/8
127.0.0.0/8
169.254.0.0/16
172.16.0.0/12
192.0.0.0/24
192.0.2.0/24
192.88.99.0/24
192.168.0.0/16
198.18.0.0/15
198.51.100.0/24
203.0.113.0/24
224.0.0.0/4
240.0.0.0/4
255.255.255.255/32
```

Configuration

R2:

```
ip access-list extended RFC3330
 deny ip 0.0.0.0 0.255.255.255 any log-input
 deny ip 10.0.0.0 0.255.255.255 any log-input
 deny ip 14.0.0.0 0.255.255.255 any log-input
```

https://t.me/learningnets

```
 deny ip 24.0.0.0 0.255.255.255 any log-input
 deny ip 39.0.0.0 0.255.255.255 any log-input
 deny ip 127.0.0.0 0.255.255.255 any log-input
 deny ip 128.0.0.0 0.0.255.255 any log-input
 deny ip 169.254.0.0 0.0.255.255 any log-input
 deny ip 172.16.0.0 0.15.255.255 any log-input
 deny ip 191.255.0.0 0.0.255.255 any log-input
 deny ip 192.0.0.0 0.0.0.255 any log-input
 deny ip 192.0.2.0 0.0.0.255 any log-input
 deny ip 192.88.99.0 0.0.0.255 any log-input
 deny ip 192.168.0.0 0.0.255.255 any log-input
 deny ip 198.18.0.0 0.1.255.255 any log-input
 deny ip 223.255.255.0 0.0.0.255 any log-input
 deny ip 224.0.0.0 15.255.255.255 any log-input
 deny ip 240.0.0.0 15.255.255.255 any log-input
 permit ip any any
!
ip access-list log-update threshold 1
!
interface GigabitEthernet0/0
 ip access-group RFC3330 in
```

Verification

Verify that the access-list is applied correctly inbound on R2's interface:

```
R2#show ip interface gigabitEthernet 0/0 | i RFC3330
 Inbound  access list is RFC3330
!
!R2#show ip access-lists RFC3330

Extended IP access list RFC3330
    10 deny ip 0.0.0.0 0.255.255.255 any log-input
    20 deny ip 10.0.0.0 0.255.255.255 any log-input
    30 deny ip 14.0.0.0 0.255.255.255 any log-input
    40 deny ip 24.0.0.0 0.255.255.255 any log-input
    50 deny ip 39.0.0.0 0.255.255.255 any log-input
    60 deny ip 127.0.0.0 0.255.255.255 any log-input
    70 deny ip 128.0.0.0 0.0.255.255 any log-input
    80 deny ip 169.254.0.0 0.0.255.255 any log-input
    90 deny ip 172.16.0.0 0.15.255.255 any log-input
    100 deny ip 191.255.0.0 0.0.255.255 any log-input
    110 deny ip 192.0.0.0 0.0.0.255 any log-input
    120 deny ip 192.0.2.0 0.0.0.255 any log-input
    130 deny ip 192.88.99.0 0.0.0.255 any log-input
```

https://t.me/learningnets

```
    140 deny ip 192.168.0.0 0.0.255.255 any log-input
    150 deny ip 198.18.0.0 0.1.255.255 any log-input
    160 deny ip 223.255.255.0 0.0.0.255 any log-input
    170 deny ip 224.0.0.0 15.255.255.255 any log-input
    180 deny ip 240.0.0.0 15.255.255.255 any log-input
    190 permit ip any any (22 matches)
```

Ensure that detailed logging information is generated for each policy violation. For this, we generate packets from R1, sourced from one of the RFC 3330-defined IPv4 address ranges:

```
R1#ping 136.1.23.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R1#ping 136.1.23.2 source loopback10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:
Packet sent with a source address of 10.10.10.1 U.U.U
Success rate is 0 percent (0/5)
!
!R2#show logging | i RFC3330
%SEC-6-IPACCESSLOGDP: list RFC3330 denied icmp 10.10.10.1 (GigabitEthernet0/0 001e.f779.4771) -> 136.1.23.2 (0/0), 1
  %SEC-6-IPACCESSLOGDP: list RFC3330 denied icmp 10.10.10.1(GigabitEthernet0/0 001e.f779.4771)
 -> 136.1.23.2 (0/0), 1 packet  %SEC-6-IPACCESSLOGDP: list RFC3330 denied icmp 10.10.10.1
(GigabitEthernet0/0 001e.f779.4771)
 -> 136.1.23.2 (0/0), 1 packet  %SEC-6-IPACCESSLOGDP: list RFC3330 denied icmp 10.10.10.1
(GigabitEthernet0/0 001e.f779.4771)
 -> 136.1.23.2 (0/0), 1 packet  %SEC-6-IPACCESSLOGDP: list RFC3330 denied icmp 10.10.10.1
(GigabitEthernet0/0 001e.f779.4771)
 -> 136.1.23.2 (0/0), 1 packet
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IOS TCP Intercept

Task

- Configure R3 to protect the servers in VLAN 13 from SYN flooding attacks. Intercept connections as being initiated and close connections after 30 seconds of inactivity.
- Allow for a maximum of 100 half-open connections and drop to 80 when the threshold is crossed.
- The router should start dropping half-open connections if it reaches 60 sessions within a minute and stop when session rate is 30 within a minute.
- The router should not account for connection age when dropping it.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The TCP Intercept feature protects against TCP SYN flood attacks targeted at servers/hosts located behind the router. SYN flooding involves sending a huge number of SYN packets toward the server's host without ever finishing the TCP three-way handshake. The server exhausts its resources and may eventually refuse to accept a connection from a legitimate user. There are two functional modes for the TCP Intercept feature:

https://t.me/learningnets

- **Intercept mode**
- **Watch mode**

In the intercept mode, the router permanently acts as a proxy between client and server. For every incoming SYN packet, the router immediately answers with a SYN-ACK packet to the client in the server's name, and also sends the SYN packet to the server. The router then waits for a final ACK from the client and a SYN-ACK from the server. When both of these are received, the router makes the join and leaves the proxy mode for the session, just performing routing. If one of the endpoints does not respond, the router probes it repeatedly using an exponential backoff algorithm. The router retransmits a second SYN packet in 1 second, a third in 2 seconds, a fourth in 4 seconds, a fifth retransmit in 8 seconds, and the sixth retransmit in 16 seconds. Altogether, this takes 1+2+4+8+16=31 seconds. If no answer is received, it sends an RST and terminates the session. You can enable the TCP intercept mode with the global command `ip tcp intercept mode intercept .`

In this mode, the router limits the total number of half-open (three-way handshake unfinished) connections from clients. The limit is set with the global command `ip tcp intercept max-incomplete {low|high} <threshold> .` This limit only applies to the half-open connections and does not affect established connections. When the total number of half-open sessions reaches the high threshold, the router starts dropping the half-open connections until their number falls to the low threshold.

You may also limit the rate of incoming SYN-packets by using a per-minute value. Specifically, you can set the maximum number of half-open sessions allowed per minute using the global command `ip tcp intercept one-minute {low|high} <threshold> .` When the number of half-open sessions opened during the last minute exceeds the high threshold, the router starts dropping it until the measured rate drops below the low threshold. The router keeps track of both thresholds at the same time, and drops the connections until their number reaches the low threshold. Depending on the drop-mode, the router may drop either the oldest half-open session or any random session. The command to control this behavior is `ip tcp intercept drop-mode {random|oldest} .`

When high threshold values are reached, in addition to deleting the excessive connection states, the router also changes the backoff algorithm behavior. It reduces all retransmit timers by half, so the additional probes are sent in 0.5 seconds, 1 second, 2 seconds, 4 seconds, and 8 seconds. This results in 15 seconds of probing for a non-answering endpoint. Altogether, dropping the connections and reducing the backoff timers is called **aggressive** mode behavior.

Some other configuration parameters of TCP Intercept mode are:

- **FIN/RST timeout:** This is the amount of time for which the router keeps a connection

https://t.me/learningnets

entry after it receives a FIN packet from a client or server. After this interval, the router sends forceful RST packets to the server (causing it to free the resources) and removes the entry. The command to set this timeout is `ip tcp intercept finrst-timeout <seconds>` .

- **Scope of TCP interception:** You can limit this scope by using the command `ip tcp intercept list [<acl-number|acl-name>]` . Only the connections matching the access-list are subject to TCP interception. By default, TCP Intercept does not apply to any transit connection, unless an access-list for traffic to be intercepted is configured. Configuring the router to intercept all TCP connections will eventually starve router resources and all traffic will be affected. Because we intercept TCP sessions, our access-list can contain IPv4-based statements, or TCP-based statements where we can specify which services are subject to interception based on port numbers. In our configuration, we use IPv4-based statements, so any TCP traffic toward VLAN 13 is subject to TCP intercept.
- **Timeout for inactive connections:** This defines the amount of time that TCP Intercept maintains an inactive connection and is configured with command `ip tcp intercept connection-timeout <seconds>` . After the timeout expires, the router sends RST packets to both sides of the connection, effectively terminating it.

Configuration

R3:

```
ip access-list extended TCP_INTERCEPT
 permit ip any 136.1.13.0 0.0.0.255
!
ip tcp intercept list TCP_INTERCEPT
!
ip tcp intercept drop-mode random
ip tcp intercept connection-timeout 30
ip tcp intercept one-minute low 30 high 60
ip tcp intercept max-incomplete low 80 high 100
```

Verification

We enable TCP Intercept debug on R3 and initiate a telnet session from R2 to R1's VLAN 13 interface, which will trigger the TCP intercept function:

R3#debug ip tcp intercept

```
TCP intercept debugging is on
```

https://t.me/learningnets

```
!
!R2#telnet 136.1.13.1
Trying 136.1.13.1 ... Open


User Access Verification


Password:
R1>
!
!INTERCEPT: new connection (136.1.23.2:41667 SYN -> 136.1.13.1:23)

INTERCEPT(*): (136.1.23.2:41667 <- ACK+SYN 136.1.13.1:23)
INTERCEPT: 1st half of connection is established (136.1.23.2:41667 ACK -> 136.1.13.1:23)
INTERCEPT(*): (136.1.23.2:41667 SYN -> 136.1.13.1:23)
INTERCEPT: 2nd half of connection established  (136.1.23.2:41667 <- ACK+SYN 136.1.13.1:23)
INTERCEPT(*): (136.1.23.2:41667 ACK -> 136.1.13.1:23)
INTERCEPT(*): (136.1.23.2:41667 <- WINDOW 136.1.13.1:23)
```

Verify that the session appears in the TCP intercept state table, and that the mode is intercept, not watch:

```
R3#show tcp intercept connections
Incomplete:
Client                Server                State    Create   Timeout  Mode
!
!
Established:
Client                Server                State    Create   Timeout  Mode
136.1.23.2:65495      136.1.13.1:23         ESTAB    00:00:18 00:00:12 I
```

First, the router adds the connection to the intercept list and ensures that both parts of the connections (client <-> router and router <-> server) have been established. After that, the router counts the time of inactivity. After 30 seconds, the router terminates the connection by sending RST packets to both the client and the server:

```
R3:
INTERCEPT: ESTAB timing out (136.1.23.2:38054 <-> 136.1.13.1:23)

INTERCEPT(*): (136.1.23.2:38054 <- RST 136.1.13.1:23)
INTERCEPT(*): (136.1.23.2:38054 RST -> 136.1.13.1:23)
```

Now we block SYN-ACK packets going from the server to the client, effectively

https://t.me/learningnets

preventing connection to reach an established state:

```
R3:

ip access-list extended NO_ACK
 deny tcp any any established
 permit ip any any
!
interface FastEthernet0/1.13
 ip access-group NO_ACK in
```

Let's initiate a TCP connection from R2 to R1. It seems to open from R2's perspective, because R3 proxies the connection. However, the SYN-ACK from R1 is blocked by R3. After this, R3 tries to re-send the SYN to R1 five times and gives up after 31 seconds. After that, it sends the RST and terminates the connection:

```
R2#telnet 136.1.13.1
Trying 136.1.13.1 ... Open
!
!
R3#INTERCEPT: new connection (136.1.23.2:27065 SYN -> 136.1.13.1:23)
*INTERCEPT(*): (136.1.23.2:27065 <- ACK+SYN 136.1.13.1:23)
*INTERCEPT: 1st half of connection is established (136.1.23.2:27065 ACK -> 136.1.13.1:23)
*INTERCEPT(*): (136.1.23.2:27065 SYN -> 136.1.13.1:23)
*INTERCEPT(*): SYNSENT retransmit 1 (136.1.23.2:27065 SYN -> 136.1.13.1:23)
R3#
*INTERCEPT: client packet dropped in SYNSENT (136.1.23.2:27065 -> 136.1.13.1:23)
*INTERCEPT(*): SYNSENT retransmit 2 (136.1.23.2:27065 SYN -> 136.1.13.1:23)
R3#
*INTERCEPT(*): SYNSENT retransmit 3 (136.1.23.2:27065 SYN -> 136.1.13.1:23)
*INTERCEPT: client packet dropped in SYNSENT (136.1.23.2:27065 -> 136.1.13.1:23)
R3#
*INTERCEPT(*): SYNSENT retransmit 4 (136.1.23.2:27065 SYN -> 136.1.13.1:23)
R3#
*INTERCEPT: client packet dropped in SYNSENT (136.1.23.2:27065 -> 136.1.13.1:23)
R3#INTERCEPT: SYNSENT retransmitting too long (136.1.23.2:27065 <-> 136.1.13.1:23)
INTERCEPT(*): (136.1.23.2:27065 <- RST 136.1.13.1:23)
```

Remove the configuration used for testing purposes:

```
R3:

interface FastEthernet0/1.13
```

https://t.me/learningnets

```
 no ip access-group NO_ACK in
!
no ip access-list extended NO_ACK
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IOS TCP Intercept Watch Mode

Task

- Configure R3 to reset any TCP connection toward VLAN 13 resources that does not reach the established state within 20 seconds.
- R3 should not be actively involved in the session establishment process.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

TCP Intercept Watch mode is a lightweight version of the full TCP Intercept mode. Instead of acting like a TCP proxy and sitting in the path of TCP connections, the router just watches them as the three-way handshake is established. The watch mode is based on the **watch timeout** value. If the client and server do not successfully establish a connection during this time, the router forcefully injects RST messages to the server, causing it to free up the allocated resources. Note that this is in contrast to the Intercept mode, in which the router attempts to communicate with the client and the server by itself, using an exponential backoff algorithm. The global command to configure the watch timeout is `ip tcp intercept watch-timeout <seconds>`.

The watch timeout value only makes sense when you configure the router to function in this mode using the global command `ip tcp intercept mode watch`. Watch mode uses the same timers as the Intercept mode. The half-open connection

https://t.me/learningnets

thresholds are treated in the same way, switching to aggressive mode when their number reaches the high threshold. When in aggressive mode, the configured watch timeout is reduced by half as well. There are no retransmission attempts, because the router is not acting as a proxy for the sessions.

Configuration

```
R3:

ip access-list extended TCP_INTERCEPT
 permit ip any 136.1.13.0 0.0.0.255
!
!
ip tcp intercept list TCP_INTERCEPT
ip tcp intercept mode watch
ip tcp intercept watch-timeout 20
```

Verification

We first block SYN-ACK packets going from the server to the client, effectively preventing the established connection formation:

```
R3:

ip access-list extended NO_ACK
 deny tcp any any established
 permit ip any any
!
interface FastEthernet0/1.13
 ip access-group NO_ACK in
```

We enable TCP Intercept debug on R3:

```
R3#debug ip tcp intercept
TCP intercept debugging is on
```

Try opening a connection to R1 from R2. This time it does not seem to open from R2's perspective, because R3 no longer acts as a proxy for the connection. The SYN-ACK from R1 is blocked by R3. After this, R2 tries to re-send the SYN to R1:

https://t.me/learningnets

the first SYN in 2 seconds, the second SYN in 4 seconds, the third SYN in 8 seconds. The next SYN would be re-sent in 16 seconds, but the watch-timeout of 20 seconds is elapsed and R3 sends an RST packet terminating the session. Note that the default SYN timeout for router-initiated sessions in Cisco IOS is 30 seconds, configurable with the command **`ip tcp synwait-time <seconds>`** ; so, even if the watch-timeout had been, for example, 40 seconds on R3, the SYN timeout on R2 would be reached and R2 would give up trying to establish the session.

```
R3#INTERCEPT: new connection (136.1.23.2:17217 SYN -> 136.1.13.1:23)
R3#
INTERCEPT: client packet passed in SYNSENT (136.1.23.2:17217 -> 136.1.13.1:23)
R3#
INTERCEPT: client packet passed in SYNSENT (136.1.23.2:17217 -> 136.1.13.1:23)
R3#
INTERCEPT: client packet passed in SYNSENT (136.1.23.2:17217 -> 136.1.13.1:23)
R3#INTERCEPT: SYNSENT timing out (136.1.23.2:17217 <-> 136.1.13.1:23)
INTERCEPT(*): (136.1.23.2:17217 RST -> 136.1.13.1:23)
```

Remove the configuration used for testing:

R3:

```
interface FastEthernet0/1.13
 no ip access-group NO_ACK in
!
no ip access-list extended NO_ACK
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

ASA TCP Intercept

Task

- Configure ASA3 to proxy all TCP sessions traversing the firewall if the total amount of embryonic connections is above 15, or above 1 per client.
- Ensure that telnet sessions for the **INSIDE** interface of ASA3 are TCP intercepted only if the number of embryonic connections is above 10.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

TCP Intercept on the ASA can be configured for both transit traffic (data plane) and non-transit traffic (management plane). By default, TCP Intercept is enabled for all TCP sessions terminated on the ASA, such as telnet, SSH, and HTTPS, regardless of the number of embryonic connections; basically all TCP sessions destined to the ASA are TCP intercepted. For all TCP sessions traversing the firewall, TCP Intercept is disabled by default.

https://t.me/learningnets

When TCP Intercept mode is configured on routers, the router is responsible for maintaining each TCP session state, completing the three-way handshake with the client and bridging the TCP connection back to the server when the three-way handshake is complete. With this functionality, we basically move the SYN flooding attack away from the server/host to the router, which may actually have fewer CPU and memory resources than the server/host to keep up with the attack.

The TCP Intercept function on the firewall works a bit different than on routers. The ASA does not just intercept the TCP session; instead it uses a SYN cookie mechanism that has the role of protecting the ASA from DOS attacks based on SYN flooding. ASA receives the SYN packet, creates a cookie based on selective TCP/IP header information and a password known only by itself, and uses this cookie as a sequence number for its SYN-ACK response. Then it simply forgets the sent SYN-ACK and therefore does not consume resources to maintain it. A legitimate client will send the ACK using the sequence number from the SYN-ACK of the ASA plus one, at which point the ASA checks that it created the cookie for this session using the same algorithm. If the result is positive, it proxies the TCP session to the server and creates an entry in the connection table.

TCP Intercept is configured using MPF, where relevant traffic is identified using a class-map. Within the class-map you can match on specific TCP ports, on access-lists, or using the keyword **any**. The `match default-inspection-traffic` is not an option, because only the action type of **inspect** can be configured for it; using this class for the scope of TCP intercept will result in the following error:

```
ERROR: Only 'inspect' action is allowed for the class with 'match default-inspection-traffic'
```

TCP intercept thresholds are specified using a policy-map, and there are two configurable values:

- The number of embryonic connections after which TCP intercept is activated for any upcoming TCP session, configured with the command `set connection embryonic-conn-max <threshold-value> .`
- The number of embryonic connections per client after which TCP intercept is activated for any upcoming TCP sessions of that specific client, configured with the command `set connection per-client-embryonic-max <threshold-value> .` The per-client setting is not available for management traffic destined to the ASA.

Configuration

```
ASA3:
```

https://t.me/learningnets

```
class-map type management TELNET
 match port tcp eq telnet
!
policy-map TELNET
 class TELNET
  set connection embryonic-conn-max 10
!
service-policy TELNET interface INSIDE
!
class-map TCP_INTERCEPT
 match any
!
policy-map global_policy
  class TCP_INTERCEPT
   set connection embryonic-conn-max 15 per-client-embryonic-max 1
```

Verification

Initiate a telnet session to ASA3 and R2 from SW2, but apply an access-list inbound on SW2's VLAN 213 interface to drop SYN-ACK packets, thus ensuring that the three-way handshake does not complete:

```
SW2:

ip access-list extended NO_ACK
 deny tcp any any established
 permit ip any any
!
interface vlan 213
 ip access-group NO_ACK in
```

Initiate the telnet session and verify statistics on ASA3:

```
SW2#telnet 136.1.213.13
Trying 136.1.213.13 ... % Connection timed out; remote host not responding
!
!ASA3#show service-policy interface INSIDE set connection detail

Interface INSIDE:
  Service-policy: TELNET
    Class-map: TELNET
      Set connection policy: embryonic-conn-max 10 current embryonic conns 1, drop 0
!
!       SW2#telnet 136.1.23.2
```

https://t.me/learningnets

```
Trying 136.1.23.2 ... % Connection timed out; remote host not responding

!
!ASA3#show service-policy global set connection detail

Global policy:
  Service-policy: global_policy
    Class-map: TCP_INTERCEPT
      Set connection policy: embryonic-conn-max 15 per-client-embryonic-max 1
        current embryonic conns 1, drop 0
        Per client                Embryonic     Total INSIDE     136.1.213.8     1               1
```

Some TCP intercept statistics can be visualized by enabling advanced threat detection for this specific feature on the ASA, using the global command **`threat-detection statistics tcp-intercept`** :

```
ASA3#show threat-detection statistics
Top          Name    Id    Average(eps)    Current(eps) Trigger      Total events
Top 10 protected servers under attack (sorted by average rate)
Monitoring window size: 30 mins    Sampling interval: 30 secs
<Rank> <Server IP:Port> <Interface> <Ave Rate> <Cur Rate> <Total> <Source IP (Last Attack Time)>
--------------------------------------------------------------------------------
1     136.1.23.2:23 INSIDE 0 0 1 136.1.213.8 (7 secs ago)
2     136.1.213.13:23 INSIDE 0 0 1 136.1.213.8 (1 mins ago)
```

Remove the configuration used for testing:

```
SW2:

interface vlan 213
 no ip access-group NO_ACK in
!
no ip access-list extended NO_ACK
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IOS Advanced FPM

Task

- Configure R3 to filter ICMP Echo packets with the string **AAAA** in the payload. Look no deeper than 256 bytes in the packet.
- Ensure that the filtering applies only to ICMP/IP packets from Ethernet frames on its VLAN 13 interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Flexible Packet Matching (FPM) is a feature that allows for granular packet inspection in Cisco IOS routers. Using FPM, you can match any string, byte, or even bit at any position in the IPv4 packet (or, theoretically, non-IP packet). This may greatly help in identifying and blocking network attacks of known patterns found in malicious traffic. It is basically a smart ACL. FPM has some limitations, however:

- It is completely stateless because it does not track the state/history of the packet flow. Thus, FPM cannot discover dynamic protocol ports negotiated by protocols such as H.323 or FTP, nor can it detect patterns split across multiple packets. Essentially, you are restricted to applying inspection on a per-packet basi; you cannot match across packets.
- You cannot apply FPM for control-plane traffic, because FPM is implemented purely

https://t.me/learningnets

in the CEF switching layer, thus CEF needs to be enabled.
- Fragmented traffic is not assembled for matching, and the only inspected packet is the initial fragment of the IP packet flow.
- IP packets with IP options are not matched by FPM, because they are punted to the route processor and therefore not CEF switched.
- It inspects only unicast IPv4 packets and does not apply to MPLS encapsulated packets.

Configuring an FPM filter consists of the following steps:

1. **Load protocol headers.**
2. **Define a protocol stack.**
3. **Define a traffic filter.**
4. **Apply the policy and verify.**

The following is a deep-dive into each of the FPM components:

Load a PHDF (optional). PHDF stands for Packet Header Definition File. These files use XML syntax and define the structure of various packet headers, such as Ethernet, IP, TCP, UDP, ICMP, and so on. They are very helpful in making filtering more structured; with PHDFs loaded, you may filter based on the header field names and their values, instead of matching fixed offsets in the unstructured packet body. You may load the files into the router's memory using the command `load protocol <path>` . The PDHF files could be manually created using the simple XML syntax or downloaded from CCO. Since IOS version 12.4(15)T, five basic PHDFs are included in the IOS code and located at the virtual path `system:fpm/phdf` . You could load the files directly from there. Defining a custom PHDF requires understanding of protocol header formats and field values along with basic XML formatting, which is beyond the scope of this document and CCIE lab exam.

Define a protocol stack (optional). This step uses the PDHFs loaded previously and allows specifying the protocol headers for the traffic you want to inspect. The protocol stack definition induces structure in the packets being inspected. This allows for filtering based on header field values and specifying offsets in the packet relative to the header fields. Additionally, you may define various protocol stacks (such as UDP in IP, UDP in GRE in IP) and reuse the same access control policy with various stacks.

You define the protocol stack with the command `class-map type stack <name>` . This class-map is always of type `match-all` and consists of one or more match entries. Every match entry should specify a protocol name defined in loaded PHDFs, a header field value, and the next protocol name found in stack. The below sample defines a series of headers such as: TCP in IPIP tunnel (protocol 4) encapsulated in

https://t.me/learningnets

IP and lastly in Ethernet.

https://t.me/learningnets

Configuration

Image not readable or empty
/uploads/workbooks/images/originals/4d4c28679aee8777dd7615bff2bef90a821f307d.png

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Section 2 Introduction

> **This section utilizes Firewall Services**. To start practicing this section, load the initial configuration files, namely **Threat Identification and Mitigation Configuration Files**.

Lab Topics

Threat Identification and Mitigation is the second blueprint topic of the CCIE Security Lab Exam. The high-level topics as stated by Cisco include the following:

- Identify and protect against fragmentation attacks
- Identify and protect against malicious IP option usage
- Identify and protect against network reconnaissance attacks
- Identify and protect against IP spoofing attacks
- Identify and protect against MAC spoofing attacks
- Identify and protect against ARP spoofing attacks
- Identify and protect against Denial of Service (DoS) attacks
- Identify and protect against Distributed Denial of Service (DDoS) attacks
- Identify and protect against Man-in-the-Middle (MiM) attacks
- Identify and protect against port redirection attacks
- Identify and protect against DHCP attacks
- Identify and protect against DNS attacks
- Identify and protect against MAC Flooding attacks
- Identify and protect against VLAN hopping attacks
- Identify and protect against various Layer2 and Layer3 attacks
- NBAR
- NetFlow
- Capture and utilize packet captures

These topics are further expanded in the following checklist.

https://t.me/learningnets

Technology	Mastered?
Implementing RFC 1918 Antispoofing Filtering	
Implementing RFC 2827 Antispoofing Filtering	
Implementing RFC 2401 Antispoofing Filtering	
Enabling a TCP Intercept on a Router	
Enabling a TCP Intercept on the Cisco ASA Security Appliance	
FPM (Flexible Packet Matching) and Protocol Header Definition File (PHDF) Files and Configuration of Nested Policy Maps	
Classification Using NBAR	
Understanding and Enabling NetFlow on a Router	
Port Security on a Switch	
Storm Control on a Switch	
Private VLAN (PVLAN) on a Switch	
Port Blocking on a Switch	
Port ACL on a Switch	
MAC ACL on a Switch	

https://t.me/learningnets

Technology	Mastered?
VLAN ACL on a Switch	
Spanning Tree Protocol (STP) Protection Using BPDU Guard and Loop Guard on a Switch	
DHCP Snooping on a Switch	
IP Source Guard on a Switch	
Dynamic ARP Inspection (DAI) on a Switch	
SeND for ND Protection	
IPv6 First Hop Security	
Disabling DTP on All Nontrunking Access Ports	
Concept of Proactive vs. Reactive Measures	
Knowledge of Protocols: TCP, UDP, HTTP, SMTP, ICMP, FTP	
Knowledge of Common Attacks: Network Reconnaissance, IP Spoofing, DHCP Snooping, DNS Spoofing, MAC Spoofing, ARP Snooping, Fragment Attack, Smurf Attack, TCP SYN Attack	
Understanding and Interpreting ARP Header Structure	
Understanding and Interpreting IP Header Structure	
Understanding and Interpreting TCP Header Structure	

https://t.me/learningnets

Technology	Mastered?
Understanding and Interpreting UDP Header Structure	
Understanding and Interpreting HTTP Header Structure	
Understanding and Interpreting ICMP Header structure	
Understanding and Interpreting ICMP Type Name and Codes	
Understanding and Interpreting Syslog Messages	
Understanding and Interpreting Packet Capture Outputs (Sniffer, Ethereal, Wireshark, TCPDump)	
Understanding Different Types of Attack Vectors	
Interpreting Various show and debug Outputs	
Classifying Attack Patterns Using FPM	
Memorizing Common Protocol and Port Numbers	
Preventing an ICMP Attack Using ACLs	
Preventing an ICMP Attack Using NBAR	
Preventing an ICMP Attack Using Policing	
Preventing an ICMP Attack Using the Modular Policy Framework (MPF) on the Cisco ASA Security Appliance	

https://t.me/learningnets

Technology	Mastered?
Preventing a SYN Attack Using ACLs	
Preventing a SYN Attack Using NBAR	
Preventing a SYN Attack Using Policing	
Preventing a SYN Attack Using CBAC	
Preventing a SYN Attack Using CAR	
Preventing a SYN Attack Using a TCP Intercept	
Preventing a SYN Attack Using the Modular Policy Framework (MPF) on the Cisco ASA Security Appliance	
Preventing Application Protocol–Specific Attacks Using FPM (e.g., HTTP, SMTP)	
Preventing Application Protocol–Specific Attacks Using NBAR (e.g., HTTP, SMTP)	
Preventing Application Protocol–Specific Attacks Using the Modular Policy Framework (MPF) on the Cisco ASA Security Appliance (e.g., HTTP, SMTP)	
Preventing IP Spoofing Attacks Using Antispoofing ACLs	
Preventing IP Spoofing Attacks Using uRPF	
Preventing IP Spoofing Attacks Using IP Source Guard	
Preventing Fragment Attacks Using ACLs	

https://t.me/learningnets

Technology	Mastered?
Preventing MAC Spoofing Attacks Using Port Security	
Preventing ARP Spoofing Attacks Using DAI	
Preventing VLAN Hopping Attacks Using the switchport mode access Command	
Preventing STP Attacks Using the Root Guard or BPDU Guard	
Preventing DHCP Spoofing Attacks Using Port Security	
Preventing DHCP Spoofing Attacks Using DAI	
Preventing Port Redirection Attacks Using ACLs	

In the following labs, focus on learning the details. This will assist you in mastering the technologies above.

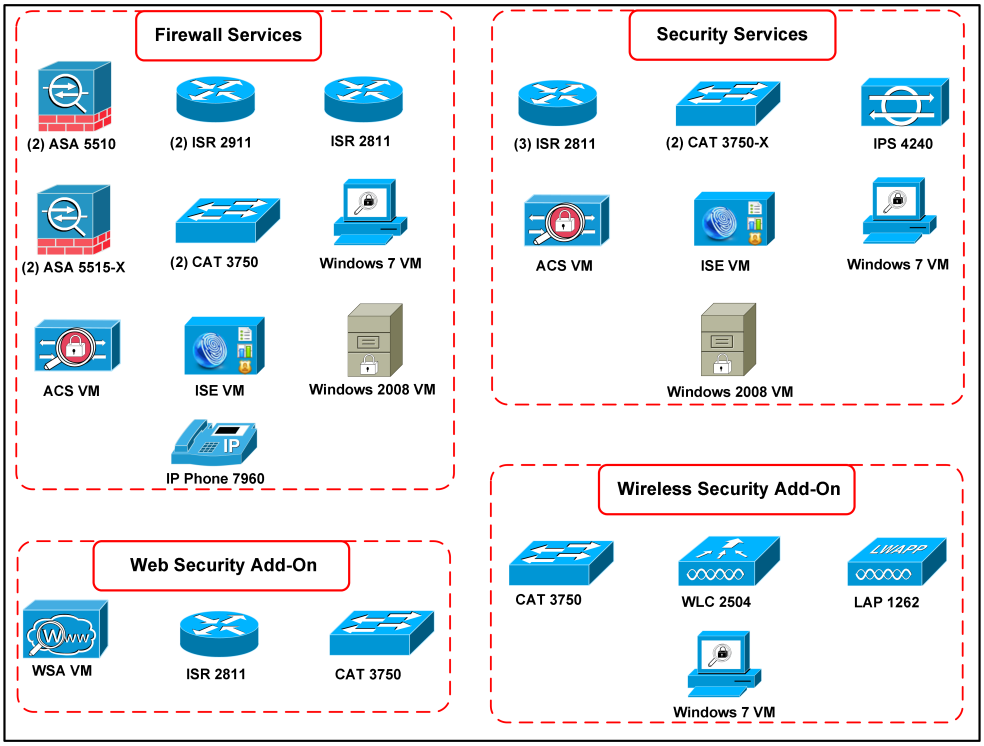
Section 2 Configuration Files

The configuration files found at the top of this page can be used to initialize your rack. Load these configurations before starting, and do each lab in sequence.

Section 2 General Lab Topology

The INE CCIE Security Racks are divided into modules as seen in the figure below. This section, Section 2, utilizes **Firewall Services**.

https://t.me/learningnets



Use the Section 2 Topology Diagram as your reference for the labs in this section.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124
- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IOS NBAR Traffic Classification

Task

- Configure R3 to drop any potentially dangerous HTTP downloads with the extensions `.exe` , `.com` , or `.bin` .
- Use a single pattern to match all the file names at once.
- Apply the configuration egress on VLAN 143 interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The NBAR (Network Based Application Recognition) protocol classification mechanism allows deep packet inspection for some protocols. One notable feature is matching various fields inside HTTP headers. Another important aspect is that NBAR matching is stateful and applies to the complete traffic flow in both directions.

Commonly, you see HTTP inspection used for matching the URL field against certain patterns. The URL matching occurs for HTTP GET/PUT/POST requests. It is important to note that the NBAR engine classifies bi-directional traffic flows. That is, if you apply a policy map matching the URL pattern inbound on an interface, the policy map will classify both inbound HTTP client requests as well as outgoing HTTP server replies. This is because NBAR classification occurs independently of policy-map direction, and the classification decision applies to both parts of a bi-directional traffic flow. The command to match a URL is `match http protocol url <pattern>`

https://t.me/learningnets

. The URL matching feature uses special wildcard characters:

```
"*"  – matches any sequence of characters (non-empty)
"?"  – matches any single character (you need to press Ctrl-V in order to enter "?")
"|"  – alternative, logical OR
"[]" – matches a range of characters, e.g. [ab] matches either "a" or "b"
"()" – grouping; delimit the logical end of a pattern; for example, you can use "*.(exe|bin)" as equivalent to "*.ex
```

All matching is case insensitive, so the pattern **text** matches **TEXT** as well. The engine matches your URL pattern against the directory path and the file name in the URL. For example, if the URL string is **http://www.cisco.com/pub/uploads/image.jpeg**, the matching procedure will only use the **pub/uploads/image.jpeg** part of the URL. When you submit a request like the above URL, it translates into the following headers (there are actually more, but this is the bare minimum):

```
GET /pub/uploads/image.jpeg HTTP/1.1
Host: www.cisco.com
```

Therefore, if you need to match a hostname, you should use the `match protocol http host` statement. You can use the same wildcard characters described above to match the HTTP host header. We are going to use the following single-line expression for our task: `*.bin|*.exe|*.com` . This pattern matches files with extension **.com**, **.exe**, or **.bin**. You can also use the `match-any` class map and match multiple URL strings on separate lines, resulting in the same effect as separating a pattern with the symbol `|` .

Configuration

R3:

```
class-map match-all EXTENSION
 match protocol http url *.bin|*.exe|*.com
!
policy-map DROP
 class EXTENSION
  drop
!
!Note that NBAR protocol classification engine needs to be enabled at the interface level separately.
!
interface FastEthernet0/0.143
 ip nbar protocol-discovery
```

https://t.me/learningnets

```
 service-policy output DROP
```

Verification

Simulate an HTTP server on SW1 and create some files in the flash with extensions “exe”, “com”, and “bin”:

```
SW1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.SW1(config)#ip http path flash:
!
!SW1#copy running-config flash:config.exe
Destination filename [config.exe]?
4619 bytes copied in 2.173 secs (2126 bytes/sec)
!
!SW1#copy running-config flash:config.bin
Destination filename [config.bin]?
4619 bytes copied in 1.132 secs (4080 bytes/sec)
!
!SW1#copy running-config flash:config.com
Destination filename [config.com]?
4619 bytes copied in 1.116 secs (4139 bytes/sec)
```

Now try downloading these files on R1 across R3; transfer of banned files should be forbidden and matches inside the class-map on R3 should occur:

```
R1#copy http://cisco:cisco@172.1.114.7/config.exe null:
Accessing http://*****:*****@172.1.114.7/config.exe...
%Error opening http://*****:*****@172.1.114.7/config.exe (I/O error)
!
!R1#copy http://cisco:cisco@172.1.114.7/config.bin null:
Accessing http://*****:*****@172.1.114.7/config.bin...
%Error opening http://*****:*****@172.1.114.7/config.bin (I/O error)
!
!R1#copy http://cisco:cisco@172.1.114.7/config.com null:
Accessing http://*****:*****@172.1.114.7/config.com...
%Error opening http://*****:*****@172.1.114.7/config.com (I/O error)
!
!R1#copy http://cisco:cisco@172.1.114.7/config.text null:
Accessing http://*****:*****@172.1.114.7/config.text...
Loading http://*****:*****@172.1.114.7/config.text 3275 bytes copied in 0.036 secs (90972 bytes/sec)
!
!R3#show policy-map interface fastEthernet0/0.143
```

https://t.me/learningnets

```
FastEthernet0/0.143

 Service-policy output: DROP

   Class-map: EXTENSION (match-all) 53 packets, 6218 bytes
     5 minute offered rate 0 bps, drop rate 0 bps Match: protocol http url "*.bin|*.exe|*.com"

     drop

   Class-map: class-default (match-any)
     134 packets, 9478 bytes
     5 minute offered rate 0 bps, drop rate 0 bps
     Match: any
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IOS Netflow Statistics

Task

- Configure R1 and R3 to export Netflow version 5 information to host 136.1.13.20 on port 9999.
- R1 should collect ingress flow information on its VLAN 13 interface, and R3 should collect outgoing flow information on its VLAN 23 interface.
- Limit the size of the flow cache to 4096 entries and configure to export BGP AS of Origin if available.
- Collect additional information about packet size statistics in flows, ICMP messages information, and layer 2 VLAN number where applicable.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Netflow evolved from a route-caching technology to a powerful statistics-gathering tool. You may enable flow collection on interfaces tracking either ingress or egress flow with interface-level command `ip flow {ingress|egress}` . The difference between those two is that ingress flow monitoring will capture flows received by the router, including the packets destined to the router itself, whereas egress flow collects packets switched across the router, exiting a router interface (does not include packets that originated from router itself). Egress flow collection has some performance penalties, but it is useful when you need to collect netflow information

https://t.me/learningnets

on MPLS traffic, if it exits the router untagged (netflow does not understand MPLS tagged packets).

You may distinguish egress flows in the show command's output by recognizing a `*` sign near the destination interface. Another difference between ingress and egress netflow is their interaction with other IOS features, such as NAT, rate-limiting, encryption. Basically, ingress netflow is applied before ACLs and rate-limiting, and egress netflow is applied after rate-limiting and encryption features.

Different versions of Netflow are available on IOS routers. The most basic and widely available is version 5, which is also the default. Additional versions include version 8, which supports aggregation caches, and version 9, which is incompatible with previous versions but allows for the most flexibility with choosing the export template format. Note that aggregation caches are also supported by version 9, so version 8 is becoming deprecated. To configure the version, use the global command `ip flow-export version [1|5|9] {origin-as|peer-as|bgp-nexthop}` ; the additional BGP attributes that can be recorded by Netflow and exported are available only if version 5 or 9 is used. To verify your configuration, generate some traffic and check flow-cache output using the command `show ip cache verbose flow .`

Netflow can be configured to collect additional layer 2 or layer 3 information for the traffic such as: ICMP types/codes, length of the IP datagram, TTL values, source/destination MAC addresses. This is configured using the global command `ip flow-capture [vlan-id|icmp|packet-length|ttl|mac-address] .`

Configuration

```
R1:
ip flow-capture vlan-id
ip flow-capture icmp
ip flow-capture packet-length
ip flow-export version 5 origin-as
ip flow-export destination 136.1.13.20 9999
ip flow-cache entries 4096
!
interface GigabitEthernet0/0
 ip flow ingress


R3:

ip flow-capture vlan-id
ip flow-capture icmp
ip flow-capture packet-length
```

https://t.me/learningnets

```
ip flow-export version 5 origin-as
ip flow-export destination 136.1.13.20 9999
ip flow-cache entries 4096
!
interface FastEthernet0/1.23
 ip flow egress
```

Verification

Generate some traffic between R1 and R2 and check that both R1 and R3 have the flow in the cache. Notice that flow is marked as egress in R3:

```
R1#ping 136.1.23.2 size 1500 repeat 200
Type escape sequence to abort.
Sending 200, 1500-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (200/200), round-trip min/avg/max = 1/2/4 ms
!
!R1#show ip cache flow
IP packet size distribution (1229 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .025 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 .974 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 278544 bytes 1 active, 4095 inactive, 3 added
  178 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 34056 bytes
  2 active, 1022 inactive, 6 added, 3 added to flow
  0 alloc failures, 0 force free
  1 chunk, 1 chunk added
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
ICMP                 2      0.0       599  1500      0.0       3.3      15.2
Total:               2      0.0       599  1500      0.0       3.3      15.2


SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Gi0/0         136.1.13.3      Null          224.0.0.10      58 0000 0000    31
```

https://t.me/learningnets

```
!
!R3#show ip cache flow
IP packet size distribution (1398 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 1.00 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 278544 bytes 1 active, 4095 inactive, 3 added
  52 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 34056 bytes
  2 active, 1022 inactive, 6 added, 3 added to flow
  0 alloc failures, 0 force free
  1 chunk, 1 chunk added
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
ICMP                 2      0.0       599  1500      0.0       3.3      15.7
Total:               2      0.0       599  1500      0.0       3.3      15.7

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Fa0/1.13      136.1.13.1      Fa0/1.23*     136.1.23.2      01 0000 0800   200
```

Check the flow export statistics:

```
R1#show ip flow export
Flow export v5 is enabled for main cache
  Export source and destination details :
  VRF ID : Default Destination(1)  136.1.13.20 (9999)
  Version 5 flow records, origin-as
  3 flows exported in 3 udp datagrams
  0 flows failed due to lack of export packet
  0 export packets were sent up to process level
  0 export packets were dropped due to no fib
  0 export packets were dropped due to adjacency issues
  0 export packets were dropped due to fragmentation failures
  0 export packets were dropped due to encapsulation fixup failures
!
!R3#show ip flow export
Flow export v5 is enabled for main cache
  Export source and destination details :
```

https://t.me/learningnets

```
  VRF ID : Default Destination(1)  136.1.13.20 (9999)
Version 5 flow records, origin-as

  3 flows exported in 3 udp datagrams
  0 flows failed due to lack of export packet
  0 export packets were sent up to process level
  0 export packets were dropped due to no fib
  3 export packets were dropped due to adjacency issues
  0 export packets were dropped due to fragmentation failures
  0 export packets were dropped due to encapsulation fixup failures
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing ICMP Attacks Using ACL

Task

- Hosts behind SW2 are under an ICMP attack sourced from VLAN 114.
- Configure SW2 to protect hosts from the attack, but allow traffic required for troubleshooting purposes: ping, traceroute, path MTU discovery, router buffer drops.
- Configure ASA3 to allow only ICMP Echo-Reply traffic destined to its VLAN 133 interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

We are asked to drop most of the ICMP traffic coming from VLAN 114, so we need to match specifically on the source. PING uses Echo and Echo-Reply; the Microsoft version of traceroute uses Echo, Echo-Reply, and TTL-Exceeded, and the Linux version of traceroute uses UDP/TCP traffic plus TTL-Exceeded and Port-Unreachable. Path MTU sets the DF bit to discover the end-to-end MTU and therefore uses Packet-Too-Big; in accordance with to RFC 792, when a router drops packets because of insufficient buffer to queue the packet on its outgoing interface toward the destination (or when it approaches the buffer capacity), it may send a Source-Quench ICMP message back to the source, basically telling it to slow the sending rate. Remember that IOS inbound access-lists apply to both transit traffic

https://t.me/learningnets

and non-transit (control-plane) traffic.

Configuration

```
SW2:
ip access-list extended ICMP_ATTACK
 permit icmp 172.1.114.0 0.0.0.255 any echo
 permit icmp 172.1.114.0 0.0.0.255 any echo-reply
 permit icmp 172.1.114.0 0.0.0.255 any ttl-exceeded
 permit icmp 172.1.114.0 0.0.0.255 any port-unreachable
 permit icmp 172.1.114.0 0.0.0.255 any packet-too-big
 permit icmp 172.1.114.0 0.0.0.255 any source-quench
 deny icmp 172.1.114.0 0.0.0.255 any
 permit ip any any
!
interface Vlan213
 ip access-group ICMP_ATTACK in

ASA3:

icmp permit any echo-reply OUTSIDE
```

Verification

Verify that all but ICMP Echo-Reply traffic is blocked if destined to the ASA3 VLAN 133 interface. For example, verify that ICMP Echo is not allowed:

```
ASA3# ping 136.1.133.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.133.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
!
!R3#ping 136.1.133.13

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.133.13, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)
```

https://t.me/learningnets

Verify that ICMP traffic required for troubleshooting purposes is allowed from VLAN 114:

```
SW2#ping 172.1.114.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.1.114.7, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!SW2#traceroute 172.1.114.7

Type escape sequence to abort.
Tracing the route to 172.1.114.7
 1 136.1.133.3 0 msec 8 msec 0 msec
2 172.1.114.7 9 msec *  0 msec
!
!SW1#ping 136.1.213.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.213.8, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing ICMP Attacks Using NBAR

Task

- Using NBAR, deny any ICMP traffic traversing R3 toward VLAN 23, outbound on its Ethernet interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

NBAR in general allows only matching the protocol type, so it does not allow matching on specific ICMP types or codes; it does not have any granularity for ICMP traffic, so we will actually drop all ICMP traffic types. If specific ICMP traffic must be matched, inside the class-map we can combine NBAR protocol matching with access-lists.

Configuration

R3:

```
class-map match-all ICMP
 match protocol icmp
!
policy-map ICMP
```

https://t.me/learningnets

```
 class ICMP
  drop
!
!Note that NBAR classification engine needs to be enabled at the interface level for NBAR traffic policies to be mat
!
interface FastEthernet0/1.23
 ip nbar protocol-discovery
 service-policy output ICMP
```

Verification

Generate ICMP traffic toward VLAN23 and verify that it is dropped by configured NBAR policy; also verify that ICMP shows up as a discovered protocol on VLAN23 interface:

```
R1#ping 136.1.23.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R3#show policy-map interface fastEthernet0/1.23
 FastEthernet0/1.23

  Service-policy output: ICMP

    Class-map: ICMP (match-all) 5 packets, 590 bytes
      5 minute offered rate 0 bps, drop rate 0 bps Match: protocol icmp
      drop

    Class-map: class-default (match-any)
      23 packets, 1902 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
!
!R3#show ip nbar protocol-discovery interface fastEthernet0/1.23

 FastEthernet0/1.23

 Last clearing of "show ip nbar protocol-discovery" counters 00:00:58


                            Input                  Output
                            -----                  ------
   Protocol                 Packet Count           Packet Count
```

https://t.me/learningnets

```
                          Byte Count                Byte Count
                          5min Bit Rate (bps)       5min Bit Rate (bps)
                          5min Max Bit Rate (bps)   5min Max Bit Rate (bps)
    ----------------------- ------------------------ ------------------------
    eigrp                   0                        13
                            0                        1014
                            0                        0
                            0                        0
icmp                        0                        5


                            0                        590
                            0                        0
                            0                        0
    Total                   0                        18
                            0                        1604
                            0                        0
                            0                        0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing ICMP Attacks Using Policing

Task

- Configure R3 to limit the scope of an ICMP attack based on Echo-Reply packets, inbound on its VLAN 13 connection.
- If the traffic rate is less than 128 Kbps, mark the packets with an IP precedence of 1. If it exceeds 128 Kbps, mark the packets with an IP precedence of 0.
- Drop any violating traffic.
- Ensure that the burst size is large enough to accommodate normal and excess burst durations of 200 ms and 300 ms at a rate of 128 Kbps.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

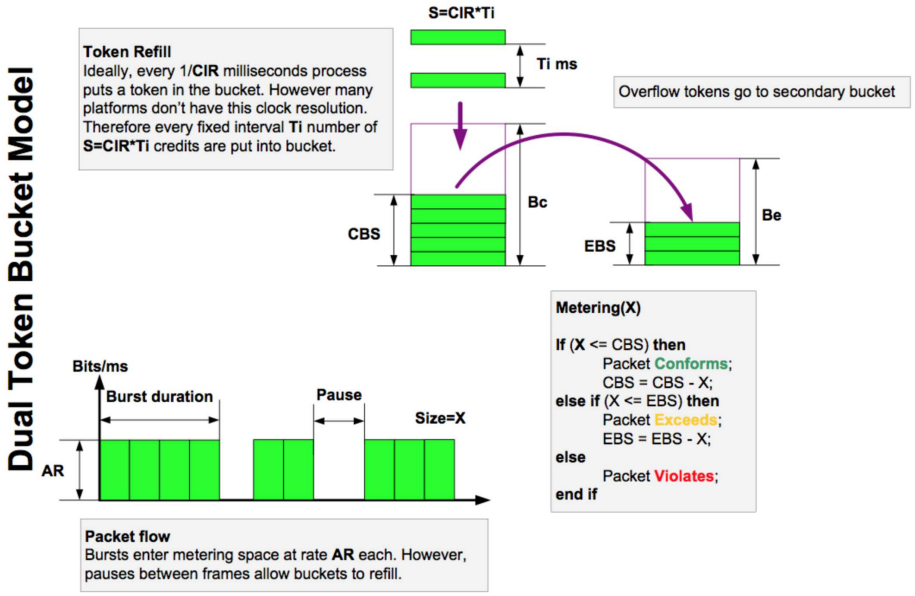
Overview

A Single Rate Three Color Marker (Policer) or srTCM is the RFC-based implementation of the metering process. Compared to legacy CAR, functionality is changed, although scope is the same to `rate-limit` traffic accordingly to configured policies.

First, the concept of Excess Burst (Be) is no longer tied to a pseudorandom marking behavior. Rather, **Be** is used as extra credit for periods of inactivity. Second, the policing process uses a Dual Token Bucket model, but the implementation is different. **Three color** in this case means that incoming traffic is metered against

https://t.me/learningnets

CIR using the configured burst sizes, and the result is a marking using one of three colors: Green (Conform), Yellow (Exceed), Red (Violate). The token bucket exhibits the concept of a sliding metering window over the packet flow and uses a secondary bucket called Excess Burst.



The Excess Burst bucket accumulates extra credit, spilled over from the Normal Bucket when it overfills. This only happens during long periods of pause that exceed the averaging interval **Tc**. With just a single bucket, the process does not account for extended periods of silence. The Excess Burst bucket allows for accumulating credits and using them when incoming burst exceeds the configured **Bc** value. Thus, **Be** in this usage serves a purpose similar to that of the **Be** functionality with traffic-shaping. In fact, these two are complementary in the sense that srTCM's **Be** properly recognizes the excessive bursting of traffic shaping. In a classic token bucket model, a special timer fires every **Ti** interval, triggering the token refresh procedure. This method, however, limits the resolution of metering. At the same time, the amount of CPU operations needed to meter a packet is small. This is because the procedure uses only addition and subtraction operations, which consume few CPU cycles.

Cisco's implementation of the token bucket for srTCM uses a different algorithm. This algorithm has a so-called unlimited precision, but it consumes more CPU cycles per packet. The process keeps track of three variables:

- CBS – current normal burst size
- EBS – current excess burst size
- T0 – last packet arrival time

https://t.me/learningnets

Now imagine that a new packet of size **S** arrives at time **T**.

```
Step 1:
The process computes accumulated credit:

Cr = CIR*(T-T0)

then adds this values to CBS:

if ((CBS + Cr) <= Bc) then
       CBS = CBS + Cr;
else
       CBS = Bc;
end if

Step 2:
The process computes a new EBS size if there are enough credits:

If (CBS + Cr > Bc) then
       EBS = CBS + Cr – Bc;
end

if (EBS > Be) then
       EBS = Be;

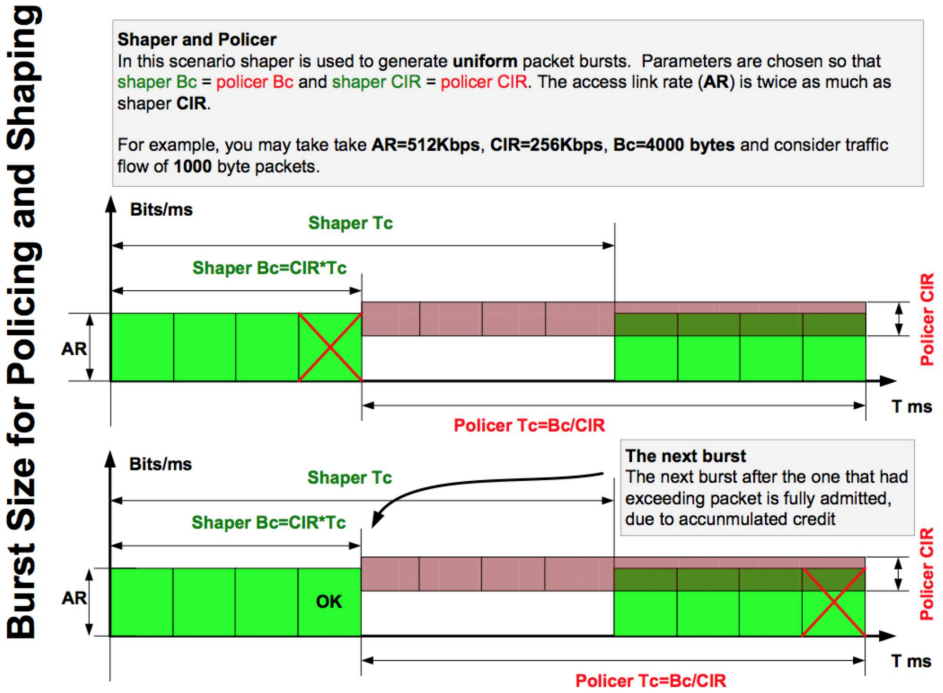
Step 3:
The process compares packet size to accumulated credits and performs marking:

if ( S < CBS ) then
     Packet Conforms;
     CBS = CBS – S;
else if (S < EBS) then
     Packet Exceeds;
     EBS = EBS – S;
else
     Packet Violates;
end if
```

This implementation does not use a hardware timer, so the precision is not limited. However, additional multiplication operations may become CPU consuming on intense packet flows. For an optimal burst size selection, if you are using srTCM to perform traffic admission at the network edge, ensure that your policing burst sizes match at least one packet size more than the burst size configured at the customer side. This is necessary because now the router meters traffic with extra precision,

https://t.me/learningnets

and the effect illustrated in the following figure (shaper Bc = policer Bc) may occur.



Note that the policer may mark the beginning of the second burst as exceeding. This is because the policer **Bc** is exactly equal to the shaper **Bc**. In this situation, when the sliding window hits the beginning of the next burst, it has no capacity to accommodate another packet. The packet marked as exceeding does not account for the sliding window content, so the next burst will have all four packets admitted. For this ideal simulation, the policer **Bc** should be larger than traffic burst size by at least one packet. This brings us to the formula:

Policer Bc = Shaper Burst Size + Average Packet Size

In addition, it makes sense to configure burst size proportional to the average packet size as well, to ensure more precise matching. However, this is not necessary, especially with larger bursts. In this particular example, the burst sizes are based on the time intervals as follows:

```
Bc = (128000*0,2)/8=3200 bytes
Be = (128000*0,3)/8=4800 bytes
```

Configuration

```
R3:

ip access-list extended ICMP_ATTACK
```

https://t.me/learningnets

```
 permit icmp any any echo-reply
!
class-map ICMP_ATTACK
 match access-group name ICMP_ATTACK
!
policy-map POLICE_VLAN13
 class ICMP_ATTACK
   police 128000 3200 4800 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action drop
!
interface FastEthernet0/1.13
  service-policy input POLICE_VLAN13
```

Verification

To test the policer, simulate ICMP Echo-Reply traffic flow from R1 to R3:

https://t.me/learningnets

```
R3#ping 136.1.13.1 repeat 50 size 1500 timeout 1


Type escape sequence to abort.
Sending 50, 1500-byte ICMP Echos to 136.1.13.1, timeout is 1 seconds:
!!!!!.!!!!!.!!!!!.!!!!!.!!!!!.!!!!!.!!!!!.!!!!!.!!

Success rate is 84 percent (42/50), round-trip min/avg/max = 1/2/4 ms
!
!R3#show policy-map interface fastEthernet0/1.13

 FastEthernet0/1.13

  Service-policy input: POLICE_VLAN13

    Class-map: ICMP_ATTACK (match-all)
      50 packets, 75700 bytes
      30 second offered rate 12000 bps, drop rate 2000 bps
      Match: access-group name ICMP_ATTACK
      police:
          cir 128000 bps, bc 3200 bytes, be 4800 bytes conformed 18 packets, 27252 bytes; actions:

          set-prec-transmit 1 exceeded 24 packets, 36336 bytes; actions:

          set-prec-transmit 0 violated 8 packets, 12112 bytes; actions:

          drop
        conformed 4000 bps, exceed 5000 bps, violate 2000 bps

    Class-map: class-default (match-any)
      10 packets, 740 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing ICMP Attacks Using MPF

Task

- Configure ASA4 to limit ICMP Echo traffic toward VLAN 143 to 128 Kbps.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Policing configuration on ASA uses MPF syntax, where traffic is identified within a class-map, policing action and direction is defined within a policy-map, and policy gets applied at global or interface level. With MPF in general, depending on the configuration, there can be differences between final results when applying a policy at interface-level or global-level. This does not, however, apply to policing; regardless of where the point of attachment is, the end result is the same. On the ASA, policing can be applied both in the inbound and the outbound directions.

Configuration

ASA4:

```
access-list ICMP_VLAN143 extended permit icmp any 172.1.143.0 255.255.255.0 echo
!
class-map ICMP_VLAN143
```

https://t.me/learningnets

```
 match access-list ICMP_VLAN143
!
policy-map ICMP_VLAN143
 class ICMP_VLAN143
  police output 128000
!
service-policy ICMP_VLAN143 interface OUTSIDE
```

Verification

Issue an extended PING from SW1 to R3 and note the dropped ICMP packets:

```
SW1#ping 172.1.143.3 size 1500 repeat 100 timeout 1

Type escape sequence to abort.
Sending 100, 1500-byte ICMP Echos to 172.1.143.3, timeout is 1 seconds:
!!!!!!!!!!!!.!!!!!!!!!!!!.!!!!!!!!!!!!.!!!!!!!!!!!!.!!!!!!!!!!.!!!!!!!!!!!!.!!!!!!!!!!!!.!!
Success rate is 92 percent (92/100), round-trip min/avg/max = 1/4/9 ms
!
!ASA4#show service-policy police

Interface OUTSIDE:
  Service-policy: ICMP_VLAN143
    Class-map: ICMP_VLAN143
      Output police Interface OUTSIDE:
        cir 128000 bps, bc 4000 bytes conformed 92 packets, 139288 bytes; actions:  transmit
 exceeded 8 packets, 12112 bytes; actions:  drop

        conformed 0 bps, exceed 0 bps
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing SYN Attacks Using ACL

Task

- Assuming that R3's VLAN 13 interface represents your network edge segment, ensure that no TCP sessions are allowed inbound, unless they have been initiated from inside of your network.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Using an ACL entry with the `established` keyword was one of the first steps toward today's stateful firewalls. Although it offers weak security and is simple to bypass, you need to know it as an alternative. The keyword ensures that the TCP packet has either the **ACK** or **RST** bits set, so it is not a new TCP session initiation that only contains the **SYN** bit, thus it is a response to a TCP session. The alternative to using the keyword is to actually match the **RST** or **ACK** bits. If you choose to type the command `permit tcp any any ack rst` inside an ACL, IOS will actually convert it to `permit tcp any any established`. This is not truly a stateful firewall because it does not track a session state, but it still provides a basic level of security.

Configuration

https://t.me/learningnets

R3:

```
ip access-list extended TCP_ESTABLISHED
 permit tcp any any established
 deny tcp any any
 permit ip any any
!
interface FastEthernet0/1.13
 ip access-group TCP_ESTABLISHED in
```

Verification

Verify that TCP sessions initiated from behind R3 are allowed to form, and that TCP sessions initiated from R1 are dropped. Also verify the matched ACL entries on R3:

```
R2#telnet 136.1.13.1
Trying 136.1.13.1 ... Open

User Access Verification

Password: R1>
!
!R1#telnet 136.1.23.2
Trying 136.1.23.2 ... % Destination unreachable; gateway or host down
!
!R3#show ip access-lists TCP_ESTABLISHED
Extended IP access list TCP_ESTABLISHED 10 permit tcp any any established (20 matches)
 20 deny tcp any any (1 match)

    30 permit ip any any (200 matches)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing SYN Attacks Using Policing

Task

- There have been several SYN flood attacks targeting servers in VLAN 13. Although the source is not exactly known, configure R3 to allow no more than 16 Kbps of SYN traffic inbound on its VLAN 133 interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

We'll match the specific traffic using an access-list, and then use MQC syntax to police the traffic to the specified rate.

Configuration

R3:

```
ip access-list extended SYN_VLAN13
 permit tcp any 136.1.13.0 0.0.0.255 syn
!
class-map SYN_VLAN13
 match access-group name SYN_VLAN13
!
```

https://t.me/learningnets

```
policy-map SYN_VLAN13
 class SYN_VLAN13
  police 16000 conform-action transmit exceed-action drop violate-action drop
!
interface FastEthernet0/0.133
 service-policy input SYN_VLAN13
```

Verification

Initiate a TCP session from SW2 toward VLAN 13, and verify that traffic matches on R3's configuration:

```
SW2#telnet 136.1.13.1
Trying 136.1.13.1 ... Open

User Access Verification

Password: R1>
!
!R3#show policy-map interface fastEthernet0/0.133
 FastEthernet0/0.133

  Service-policy input: SYN_VLAN13

    Class-map: SYN_VLAN13 (match-all) 1 packets, 64 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name SYN_VLAN13
      police:
          cir 16000 bps, bc 1500 bytes, be 1500 bytes conformed 1 packets, 64 bytes; actions:
          transmit  exceeded 0 packets, 0 bytes; actions:
          drop  violated 0 packets, 0 bytes; actions:

          drop
        conformed 0 bps, exceed 0 bps, violate 0 bps

    Class-map: class-default (match-any)
      24 packets, 1685 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing SYN Attacks Using CBAC

Task

- Configure R3 to allow TCP and UDP traffic from hosts in VLAN 13 using stateful inspection.
  - Apply the inspection inbound on VLAN 13.
- Ensure that R3 monitors TCP sessions as follows:
  - Limit half-open sessions to 50 per host and block a host for 5 minutes when the threshold is reached.
  - Reset a TCP session if not established within 10 seconds.
  - Generate a log for CBAC dropped packets.
- Ensure that no TCP sessions can be initiated toward VLAN 13 hosts.
  - Do not use any access-lists to accomplish this.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

CBAC functionality is explained in detail in Section 6, **Perimeter Security and Services**. CBAC, or context-based access-control, introduces true stateful inspection to IOS code. It is an upgraded version of reflexive ACLs, which did not support protocols negotiating secondary data channels. The general idea is to inspect protocol-specific information in traffic flows going across the router and

https://t.me/learningnets

dynamically open holes in access-lists for returning traffic. CBAC configurations commonly consist of two parts:

- The CBAC inspection rule (defines protocols to be inspected) applied to an interface in the direction matching the initial flow of traffic (usually from the protected network to the outside world or unprotected network).
- An access-list applied to an interface in the inbound direction, matching the returning traffic flow (usually from the unprotected network to the protected network).

By default, with every inspection rule, CBAC enables DoS prevention features very similar to the TCP Intercept feature. CBAC applies DoS prevention to UDP sessions as well. A UDP session is considered half-open if no packets have been sent in response to the initial packet. Both UDP and TCP half-open sessions count against maximum connection thresholds at the same time; therefore, if you configure a limit of 100 half-open sessions, the router will count both UDP and TCP sessions against this threshold.

CBAC supports both maximum thresholds and per-minute rates for half-open sessions. When the current number of connections exceeds any of the thresholds, the router starts deleting half-open connection entries and sending RST packets to TCP session endpoints; it does not delete any sessions that have already been established. For TCP sessions, the only supported mode of operations is **watch**. The router does not actively proxy TCP sessions, it just watches as being established. Offending sessions (sessions that did not complete in time or that exceed thresholds) receive RST messages as well when exceeding configured timeouts. The amount of time CBAC waits for TCP connection establishment is configured using the command `ip inspect tcp synwait-time <time-interval>` , which is the equivalent of `watch-timeout` for the TCP Intercept feature.

One feature specific to CBAC is per-host TCP connection limits. CBAC bounds the number of half-open connections per-host, in addition to having the global thresholds. As the number of half-open sessions destined to the same IP address exceeds the per-host threshold, the router can block any further connections from establishing for the amount of time specified. The command syntax is `ip inspect tcp max-incomplete host <session-number> block-time <time-interval-minutes>` . CBAC TCP Intercept functionality is enabled by default when TCP sessions are inspected and it cannot be disabled.

You can enable CBAC inspection in the inbound or outbound direction of an interface. For example:

- If you apply CBAC inbound on interface Fa0/0, traffic matched by CBAC rules is inspected when entering Fa0/0 and return traffic is allowed automatically, regardless

https://t.me/learningnets

of any access-lists configured inbound on other interfaces.

- If you apply CBAC outbound on interface Fa0/1, traffic matched by CBAC rules is inspected when exiting Fa0/1 and return traffic is allowed automatically, regardless of any access-list configured inbound only on Fa0/1.

You can configure the router so that no TCP sessions are allowed across the router, unless inspected by CBAC and therefore present in the state table, regardless of any access-lists configured. We can say that with this function enabled, inbound access-lists are ignored. This is configured with the global-level command **`ip inspect tcp block-non-session`**. If we consider the above two examples presented, the following behavior is expected:

- No TCP traffic is allowed to exit interface Fa0/0, unless present in the session table of CBAC, so traffic was initiated from behind Fa0/0.
- No TCP traffic is allowed to enter interface Fa0/1, unless present in the session table of CBAC, so traffic was initiated from behind any router interface except for Fa0/1.

Configuration

```
R3:


ip inspect tcp block-non-session
ip inspect tcp synwait-time 10
ip inspect tcp max-incomplete host 50 block-time 5
ip inspect log drop-pkt
!
ip inspect name CBAC_VLAN13 tcp
ip inspect name CBAC_VLAN13 udp
!
interface FastEthernet0/1.13
 ip inspect CBAC_VLAN13 in
```

Verification

Verify CBAC configuration:

```
R3#show ip inspect config
Dropped packet logging is enabled
Drop non-session initiated traffic is enabled
Session audit trail is disabled
Session alert is enabled
one-minute (sampling period) thresholds are [unlimited : unlimited] connections
```

https://t.me/learningnets

```
max-incomplete sessions thresholds are [unlimited : unlimited]
max-incomplete tcp connections per host is 50. Block-time 5 minutes.
tcp synwait-time is 10 sec -- tcp finwait-time is 5 sec

tcp idle-time is 3600 sec -- udp idle-time is 30 sec
tcp reassembly queue length 16; timeout 5 sec; memory-limit 1024 kilo bytes
dns-timeout is 5 sec
Inspection Rule Configuration
 Inspection name CBAC_VLAN13
    tcp alert is on audit-trail is off timeout 3600
    udp alert is on audit-trail is off timeout 30
```

Ensure that TCP sessions initiated from VLAN 13 are established:

```
R1#telnet 136.1.23.2
Trying 136.1.23.2 ... Open

User Access Verification

Password: R2>
!
!R3#show ip inspect sessions
Established Sessions Session 47508200 (136.1.13.1:65310)=>(136.1.23.2:23) tcp SIS_OPEN
```

Verify that TCP sessions initiated toward VLAN 13 are dropped:

```
R2#telnet 136.1.13.1
Trying 136.1.13.1 ... % Connection timed out; remote host not responding
!
!R3#show logging | i DROP
%FW-6-DROP_PKT: Dropping tcp session 136.1.13.1:23 136.1.23.2:44361  due to  Segment matching no TCP connection with
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing SYN Attacks Using CAR

Task

- Configure R2 to defend against SYN flooding attacks inbound on its VLAN 23 interface.
  - Drop malicious traffic exceeding 64 Kbps.
  - Use a committed burst of 64 Kbps and an excess burst of 128 Kbps.
- Do not use MQC syntax or the **syn** keyword to accomplish this.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The original idea behind tools such as Committed Access Rate (CAR) or policing was to enforce traffic contracts with customers at the edge of an oversubscribed network. Oversubscription is one of the prominent features of packet switched networks (PSNs). Statistical multiplexing used in PSNs allows multiple users to share the same resources in dynamic proportions. Based on the fact that not all users may need the same resources at the same time, it becomes possible to grant some users the ability to temporarily use more resources than the network can guarantee. This is in clear contrast with time-division multiplexing (TDM) networks, in which users cannot exceed the guaranteed share of resources, such as a group of TDM slots.

https://t.me/learningnets

In PSNs, service providers may sell more bandwidth than they can actually provide. For example, a 10 Mbps link could be shared between 10 users, with the agreement that each user can pull down up to 10 Mbps if nobody else uses the link, but users are only guaranteed 1 Mbps of link bandwidth in the case of heavy congestion. The actual implementation of this downstream sharing may vary, but most often it uses some sort of fair-queuing (for slow links) or round-robin (for high-speed links) scheduling, which implements the concept of max-min fairness for resource sharing.

To implement the upstream oversubscription, the provider uses a traffic metering procedure, such as CAR or policing, to mark the incoming traffic at the network edge. Theoretically, a policer may sit at the ingress edge of the service provider's network and meter traffic received from users, comparing it with the contracted CIR and PIR values.

By using CAR with the exceed action set to drop, it's possible to limit the rate of both incoming and outgoing traffic. Although CAR was not designed for this purpose originally, this feature allows for effective limiting of intense flows and prevention of DoS attacks. Finding the optimal **Bc** value is key to obtaining desired results and becomes a challenge. Also, note that the **Be** value with CAR is not used to define the Peak Information Rate (PIR), but rather to implement a pseudo-random exceed action. This feature simulates random early detection's drop behavior for TCP flows and prevents flow synchronization.

To rate-limit the SYN packets, we need to match on the TCP SYN bit, but we are not allowed to use the **syn** keyword. The remaining option is to exclude the TCP sessions that are already established from rate-limiting.

> Note that **target-bit-rate** value needs to be configured in bits, while **Bc** and **Be** values need to be configured in bytes.

Configuration

R2:

```
ip access-list extended 100
 deny tcp any any established
 permit tcp any any
!
interface GigabitEthernet0/0
 rate-limit input access-group 100 64000 8000 16000 conform-action transmit exceed-action drop
```

https://t.me/learningnets

Verification

Verify that SYN packets modify the rate-limit statistics:

```
R1#telnet 136.1.23.2
Trying 136.1.23.2 ... Open



User Access Verification

Password: R2>
!
!R2#show interfaces gigabitEthernet0/0 rate-limit
GigabitEthernet0/0 *** TO SW1 Fa0/2***
  Input
    matches: access-group 100 params:  64000 bps, 8000 limit, 16000 extended limit
conformed 2 packets, 134 bytes; action: transmit


      exceeded 0 packets, 0 bytes; action: drop
      last packet: 17816ms ago, current burst: 0 bytes
      last cleared 00:01:04 ago, conformed 16 bps, exceeded 0 bps
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing Application Attacks Using MPF

Task

- Configure FTP inspection on ASA3 as follows:
  - Deny the use of FTP command `DELETE` .
  - Deny the download of files with names that start with **cisco**, **c100**, or **c200**.
  - To prevent hackers from using known exploits, mask the FTP server banner and the system information reply.
  - Do not use the default policy-map.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

To configure application-specific restrictions on the ASA, Modular Policy Framework (MPF) is being used. Traffic inspection is the core of the stateful firewall, because it accomplishes the following main purposes:

- Inspects and understands the application layer; therefore, it dynamically modifies firewall rules to permit secondary data channels used by protocols such as H.323 and FTP.
- Performs deep traffic inspection to detect protocol violations and RFC conformance, anomalies, and attacks.

https://t.me/learningnets

- Dynamically modifies packet contents to remove potential threats or hide sensitive information.

MPF replaced the legacy `fixup` traffic inspection command from PIX with a totally new configuration and functional concepts. MPF configuration is similar to Modular QoS CLI (MQC) configuration from IOS routers. You may find many similarities between the IOS Zone Based Firewall and the ASA MPF configuration. To perform protocol inspection the ASA firewall, you should perform the following basic steps:

Define layer 3/layer 4 classification criteria.
For example, if you want to inspect FTP traffic, you may define an access-list to match it and then create a layer3/layer4 class-map that matches the traffic flow:

```
access-list FTP_TRAFFIC_ACL extended permit tcp any any eq 21
!
class-map FTP_TRAFFIC_CMAP
 match access-list FTP_TRAFFIC_ACL
```

This is very similar to MQC configuration from IOS routers in which class-maps are used for traffic classification based on layer 3/layer 4 criteria. As usual, you may create `match-any` or `match-all` class-maps:

- `math-any` requires at least one condition from the class-map to be matched.
- `match-all` requires all conditions from the class-map to be matched.

In addition to using ACLs, you can match traffic based on DSCP/Precedence numbers, although using the ACLs is more flexible. The ASA firewall also allows use of some unique classification criteria, such as matching VPN traffic flows.

Define a policy-map that associates certain actions with layer3/layer4 class-maps.
For example, here is a policy that inspects FTP traffic:

```
policy-map INSPECT_FTP_PMAP
 class FTP_TRAFFIC_CMAP
  inspect ftp
```

https://t.me/learningnets

This will apply FTP inspection engine to the traffic matching the layer 3/layer 4 class map. In addition to inspection, another common setting is the command `set connection` that allows changing various TCP session parameters. Most notably, you can change the maximum number of connections and the maximum number of half-open (embryonic) connections for the particular traffic class.

Apply the policy map to an interface or globally.
When you apply the service policy to an interface using the command `service-policy <policy-name> interface <nameif>`, all ingress or egress traffic on the specified interface is inspected according to the policy settings. When you apply the policy-map globally using the command `service-policy <policy-name> global`, it applies to all inbound traffic flows on all interfaces. If you have both the global and interface policies applied, the interface policy takes precedence for the same feature:

- If the interface-level policy inspects HTTP and the global policy inspects FTP, the end result is combined, both inspections being active.
- If both the interface and global policy inspect the same protocol (with potentially different settings), the interface-level policy-map takes precedence.
- The default inspection policy uses a special match statement, `match default-inspection-traffic`, which allows matching multiple default protocols at the same time. The respective class has a group of inspect statements associated in the default global policy-map. This was needed to replicate the default behavior of the old fixup commands.

Define application-level settings.
What if we would like to tune the inspection engine parameters, instead of simply relying on the default values? For example, we would like to tune HTTP and FTP inspection. To accomplish this, the firewall supports the concepts of application inspection class-maps and policy-maps. Defined using the syntax `class-map type inspect <protocol> <name>` and `policy-map type inspect <protocol> <name>`, these are designed to be used together. If we want to create a custom FTP inspection rule, we could use the following syntax:

```
policy-map type inspect ftp FTP_INSPECT
 parameters
  mask-banner
```

The custom inspection type policy map is bound to the inspect statement under the layer3/layer4 class in the regular policy-map:

https://t.me/learningnets

```
policy-map type inspect INSPECT_FTP_PMAP
 class FTP_INSPECT_CMAP
  inspect ftp FTP_INSPECT
```

Inspect type class-maps and policy-maps are available for a restricted number of protocols (HTTP, ESMTP, FTP, DNS, and few others); based on the protocol used, unique parameters and options are available. Some advanced inspection settings might require protocol-specific classifications. You cannot use regular layer3/layer4 class-maps with the inspect type policy-map. Remember that inspect type policy-maps have no effect until they are associated with the inspect command in a regular policy-map.

Configuration

ASA3:

```
regex REG_cisco "^cisco.*"
regex REG_c100 "^c100.*"
regex REG_c200 "^c200.*"
!
class-map FTP_PORT
 match port tcp eq 21
!
class-map type regex match-any FTP_REGEX
 match regex REG_cisco
 match regex REG_c100
 match regex REG_c200
!
class-map type inspect ftp match-all FTP_COMMAND
 match request-command dele
!
policy-map type inspect ftp FTP_INSPECT
 parameters
  mask-banner
  mask-syst-reply
 match filename regex class FTP_REGEX
  reset
 class FTP_COMMAND
  reset
!
policy-map VLAN133
 class FTP_PORT
```

https://t.me/learningnets

```
  inspect ftp strict FTP_INSPECT
!
service-policy VLAN133 interface OUTSIDE
```

Verification

Verify that FTP inspection is applied on the VLAN133 interface:

```
ASA3# show service-policy interface OUTSIDE inspect ftp

Interface OUTSIDE:
  Service-policy: VLAN133
    Class-map: FTP_PORT
      Inspect: ftp strict FTP_INSPECT, packet 0, lock fail 0, drop 0, reset-drop 0 mask-banner enabled
mask-syst-reply enabled
match filename regex class FTP_REGEX

          reset, packet 0
        class FTP_COMMAND
          reset, packet 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing IP Spoofing Attacks Using uRPF

Task

- You have been notified about IP spoofing attacks generated from behind R1.
  - Configure R3 to perform IP spoofing checks on its VLAN 13 interface.
  - A syslog message should be generated for spoofed packets.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

uRPF (Unicast Reverse Path Forwarding) is the concept of verifying the routing path for the source IP address found in an IPv4 unicast packet. Usually routers only use destination IP addresses to look up the next hop for an IP packet and route it, whereas RPF (Reverse-Path Forwarding) is used with multicast routing. Generally (except for asymmetrical routing scenarios), packets arrive on the interfaces that are on the shortest path to the source of the packets. However, with IP spoofing attacks, a malicious user may inject packets with IP addresses not belonging to its segment or network area. uRPF is a very useful security feature to prevent spoofing attacks, especially on routers that have a diverse number of connections, such as a system at the Internet peering point. The feature has two modes of operation:

- **Strict mode**, which you enable on the interface by using the `ip verify unicast source reachable-via rx` command (or the legacy format `ip verify unicast reverse-path`

https://t.me/learningnets

), and the router applies the uRPF check to the source IP addresses of incoming packets. The source IP address must match an explicit (non-default) IP route in the routing table, and the next hop for this entry should point out the interface the packet was received from. Otherwise, the router will drop the packets.

- **Loose mode**, which you enable on the interface by using the `ip verify unicast source reachable-via any` command. When a router receives a packet on the interface with loose uRPF enabled, it just checks that it has an IP route matching the source address of the packet. It does not matter whether the next hop for this route points out the interface the packet was received on or not. The exception to this verification is that if the route points to Null0, the packet is dropped. This feature is useful in enterprises that have more than one ISP uplink and possibly use asymmetric routing, and it is also used in BGP RTBH scenarios.

uRPF does have additional functionalities:

- uRPF exemptions and violation logging: if you specify a standard or extended access-list with the interface command `ip verify unicast source reachable-via [rx|any] <acl-number>` , the router consults this access-list for packets violating the uRPF condition. If the ACL permits the packet, it's allowed through; if it denies the packet, the router drops it. You may use the `log` keyword in the access-list to log the packets allowed or denied by the uRPF access-list.
- Allowing the router to do a self-ping with strict uRPF enabled on the interface using the interface command `ip verify unicast source reachable-via rx allow-self-ping` (accepting packets sourced from the router itself on the interface from which they were sent).
- Allowing the use of a default route with strict uRPF by using the command `ip verify unicast source reachable-via rx allow-default` . In this case, the router will match source IP addresses against the default-route entry (0.0.0.0/0) as well.
- Enabling SNMP traps for violating packets with the interface command `snmp trap ip verify drop-rate` and configuring notification thresholds with the command `ip verify unicast notification threshold <pps>` ; this assumes that SNMP is also correctly configured.

Configuration

```
R3:


access-list 101 deny ip any any log
!
interface FastEthernet0/1.13
```

https://t.me/learningnets

```
 ip verify unicast source reachable-via rx 101
```

Verification

Configure a Loopback on R1 with a /32 IP address from VLAN 23 range:

```
R1:

interface Loopback100
 ip address 136.1.23.1 255.255.255.255
```

Initiate ICMP packets sourced from this loopback address and notice packet drops on R3:

```
R1#ping 136.1.23.2 source loopback100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:
Packet sent with a source address of 136.1.23.1 .....
Success rate is 0 percent (0/5)
!
!R3#show ip interface fastEthernet0/1.13 | section IP verify
IP verify source reachable-via RX, ACL 101
5 verification drops
   0 suppressed verification drops
   0 verification drop-rate
!
!R3#show logging | i 101
%SEC-6-IPACCESSLOGDP: list 101 denied icmp 136.1.23.1 -> 136.1.23.2 (0/0), 1 packet
```

Remove the configuration used for testing:

```
R1:

no interface Loopback100
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

Preventing Fragment Attacks Using ACL

Task

- Configure R2 to drop any fragmented ICMP traffic received inbound on its VLAN 23 interface.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

When a packet is fragmented, you should be aware of the following:

https://t.me/learningnets

- The initial fragment is the first part of the packet. This fragment has the **M** (more) flag set, signaling that more fragments will follow. It also has the **FO** (fragment offset) value of zero, signaling the position of the fragment inside the original datagram. Usually, this initial fragment contains the upper-level protocol header (such as TCP/UDP) and thus bears the port number's information. Note that normal non-fragmented packets have FO=0, M=0.
- Non-Initial fragments are packets with non-zero **FO** field, meaning these fragments follow the initial one. These fragments also have the **M** bit set, with the exception of the final fragment. The final fragment has the **M** bit set to zero, signaling the end of the sequence. Non-initial fragments do not contain any upper-level protocol header; only layer 3 information is present.

Because many attacks are based on fragment packets, you may want to disallow such packets by using access-lists with the special keyword `fragments`. Because non-initial fragments do not contain layer 4 information, the matching rules for non-initial fragments are:

- If the packet matches on a deny statement containing layer 3/layer 4 information and the `fragment` keyword is present, the next ACL entry is being processed.
- If the packet matches on a permit statement containing layer 3/layer 4 information and the `fragment` keyword is present, the packet is allowed.

Configuration

```
R2:

ip access-list extended ICMP_FRAGMENTS
 deny icmp any any fragments
 permit ip any any
!
interface GigabitEthernet0/0
 ip access-group ICMP_FRAGMENTS in
```

Verification

Verify that fragmented ICMP packets are dropped by the access-lits, and check ACL statistics:

```
R1#ping 136.1.23.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:!!!!!
```

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R1#ping 136.1.23.2 size 1600
Type escape sequence to abort.
Sending 5, 1600-byte ICMP Echos to 136.1.23.2, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R2#show ip access-lists ICMP_FRAGMENTS
Extended IP access list ICMP_FRAGMENTS 10 deny icmp any any fragments (5 matches)

    20 permit ip any any (318 matches)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IPv6 RA Guard

Task

- Configure IPv6 RA Guard on SW3 for all current and future hosts in VLAN 11.
  - Allow R4's FE80::4 link-local address to send RA messages but only for 2001::/64 prefix.
  - R6 should not be affected by the above configurations.
- Configure R5 for SLAAC.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IPv6 Neighbor Discovery Protocol (NDP) can be seen as a nonauthenticated application running on top of ICMPv6; it uses link-local addresses for source (required to accomplish its role), and it is multicast based for efficiency. It is a replacement for ARP protocol from IPv4, but it has additional functionalities:

- IPv6 router discovery
- IPv6 address autoconfiguration, also known as stateless address autoconfiguration (SLAAC)
- IPv6 address resolution, which replaces Address Resolution Protocol (ARP)
- IPv6 neighbor reachability, also known as neighbor unreachability detection (NUD)

https://t.me/learningnets

- IPv6 duplicate address detection (DAD), which works for both link-local and global assigned through SLAAC
- IPv6 redirection, which is the ICMP redirect from IPv4

In IPv6 hosts can receive IP addresses by way of four methods:

- Manual/static configuration.
- Dynamic stateless autoconfiguration (through SLACC); this is similar to APIPA from IPv4, except that subnet is not fixed—it is routable and dynamically learned through RA (router advertisement) messages from available routers on the segment.
- Dynamic stateful DHCPv6 configuration; DHCPv6 is used for most IPv6 client configuration, except for default gateway.
- Dynamic stateless DHCPv6 configuration; SLAAC is used for IPv6 address configuration and DHCPv6 for other parameters, such as domain-name and DNS servers.

Several pieces of information can be found within RA messages, but the most important and relevant are:

- IPv6 global unicast prefixes (there can be one or more, required to be /64 for SLAAC).
- IPv6 global unicast address of the router (one for each advertised prefix).
- IPV6 SLAAC (autoconfiguration) flag, signaling hosts if the received information can be used for SLAAC.

For IPv6 /64 prefixes advertised with the SLAAC flag, hosts will autoconfigure with a unique IPv6 address (/64 host portion will be based on the MAC address and the **FFFE** string inserted in the middle) and set the default gateway to be the respective router. Based on this behavior, the segment is vulnerable to some attacks:

- Spoofed RA messages with false /64 prefixes for SLAAC; basically the prefix is spoofed.
- Spoofed RA messages with false advertiser information (router); basically the router is spoofed.

https://t.me/learningnets

To protect the infrastructure from spoofed RA messages, a feature named IPv6 RA Guard can be configured on the layer2/layer3 switch where hosts and routers are attached. Basically you would want to restrict/drop receiving router type NDP messages (like router advertisement and redirects) on ports where hosts are connected. You can configure three types of RA Guard policies:

- **host**, where all inbound NDP messages are inspected and router type messages are dropped; this is the default state of a policy.
- **router**, where all inbound NDP messages are inspected but router type messages are allowed; optionally, you can restrict for which prefixes are RA messages allowed (achieved via prefix-list filtering) and which sources are allowed to send ICMPv6 Type 133, 134, and 137 (achieved via access-lists matching on link-local addresses).
- **trust**, where all inbound NDP messages are allowed and not inspected.

The RA Guard policies can be applied at the VLAN level or port level; if both are enabled, the port-level policy will take precedence over the VLAN-wide-level policy. In most real-life cases, it is faster to deploy a host policy at VLAN level and router/trust policy at port level, because there will be few router/trusted ports. To enable a Cisco router for SLAAC, the interface-level command `ipv6 address autoconfig` is used. The interface-level command `ipv6 enable` is used when you want the router to autoconfigure itself with just a link-local address on the interface (which does not require SLAAC).

Configuration

SW3:

```
ipv6 prefix-list RA_PREFIX permit 2001::/64
!
ipv6 access-list RA_SOURCE
 permit ipv6 host FE80::4 any
!
ipv6 nd raguard policy HOST_INSPECT
 device-role host
!
ipv6 nd raguard policy HOST_NON_INSPECT
 trusted-port
!
ipv6 nd raguard policy ROUTER
 device-role router
 match ra prefix-list RA_PREFIX
 match ipv6 access-list RA_SOURCE
!
```

https://t.me/learningnets

```
vlan configuration 11
 ipv6 nd raguard attach-policy HOST_INSPECT
!
interface GigabitEthernet1/0/1
 ipv6 nd raguard attach-policy ROUTER
!
interface GigabitEthernet1/0/3
 ipv6 nd raguard attach-policy HOST_NON_INSPECT
```

R5:

```
interface FastEthernet0/0
 ipv6 address autoconfig default
```

Verification

Verify that R4 sends RA messages:

```
R4#show ipv6 interface fastEthernet0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::4
  No Virtual link-local address(es):
  Global unicast address(es):
    2001::4, subnet is 2001::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:2
    FF02::1:FF00:4
    FF05::1:3
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ICMP unreachables are sent
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds (using 30000)
  ND advertised reachable time is 0 (unspecified)
  ND advertised retransmit interval is 0 (unspecified)
 ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  ND advertised default router preference is Medium
  Hosts use stateless autoconfig for addresses.
!
!R4#show ipv6 interface fastEthernet0/0 prefix
```

https://t.me/learningnets

```
IPv6 Prefix Advertisements FastEthernet0/0
Codes: A - Address, P - Prefix-Advertisement, O - Pool
       U - Per-user prefix, D - Default
       N - Not advertised, C - Calendar

PD default [LA] Valid lifetime 2592000, preferred lifetime 604800
AD 2001::/64 [LA] Valid lifetime 2592000, preferred lifetime 604800
```

Verify that IPv6 RA Guard is enabled globally and policies are attached to both port and VLAN level:

```
SW3#show ipv6 snooping features
Feature name   priority state RA guard          192   READY
!
!SW3#show ipv6 snooping policies

Target              Type  Policy              Feature        Target range
Gi1/0/1             PORT  ROUTER              RA guard       vlan all
Gi1/0/3             PORT  HOST_NON_INSPECT    RA guard       vlan all
vlan 11             VLAN  HOST_INSPECT        RA guard       vlan all
```

Further verify that IPv6 RA Guard is enabled; note that ICMPv6 Type 133 (router solicitation), 134 (router advertisement), and 137 (neighbor redirect) are punted for inspection:

```
SW3#show ipv6 snooping capture-policy
HW Policy DB:
        HW Policy 0000001C #targets:3
                Target Gi1/0/1 type 0 handle 1
                Target Gi1/0/3 type 0 handle 3
                Target vlan 11   type 1 handle B0000
Target DB: HW Target Gi1/0/1 HW policy signature 0000001C policies#:1 rules#:3 sig 0000001C
 SW policy ROUTER feature RA guard

                Rule RS Protocol ICMPV6 mask 00000004 action PUNT match 133 #feat:1
                        feature RA guard
                Rule RA Protocol ICMPV6 mask 00000008 action PUNT match 134 #feat:1
                        feature RA guard
                Rule REDIR Protocol ICMPV6 mask 00000010 action PUNT match 137 #feat:1
                        feature RA guard
HW Target Gi1/0/3 HW policy signature 0000001C policies#:1 rules#:3 sig 0000001C
 SW policy HOST_NON_INSPECT feature RA guard

                Rule RS Protocol ICMPV6 mask 00000004 action PUNT match 133 #feat:1
```

https://t.me/learningnets

```
                        feature RA guard
                Rule RA Protocol ICMPV6 mask 00000008 action PUNT match 134 #feat:1
                        feature RA guard
                Rule REDIR Protocol ICMPV6 mask 00000010 action PUNT match 137 #feat:1
                        feature RA guard
HW Target vlan 11    HW policy signature 0000001C policies#:1 rules#:3 sig 0000001C
 SW policy HOST_INSPECT feature RA guard


                Rule RS Protocol ICMPV6 mask 00000004 action PUNT match 133 #feat:1
                        feature RA guard
                Rule RA Protocol ICMPV6 mask 00000008 action PUNT match 134 #feat:1
                        feature RA guard
                Rule REDIR Protocol ICMPV6 mask 00000010 action PUNT match 137 #feat:1
                        feature RA guard
```

Verify that R5 received RA messages from both R4 and R6 (because R6's port is configured as trusted so no RA Guard inspections are performed):

```
R5#show ipv6 routers
Router FE80::6 on FastEthernet0/0, last update 0 min
  Hops 64, Lifetime 1800 sec, AddrFlag=0, OtherFlag=0, MTU=1500
  HomeAgentFlag=0, Preference=Medium
  Reachable time 0 (unspecified), Retransmit time 0 (unspecified) Prefix 2001::/64 onlink autoconfig
    Valid lifetime 2592000, preferred lifetime 604800
Router FE80::4 on FastEthernet0/0, last update 0 min
  Hops 64, Lifetime 1800 sec, AddrFlag=0, OtherFlag=0, MTU=1500
  HomeAgentFlag=0, Preference=Medium
  Reachable time 0 (unspecified), Retransmit time 0 (unspecified) Prefix 2001::/64 onlink autoconfig
    Valid lifetime 2592000, preferred lifetime 604800
!
!R5#show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
FE80::4                                    11 0019.0653.2a18  STALE Fa0/0
FE80::6                                    11 0017.5aed.28b0  STALE Fa0/0
```

Verify that R5 has been automatically assigned an IPv6 global address through SLAAC; the default route has also been installed through SLAAC:

```
R5#show ipv6 interface brief fastEthernet0/0
FastEthernet0/0            [up/up]
    FE80::217:E0FF:FE4A:9678 2001::217:E0FF:FE4A:9678
!
!R5#show ipv6 route static
```

https://t.me/learningnets

```
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       D - EIGRP, EX - EIGRP external, NM - NEMO, ND - Neighbor Discovery
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S   ::/0 [2/0]
     via FE80::6, FastEthernet0/0
```

For testing, enable logging for packet drops and configure the port toward R6 in host mode, so that RA messages will be dropped:

```
SW3:

ipv6 snooping logging packet drop
!
interface GigabitEthernet1/0/3
 ipv6 nd raguard attach-policy HOST_INSPECT
```

On SW3's console, messages similar to the following will be logged (each 200 seconds), signaling the problem that R6's RA messages are now dropped:

```
%SISF-4-PAK_DROP: Message dropped A=FE80::6 G=- V=11 I=Gi1/0/3 P=NDP::RA Reason=Message unauthorized on port
%SISF-4-PAK_DROP: Message dropped A=FE80::6 G=- V=11 I=Gi1/0/3 P=NDP::RA Reason=Message unauthorized on port
%SISF-4-PAK_DROP: Message dropped A=FE80::6 G=- V=11 I=Gi1/0/3 P=NDP::RA Reason=Message unauthorized on port
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IPv6 DHCPv6 Guard

Task

- Configure IPv6 DHCPv6 Guard on SW3 for all current and future hosts in VLAN 11.
  - Allow R4's FE80::4 link-local address to act as DHCP server but only for 2001::/64 prefix.
  - R6 should not be affected by the above configurations.
- Activate DHCPv6 server on R4 for 2001::/64 subnet; assign the domain **inelab.local** and the DNS server 2001::4 with maximum lifetime.
- Configure R5 to receive an IPv6 address via DHCPv6.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

DHCPv6 Guard achieves almost the same thing for IPv6 as DHCP snooping did for IPv4. It blocks DHCP reply and advertisement messages that originate from unauthorized DHCP servers and relay agents that forward DHCP packets from servers to clients. It lacks the advanced functionalities of deep DHCP packet inspection to prevent DHCP starvation attacks, but this problem is solved with IPv6 snooping, with which you can limit the number of IPv6 addresses that a neighbor can use; it also does not support rate-limiting the DHCPv6 clients for the number of requests. You can configure three types of DHCPv6 Guard policies:

- **client**, where all inbound DHCP messages are inspected and DHCP server type

https://t.me/learningnets

messages are dropped; this is the default state of the policy.

- **server**, where all inbound DHCP messages are inspected but router typed messages are allowed; optionally, you can restrict for which prefixes can the router acts as a DHCP server (achieved via prefix-list filtering), and which sources are allowed to send DHCP server messages (achieved via access-lists matching on link-local addresses).
- **trust**, where all inbound DHCP messages are allowed and not inspected.

The DHCPv6 guard policies can be applied at the VLAN level or port level; if both are enabled, the port-level policy will take precedence over the VLAN-wide-level policy. In most real-life cases, it is faster to deploy a client policy at VLAN level and server policy at port level, because there will be few server ports. When DHCPv6 Guard is enabled on a VLAN, the switch will silently drop all RA messages on that VLAN and SLAAC will no longer function, so all hosts need to get IPv6 addresses statically or through DHCPv6. For RA messages to be allowed, IPv6 RA Guard must be activated on the VLAN as well.

DHCPv6 is seen in IPv6 as an addition to SLAAC, but in general SLAAC should be used; this is confirmed by the fact that DHCPv6 does not assign a default-router like DHCPv4 did, and this is not configurable; default-router should be automatically selected by the host through SLAAC (from RA messages).

Configuration

SW3:

```
ipv6 prefix-list DHCP_SCOPE permit 2001::/64 le 128
!
ipv6 access-list DHCP_SOURCE
 permit ipv6 host FE80::4 any
!
!
ipv6 dhcp guard policy DHCP_SERVER
 device-role server
 match server access-list DHCP_SOURCE
 match reply prefix-list DHCP_SCOPE
!
ipv6 dhcp guard policy DHCP_CLIENT_INSPECT
 device-role client
!
ipv6 dhcp guard policy DHCP_CLIENT_NON_INSPECT
 trusted-port
!
vlan configuration 11
 ipv6 dhcp guard attach-policy DHCP_CLIENT_INSPECT
```

https://t.me/learningnets

```
!
interface GigabitEthernet1/0/1
 ipv6 dhcp guard attach-policy DHCP_SERVER
!
interface GigabitEthernet1/0/3
 ipv6 dhcp guard attach-policy DHCP_CLIENT_NON_INSPECT
```

R4:

```
ipv6 dhcp pool VLAN11
 address prefix 2001::/64 lifetime infinite infinite
 link-address 2001::4/64
 dns-server 2001::4
 domain-name inelab.local
!
interface FastEthernet0/0
 ipv6 dhcp server VLAN11
```

R5:

```
interface FastEthernet0/0
 ipv6 address dhcp
 ipv6 enable
```

Verification

Verify DHCP server configuration:

```
R4#show ipv6 dhcp interface fastEthernet0/0
FastEthernet0/0 is in server mode
 Using pool: VLAN11
   Preference value: 0
   Hint from client: ignored
   Rapid-Commit: disabled
!
!R4#show ipv6 dhcp pool
DHCPv6 pool: VLAN11
Address allocation prefix: 2001::/64 valid 4294967295 preferred 4294967295 (1 in use, 0 conflicts)
   Link-address prefix: 2001::4/64 DNS server: 2001::4
 Domain name: inelab.local

   Active clients: 1
```

https://t.me/learningnets

Verify that R4 assigned an IPv6 address to R5:

```
R4#show ipv6 dhcp binding
Client: FE80::217:E0FF:FE4A:9678
  DUID: 000300010017E04A9678
  Username : unassigned
  IA NA: IA ID 0x00030001, T1 43200, T2 69120 Address: 2001::81EF:5B7D:5E78:A1EA
            preferred lifetime INFINITY, , valid lifetime INFINITY,
!
!R5#show ipv6 interface brief fastEthernet0/0

FastEthernet0/0            [up/up]
    FE80::217:E0FF:FE4A:9678 2001::81EF:5B7D:5E78:A1EA
!
!R5#show ipv6 dhcp interface fastEthernet0/0
FastEthernet0/0 is in client mode
  Prefix State is IDLE
  Address State is OPEN
  Renew for address will be sent in 11:30:25
  List of known servers:
    Reachable via address: FE80::4
    DUID: 00030001001906532A18
    Preference: 0
    Configuration parameters:
      IA NA: IA ID 0x00030001, T1 43200, T2 69120 Address: 2001::81EF:5B7D:5E78:A1EA/128
 preferred lifetime INFINITY, valid lifetime INFINITY
 DNS server: 2001::4
 Domain name: inelab.local

      Information refresh time: 0
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled
```

Verify that IPv6 DHCPv6 Guard is enabled globally and policies are attached to both port and VLAN level:

https://t.me/learningnets

```
SW3#show ipv6 snooping features

Feature name   priority state DHCP Guard        200   READY

!
!SW3#show ipv6 snooping policies

Target               Type  Policy               Feature        Target range
Gi1/0/1              PORT  DHCP_SERVER          DHCP Guard     vlan all
Gi1/0/3              PORT  DHCP_CLIENT_NON_INSP DHCP Guard     vlan all
vlan 11              VLAN  DHCP_CLIENT_INSPECT  DHCP Guard     vlan all
```

Further verify that IPv6 DHCPv6 Guard is enabled; note that DHCP client (546) and server (547) ports are punted for inspection:

```
SW3#show ipv6 dhcp guard policy
Dhcp guard policy: DHCP_CLIENT_INSPECT Device Role: dhcp client
Target: vlan 11

Dhcp guard policy: DHCP_CLIENT_NON_INSPECT Trusted Port
Target: Gi1/0/3

Dhcp guard policy: DHCP_SERVER Device Role: dhcp server
Target: Gi1/0/1
        Max Preference: 255
        Min Preference: 0
        Source Address Match Access List: DHCP_SOURCE
        Prefix List Match Prefix List: DHCP_SCOPE
!
!SW3#show ipv6 snooping capture-policy
HW Policy DB:
        HW Policy 00000180 #targets:3
                Target Gi1/0/1 type 0 handle 1
                Target Gi1/0/3 type 0 handle 3
                Target vlan 11   type 1 handle B0000
Target DB: HW Target Gi1/0/1 HW policy signature 00000180 policies#:1 rules#:2 sig 00000180
SW policy DHCP_SERVER feature DHCP Guard

                Rule DHCP SERVER Protocol UDP mask 00000100 action PUNT match 546 #feat:1
                        feature DHCP Guard
                Rule DHCP CLIENT Protocol UDP mask 00000080 action PUNT match 547 #feat:1
                        feature DHCP Guard
HW Target Gi1/0/3 HW policy signature 00000180 policies#:1 rules#:2 sig 00000180
```

https://t.me/learningnets

```
 SW policy DHCP_CLIENT_NON_INSPECT feature DHCP Guard


                Rule DHCP SERVER Protocol UDP mask 00000100 action PUNT match 546 #feat:1
                        feature DHCP Guard
                Rule DHCP CLIENT Protocol UDP mask 00000080 action PUNT match 547 #feat:1
                        feature DHCP Guard
 HW Target vlan 11    HW policy signature 00000180 policies#:1 rules#:2 sig 00000180
  SW policy DHCP_CLIENT_INSPECT feature DHCP Guard


                Rule DHCP SERVER Protocol UDP mask 00000100 action PUNT match 546 #feat:1
                        feature DHCP Guard
                Rule DHCP CLIENT Protocol UDP mask 00000080 action PUNT match 547 #feat:1
                        feature DHCP Guard
```

For testing, enable logging for packet drops and configure the port toward R6 in client mode, so that DHCPv6 server messages are dropped; enable DHCPv6 on R6:

```
R6:
ipv6 dhcp pool VLAN11
 address prefix 2001::/64 lifetime infinite infinite
 link-address 2001::4/64
 dns-server 2001::4
 domain-name inelab.local
!
interface FastEthernet0/0
 ipv6 dhcp server VLAN11

SW3:
ipv6 snooping logging packet drop
!
interface GigabitEthernet1/0/3
 ipv6 dhcp guard attach-policy DHCP_CLIENT_INSPECT

R5:

interface FastEthernet0/0
 no ipv6 address dhcp
 ipv6 address dhcp
```

On SW3's console, messages similar to the following will be logged, signaling the problem that R6's DHCPv6 server messages are now dropped:

https://t.me/learningnets

```
ISF-4-PAK_DROP: Message dropped A=FE80::6 G=2001::6DF2:303B:3A99:39C2 V=11 I=Gi1/0/3 P=DHCPv6::ADV Reason=Message ty
%SISF-4-PAK_DROP: Message dropped A=FE80::6 G=2001::38CB:3C76:9272:31EC V=11 I=Gi1/0/3 P=DHCPv6::ADV Reason=Message
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IPv6 Snooping

Task

- Configure IPv6 snooping on SW3 for all current and future hosts in VLAN 11.
  - Allow only R4 to send RA messages and act as DHCPv6 server but enforce message validation.
  - Ensure that no more than two IPv6 addresses are allowed for each host.
  - Log all packets violating the policy and neighbor table additions.
- Protect against non-link-local IPv6 address spoofing, and also actively monitor neighbor states.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IPv6 snooping bundles several layer2 IPv6 first-hop security features while automatically enabling RA Guard and DHCPv6 Guard if configured appropriately:

- IPv6 neighbor discovery inspection
- IPv6 address/data gleaning
- IPv6 device tracking

Through IPv6 neighbor discovery (ND) inspection, it verifies the sanity of ND messages and can enforce limits on the number of IPv6 addresses a neighbor can use; it relies heavily on the SLAAC process. Only valid ND messages will be

https://t.me/learningnets

allowed and entered in the binding table (a valid ND message is considered trustworthy if its IPv6-to-MAC mapping is verifiable). This feature mitigates some of the inherent vulnerabilities for the neighbor discovery mechanism, such as attacks on duplicate address detection (DAD), address resolution, device discovery, and the neighbor cache.

Through IPv6 address gleaning, it examines the ND and DHCPv6 packets to learn the IPv6 addresses of connected devices. Learned addresses are installed into a binding table and are used for other features, like IPv6 Source Guard; each binding contains the source of the learning (NDP, DHCPv6), the IPv6 address, the MAC address, the interface and VLAN, age of the entry, current neighbor state, and time left. To help in cases where a switch rebooted and lost its binding table, a feature known as data-gleaning (an extension to the IPv6 address gleaning) will punt data packets from unknown sources to be inspected by the switch, which will actively send DAD Neighbor Solicitation messages to verify the IPv6 address. If the host truly has the IPv6 address, it will reply with a Neighbor Advertisement message, which will allow the switch to enter the mapping in its binding table. This is recommended to remain disabled (which is by default) because it has high impact on switch CPU; it can be activated in the snooping policy using the command `data-glean`. A policy with data gleaning enabled cannot be activated at the VLAN level, only at the interface level.

IPv6 device tracking allows the switch to maintain the consistency of the gleaned address entries (an extension to the IPv6 address gleaning). Stale entries are verified using the NS-NA DAD exchange and cleaned up (aged out) upon failure. This ensures that entries in the binding table are accurate and can thus enhance the IPv6 Source Guard feature. This is also CPU intensive because the switch needs to regularly initiate NS-NA messages for each entry in its binding table. It is disabled by default and can be activated in the snooping policy using the command `tracking enable`, or globally using the command `ipv6 neighbor tracking`.

By default, both DHCPv6 and NDP messages are inspected, but behavior varies based on the security-level being configured:

- **Glean** just populates the binding table from NDP/DHCP messages.
- **Inspect** populates the binding table from NDP/DHCP messages, but also validates messages to be conforming and enforces address ownership (when an address binding is performed, it does not allow another host to claim the same IP address).
- **Guard** is an enhancement of **Inspect** mode; it additionally enables DHCP Guard and RA Guard (it drops DHCP Server messages and Router RA messages).

IPv6 Source Guard is very similar with IPv4 Source Guard. When applied on a layer 2 port, the IPv6 policy will apply a transparent IPv6 PACL on the port, in the inbound direction, allowing only traffic from IPv6 hosts found in the switch IPv6 neighbor

https://t.me/learningnets

binding table (which can be populated dynamically by the switch through NDP message inspection or manually by the administrator by configuring entries).

Configuration

SW3:

```
ipv6 snooping policy VLAN_SNOOPING
 device-role node
 protocol dhcp
 protocol ndp
 limit address-count 2
 security-level guard
 tracking enable
!
ipv6 snooping policy R4_SNOOPING
 device-role node
 protocol dhcp
 protocol ndp
 limit address-count 2
 security-level inspect
 tracking enable
!
ipv6 source-guard policy PROTECT
 permit link-local
!
vlan configuration 11
 ipv6 snooping attach-policy VLAN_SNOOPING
!
interface GigabitEthernet1/0/1
 ipv6 snooping attach-policy R4_SNOOPING
!
interface range gigabitEthernet 1/0/1 - 3
 ipv6 source-guard attach-policy PROTECT
!
ipv6 snooping logging packet drop
ipv6 neighbor binding logging
ipv6 neighbor tracking
```

Verification

Verify that IPv6 Snooping and IPv6 Source Guard is enabled globally and policies

https://t.me/learningnets

are attached to both port and VLAN level:

```
SW3#show ipv6 snooping features
Feature name   priority state Snooping          128   READY
Source guard       32   READY
!
!SW3#show ipv6 snooping policies
Target               Type  Policy               Feature        Target range
Gi1/0/1              PORT  R4_SNOOPING          Snooping       vlan all
Gi1/0/1              PORT  PROTECT              Source guard   vlan all
Gi1/0/2              PORT  PROTECT              Source guard   vlan all
Gi1/0/3              PORT  PROTECT              Source guard   vlan all
vlan 11              VLAN  VLAN_SNOOPING        Snooping       vlan all
!
!SW3#show ipv6 snooping capture-policy interface gigabitEthernet1/0/1
        HW Target Gi1/0/1 HW policy signature 0000019F policies#:2 rules 7 sig 0000019F
SW policy R4_SNOOPING feature Snooping
 SW policy PROTECT feature Source guard


                Rule DHCP SERVER Protocol UDP mask 00000100 action PUNT match 546 #feat:1
                        feature Snooping
                Rule DHCP CLIENT Protocol UDP mask 00000080 action PUNT match 547 #feat:1
                        feature Snooping
                Rule RS Protocol ICMPV6 mask 00000004 action PUNT match 133 #feat:1
                        feature Snooping
                Rule RA Protocol ICMPV6 mask 00000008 action PUNT match 134 #feat:1
                        feature Snooping
                Rule NS Protocol ICMPV6 mask 00000001 action PUNT match 135 #feat:1
                        feature Snooping
                Rule NA Protocol ICMPV6 mask 00000002 action PUNT match 136 #feat:1
                        feature Snooping
                Rule REDIR Protocol ICMPV6 mask 00000010 action PUNT match 137 #feat:1
                        feature Snooping
```

Test connectivity between R4 and R6, and verify that entries for R4 and R6 are present in the neighbor binding table of SW3:

```
R4#ping 2001::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::6, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
!
!R4#ping FE80::6
Output Interface: FastEthernet0/0
```

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::6, timeout is 2 seconds:
Packet sent with a source address of FE80::4%FastEthernet0/0!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
!
!SW3#show ipv6 neighbors binding

Binding Table has 4 entries, 4 dynamic
Codes: L - Local, S - Static, ND - Neighbor Discovery, DH - DHCP, PKT - Other Packet, API - API created
Preflevel flags (prlvl):
0001:MAC and LLA match     0002:Orig trunk            0004:Orig access
0008:Orig trusted trunk    0010:Orig trusted access   0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated    0100:Statically assigned


    IPv6 address                            Link-Layer addr Interface vlan prlvl  age   state    Time left
ND  FE80::6                                 0017.5AED.28B0  Gi1/0/3     11  0005    4mn REACHABLE  49 s try 0
ND  FE80::4                                 0019.0653.2A18  Gi1/0/1     11  0005   94s  REACHABLE  218 s
ND  2001::6                                 0017.5AED.28B0  Gi1/0/3     11  0005    4mn REACHABLE  21 s try 0
ND  2001::4                                 0019.0653.2A18  Gi1/0/1     11  0005    3mn REACHABLE  97 s try 0
```

When neighbor binding table entries are created, messages similar to the following should be logged on SW3 because **`ipv6 neighbor binding logging`** is enabled:

```
%SISF-6-ENTRY_CHANGED: Entry changed A=FE80::4 V=11 I=Gi1/0/1 P=0005 M=0019.0653.2A18
%SISF-6-ENTRY_CHANGED: Entry changed A=2001::4 V=11 I=Gi1/0/1 P=0005 M=0019.0653.2A18
```

Because R6 will send RA messages, but the applied snooping policy does not allow this, messages similar to the following should be logged on SW3 because **`ipv6 snooping logging packet drop`** is enabled:

```
%SISF-4-PAK_DROP: Message dropped A=FE80::6 G=- V=11 I=Gi1/0/3 P=NDP::RA Reason=Packet not authorized on port
%SISF-4-PAK_DROP: Message dropped A=FE80::6 G=- V=11 I=Gi1/0/3 P=NDP::RA Reason=Packet not authorized on port
```

If R6 will be configured as a DHCPv6 server, based on applied snooping policy, messages similar to the following should be logged on SW3 because **`ipv6 snooping logging packet drop`** is enabled:

https://t.me/learningnets

```
%SISF-4-PAK_DROP: Message dropped A=FE80::4 G=2001::6C5F:714D:4F74:443C V=11 I=Gi1/0/1 P=DHCPv6::ADV Reason=Packet r
%SISF-4-PAK_DROP: Message dropped A=FE80::4 G=2001::6C5F:714D:4F74:443C V=11 I=Gi1/0/1 P=DHCPv6::ADV Reason=Packet r
```

To test that R4 is allowed to function as a DHCPv6 server, enable R6 as a DHCPv6 server and R5 as a DHCPv6 client:

```
R4:
ipv6 dhcp pool VLAN11
 address prefix 2001::/64 lifetime infinite infinite
 link-address 2001::4/64
 dns-server 2001::4
 domain-name inelab.local
!
interface FastEthernet0/0
 ipv6 dhcp server VLAN11

R5:

interface FastEthernet0/0
 ipv6 address dhcp
 ipv6 enable
```

Verify that R5 has been assigned an IPv6 address:

```
R5#show ipv6 interface brief
FastEthernet0/0            [up/up]
    FE80::217:E0FF:FE4A:9678 2001::458:5F16:A117:2F5F
!
!R4#show ipv6 dhcp binding
Client: FE80::217:E0FF:FE4A:9678
  DUID: 000300010017E04A9678
  Username : unassigned
  IA NA: IA ID 0x00030001, T1 43200, T2 69120 Address: 2001::458:5F16:A117:2F5F

            preferred lifetime INFINITY, , valid lifetime INFINITY,
```

To test the limit of IPv6 addresses per port (which also accounts for link-local addresses), add another IPv6 address on R6:

```
R6:

interface FastEthernet0/0
```

https://t.me/learningnets

```
 ipv6 address 2001::66/64
```

Try connectivity with the newly configured IPv6 address from R4; it should not be functional and messages similar to the following should be logged on SW3:

```
R4#ping 2001::66

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::66, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!
%SISF-4-PAK_DROP: Message dropped A=:: G=2001::66 V=11 I=Gi1/0/3 P=NDP::NS Reason=Address limit per policy reached
%SISF-4-PAK_DROP: Message dropped A=2001::66 G=- V=11 I=Gi1/0/3 P=NDP::NA Reason=Address limit per policy reached
```

To test that address ownership is protected for a host connected on the network, move R6's IPv6 address to R4:

```
R6:
interface FastEthernet0/0
 no ipv6 address 2001::6/64

R4:

interface FastEthernet0/0
 ipv6 address 2001::6/64
```

When R4 tries to spoof R6's IPv6 address, which is still in the neighbor binding table of SW3, the following log message will be issued by SW3, signaling the problem:

```
%SISF-4-PAK_DROP: Message dropped A=2001::4 G=- V=11 I=Gi1/0/3 P=NDP::NS Reason=More trusted entry exists
!
!SW3#show ipv6 snooping counters interface gigabitEthernet1/0/1
Received messages on Gi1/0/1:
Protocol        Protocol message
NDP             RA[43] NS[72] NA[44]
DHCPv6

Bridged messages from Gi1/0/1:
Protocol        Protocol message
NDP             RA[43] NS[86] NA[27]
DHCPv6
 Dropped messages on Gi1/0/1:
Feature         Protocol Msg [Total dropped]
Snooping        NDP      NS  [5]
```

https://t.me/learningnets

```
reason:  Address limit per policy reached [2]

                reason:  Packet accepted but not forwarded [3]
 NA  [4]
reason:  Address limit per policy reached [2]
reason:  More trusted entry exists [1]

                reason:  Packet accepted but not forwarded [1]
```

To test the IPv6 source guard functionality, because due to IPV6 snooping policy SW3 will not allow R4 and R6 to allocate new IPv6 addresses, we'll add on R4 two additional IPv6 addresses (one link-local and one global unicast) and manually configure the IPv6-to-MAC mappings on both R4 and R6:

```
R4#show ipv6 neighbors

IPv6 Address                              Age Link-layer Addr State Interface
2001::6                                     0 0017.5aed.28b0  REACH Fa0/0
FE80::6                                     0 0017.5aed.28b0  REACH Fa0/0
!
!R6#show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
2001::4                                     0 0019.0653.2a18  REACH Fa0/0
FE80::4                                     0 0019.0653.2a18  REACH Fa0/0


R6:
ipv6 neighbor FE80::44 fastEthernet0/0 0019.0653.2a18
ipv6 neighbor 2001::44 fastEthernet0/0 0019.0653.2a18


R4:

interface FastEthernet0/0
 ipv6 address 2001::44/64
 ipv6 address FE80::44 link-local
 no ipv6 address 2001::4/64
!
!
ipv6 neighbor FE80::6 fastEthernet0/0 0017.5aed.28b0
ipv6 neighbor 2001::6 fastEthernet0/0 0017.5aed.28b0
```

Based on the configured IPv6 source guard policy, SW3 should allow traffic from unknown (not present in its neighbor binding able) link-local addresses, but not from

https://t.me/learningnets

unknown global unicast addresses:

```
R6#ping FE80::44
Output Interface: FastEthernet0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FE80::44, timeout is 2 seconds:
Packet sent with a source address of FE80::6%FastEthernet0/0!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
!
!R6#ping 2001::44

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::44, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!SW3#show ipv6 neighbors binding

Binding Table has 4 entries, 4 dynamic
Codes: L - Local, S - Static, ND - Neighbor Discovery, DH - DHCP, PKT - Other Packet, API - API created
Preflevel flags (prlvl):
0001:MAC and LLA match     0002:Orig trunk            0004:Orig access
0008:Orig trusted trunk    0010:Orig trusted access   0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated    0100:Statically assigned


    IPv6 address                            Link-Layer addr Interface vlan prlvl  age   state    Time left
ND  FE80::6                                 0017.5AED.28B0  Gi1/0/3     11  0005    4mn REACHABLE  42 s try 0
ND  FE80::4                                 0019.0653.2A18  Gi1/0/1     11  0005   92s  REACHABLE  211 s try 0
ND  2001::6                                 0017.5AED.28B0  Gi1/0/3     11  0005    4mn REACHABLE  29 s try 0
ND  2001::4                                 0019.0653.2A18  Gi1/0/1     11  0005   90s  REACHABLE  224 s
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 2: Threat Identification and Mitigation

IPv6 Secure Neighbor Discovery

Task

- Configure IPv6 Secure Neighbor Discovery (SeND) over VLAN 11 as follows:
  - Configure R4 as the CA and NTP server.
  - Configure R5 and R6 with CGA link-local and unicast global addresses from the 2001::/64 range.
  - R5 and R6 should reject any non-SeND messages on the link and use the highest security level available.
  - For replay attacks, threat messages within a 60-second window as valid.
  - Configure and enforce a minimum RSA key length of 1024.
  - R5 will be the IPv6 host, and R4 and R6 will be the IPv6 routers on the segment.

> **About Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 2 Introduction** task or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

To prevent ARP spoofing in IPv4, DAI (Dynamic ARP Inspection) was developed. IPv6 uses NDP (Neighbor Discovery Protocol) for layer3-to-layer2 resolution (through Neighbor Solicitation and Neighbor Advertisement) and for router discovery (through Router Advertisement), messages that are vulnerable to spoofing; we can say that ARP spoofing is now NDP spoofing, and for this reason SeND (Secure

https://t.me/learningnets

Neighbor Discovery) was developed. For securing NDP messages with SeND, the following components are required:

- PKI infrastructure, with CA hierarchy trusted by all hosts and identity certificates issued to routers.
- IPv6 addresses (link-local and unicast global) on hosts and routers are automatically generated and cryptographically secured, also known as CGA (Cryptographically Generated Addresses).
- Self-generated RSA key-pair on hosts and routers so that all NDP messages are signed with the private key.
- Timestamp and nonce options, used to prevent replay attacks.

CGA addresses requires the IPv6 subnet prefix to be 64 bits in length, and the remaining 64 bits (also known as interface identifier) are automatically generated from a cryptographic hash of the RSA public key and additional parameters (modifier, which is a random value of 128 bits in length, subnet prefix and collision counter, 8 bits in length and used for DAD to ensure the uniqueness of the address). Both routers and hosts must generate an RSA key-pair that is further used for CGA address. NS and NA messages are modified as follows:

- A host/router initiates a SeND NS (Neighbor Solicitation) message that is sourced from CGA link-local address and signed with the RSA private-key; within the NS message, CGA parameters are also sent: nonce value, timestamp, and RSA public-key.
- The receiving host/router will first validate the SeND NS message signature using the received RSA public-key.
- A host/router will respond with a SeND NA (Neighbor Advertisement) message that is sourced from CGA link-local address and signed with the RSA private-key; within the NA message, CGA parameters are also sent: nonce value, timestamp, and RSA public-key.
- The receiving host/router will first validate the SeND NA message signature using the received RSA public-key.
- At this point, layer3-to-layer2 resolution is successful.

Because RA messages are used for SLAAC and default gateway selection, these must be further secured through an external authority, so routers must be enrolled in a PKI that is trusted by all hosts on the segment.

RS and RA messages are modified as follows:

- A host initiates a SeND RS (Router Solicitation) message that is sourced from CGA link-local address and signed with the RSA private-key; within the RS message, CGA parameters are also sent: nonce value, timestamp, and RSA public-key.
- The receiving router will first validate the SeND RS message signature using the

https://t.me/learningnets

received RSA public-key.

- A router will respond with a SeND RA (Router Advertisement) message that is sourced from CGA link-local address and signed with the RSA private-key; within the RA message, CGA parameters are also sent: nonce value, timestamp, and RSA public-key.
- The receiving host will first validate the SeND RA message signature using the received RSA public-key, but RA messages are not yet accepted.
- The receiving host will perform a lookup for trusted Certificate Authorities, and if it finds any, it will generate a SeND CPS (Certification Path Solicitation) message to the router; otherwise the RA message is dropped.
- The receiving router will reply with a SeND CPA (Certification Path Answer) message, basically replying with the certificate chain (CA certificate and router identity certificate).
- The receiving host will validate the initial SeND RA message at this point, if the router-presented certificate is issued by a trusted CA and is valid (based on the revocation configuration).

If the command `ipv6 nd secured full-secure` is configured, routers no longer accept non-SeND messages, so all hosts and routers on the segment need SeND implemented; otherwise SeND devices will not be able to communicate with non-SeND devices. If the command `ipv6 nd secured key-length minimum 1024` is configured, the RSA key length from CGA parameters must be at least 1024 bits in length; otherwise the SeND messages are dropped. The security parameter, configured when CGA hash is generated with the command `ipv6 cga modifier rsakeypair <name> sec-level <0-3>` , determines the security level of the generated address against brute-force attacks. Although theoretically it can have a value between 0 and 7 (it's a 3-bit value), because of computational requirements for its generation, currently only values 0 and 1 are supported by Cisco devices; increasing the security value by 1 adds 16 bits to the length of the hash that the attacker must break.

Configuration

R4:

```
ip http server
crypto key generate rsa general-keys label SEND_KEY modulus 1024
!
crypto pki server SEND_PKI
 issuer-name CN=R4.inelab.local
 grant auto
!
crypto pki trustpoint SEND_PKI
 rsakeypair SEND_KEY
```

https://t.me/learningnets

```
!
crypto pki server SEND_PKI
 no shutdown
!
ntp master
```

R5:

```
ipv6 cga modifier rsakeypair SEND_KEY sec-level 1
!
crypto pki trustpoint SEND_PKI
 enrollment url http://136.1.11.4
 revocation-check none
!
crypto pki authenticate SEND_PKI
!
ipv6 nd secured sec-level minimum 1
ipv6 nd secured full-secure
ipv6 nd secured key-length minimum 1024
!
interface FastEthernet0/0
 ipv6 cga rsakeypair SEND_KEY
 ipv6 address FE80:: link-local cga
 ipv6 address 2001::/64 cga
 ipv6 nd secured trustanchor SEND_PKI
 ipv6 nd secured timestamp delta 60
!
ntp server 136.1.11.4
```

R6:

```
ipv6 unicast-routing
ipv6 cga modifier rsakeypair SEND_KEY sec-level 1
!
crypto pki trustpoint SEND_PKI
 subject-name CN=R6.inelab.local
 rsakeypair SEND_KEY
 enrollment url http://136.1.11.4
!
crypto pki authenticate SEND_PKI
crypto pki enroll SEND_PKI
!
ipv6 nd secured sec-level minimum 1
ipv6 nd secured full-secure
ipv6 nd secured key-length minimum 1024
```

https://t.me/learningnets

```
!
interface FastEthernet0/0
 no ipv6 address FE80::6 link-local
 no ipv6 address 2001::6/64
 ipv6 cga rsakeypair SEND_KEY
 ipv6 address FE80:: link-local cga
 ipv6 address 2001::/64 cga
 ipv6 nd secured trustanchor SEND_PKI
 ipv6 nd secured timestamp delta 60
!
ntp server 136.1.11.4
```

Verification

Verify that both R5 and R6 have synchronized time with R4:

```
R5#show ntp status
 Clock is synchronized, stratum 9, reference is 136.1.11.4
nominal freq is 250.0000 Hz, actual freq is 249.9999 Hz, precision is 2**24
reference time is D6B9D7A8.A6857F5E (15:26:32.650 UTC Thu Feb 27 2014)
clock offset is 13.2909 msec, root delay is 2.19 msec
root dispersion is 22.88 msec, peer dispersion is 5.20 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000294 s/s
system poll interval is 64, last update was 466 sec ago.
!
!R6#show ntp status
 Clock is synchronized, stratum 9, reference is 136.1.11.4

nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**24
reference time is D6B9D97F.EED35E52 (15:34:23.932 UTC Thu Feb 27 2014)
clock offset is -3.7775 msec, root delay is 2.23 msec
root dispersion is 8.87 msec, peer dispersion is 4.55 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is -0.000000130 s/s
system poll interval is 64, last update was 4 sec ago.
```

Verify that R5 has successfully imported the CA certificate to validate RA messages from R6, and that R6 has successfully enrolled with the CA so that it can respond to CPS messages:

```
R5#show crypto pki certificates
 CA Certificate
 Status: Available
  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
```

https://t.me/learningnets

```
  Issuer:
    cn=R4.inelab.local
  Subject:
    cn=R4.inelab.local
  Validity Date:
    start date: 14:26:44 UTC Feb 27 2014
    end   date: 14:26:44 UTC Feb 26 2017
  Associated Trustpoints: SEND_PKI
!
!R6#show crypto pki certificates
Certificate
 Status: Available
  Certificate Serial Number (hex): 04
  Certificate Usage: General Purpose
  Issuer:
    cn=R4.inelab.local
  Subject:
    Name: R6.ine.com
    hostname=R6.ine.com
    cn=R6.inelab.local
  Validity Date:
    start date: 15:38:32 UTC Feb 27 2014
    end   date: 15:38:32 UTC Feb 27 2015
  Associated Trustpoints: SEND_PKI
 CA Certificate
 Status: Available

  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
  Issuer:
    cn=R4.inelab.local
  Subject:
    cn=R4.inelab.local
  Validity Date:
    start date: 14:26:44 UTC Feb 27 2014
    end   date: 14:26:44 UTC Feb 26 2017
  Associated Trustpoints: SEND_PKI
```

Verify that both R5 and R6 have generated CGA addresses:

```
R5#show ipv6 cga address-db

2001::/64 ::3CC6:E16D:6E49:C346 - table 0x0
        interface:      FastEthernet0/0 (3)
        modifier:       SEND_KEY
```

https://t.me/learningnets

```
        collisions:     0
FE80::/64 ::242F:1A27:B845:9B68 - table 0x12000003
        interface:      FastEthernet0/0 (3)
        modifier:       SEND_KEY
        collisions:     0
!
!R6#show ipv6 cga address-db

2001::/64 ::2021:4516:BD70:7B8F - table 0x0
        interface:      FastEthernet0/0 (3)
        modifier:       SEND_KEY
        collisions:     0
FE80::/64 ::382E:ADD1:BD9A:934D - table 0x12000003
        interface:      FastEthernet0/0 (3)
        modifier:       SEND_KEY
        collisions:     0
```

Verify that both R5 and R6 have generated the modifier, which is used as a CGA parameter, so it is binded to the CGA addresses:

```
R5#show ipv6 cga modifier-db
 DA37:3BC5:B6A1:68B:29A4:4DDE:D8D2:F051
        label:          SEND_KEY
        sec level:      1
        Addresses:
                2001::3CC6:E16D:6E49:C346
                FE80::242F:1A27:B845:9B68
!
!R6#show ipv6 cga modifier-db
 22AE:733A:E6E:86D0:D950:898E:BEDC:4B3

        label:          SEND_KEY
        sec level:      1
        Addresses:
                2001::2021:4516:BD70:7B8F
                FE80::382E:ADD1:BD9A:934D
```

Verify the timestamp configuration on R5 and R6:

```
R5#show ipv6 nd secured timestamp-db

Total number of entries: 2
Number of unreached peer entries: 2 / 1024
 FE80::382E:ADD1:BD9A:934D on FastEthernet0/0, delta 60s, fuzz 1000ms
```

https://t.me/learningnets

```
        Time to expire: 3h 59m 52s (unreached)
        TSlast: 0x530F5F5FB415 = 15:53:03 UTC Feb 27 2014
        RDlast: 0x530F5F5FC447 = 15:53:03 UTC Feb 27 2014
2001::2021:4516:BD70:7B8F on FastEthernet0/0, delta 60s, fuzz 1000ms
        Time to expire: 3h 59m 57s (unreached)
        TSlast: 0x530F5F608B11 = 15:53:04 UTC Feb 27 2014
        RDlast: 0x530F5F609B58 = 15:53:04 UTC Feb 27 2014
!
!R6#show ipv6 nd secured timestamp-db

Total number of entries: 2
Number of unreached peer entries: 1 / 1024
 FE80::242F:1A27:B845:9B68 on FastEthernet0/0, delta 60s, fuzz 1000ms
        Time to expire: 3h 45m 54s (reached)
        TSlast: 0x530F5C3BF72D = 15:39:39 UTC Feb 27 2014
        RDlast: 0x530F5C3C1039 = 15:39:40 UTC Feb 27 2014
2001::3CC6:E16D:6E49:C346 on FastEthernet0/0, delta 60s, fuzz 1000ms

        Time to expire: 2h 56m 10s (unreached)
        TSlast: 0x530F54C65005 = 15:07:50 UTC Feb 27 2014
        RDlast: 0x530F54C6693F = 15:07:50 UTC Feb 27 2014
```

Verify that R5 has requested and received the certificate of R6, so RA messages of R6 are accepted:

```
R5#show ipv6 nd secured certificates

Total number of entries: 1 / 32

Hash                             id         RA  certcnt certrcv state
03314F62BE1CBF607E5D957D8B2BBC0B 0x00005BFC no  1       1       CERT_VALIDATED
  certificate No 0
         subject  hostname=R6.ine.com,cn=R6.inelab.local
         issuer  cn=R4.inelab.local
!
!R5#show ipv6 routers
 Router FE80::382E:ADD1:BD9A:934D on FastEthernet0/0, last update 3 min
  Hops 64, Lifetime 1800 sec, AddrFlag=0, OtherFlag=0, MTU=1500
  HomeAgentFlag=0, Preference=Medium
  Reachable time 0 (unspecified), Retransmit time 0 (unspecified) Prefix 2001::/64 onlink autoconfig

    Valid lifetime 2592000, preferred lifetime 604800
```

Verify that R5 and R6 have IPv6 connectivity, and with R4 does not, because R4's

https://t.me/learningnets

NDP messages are not secure and are therefore dropped by both R5 and R6, so layer3-to-layer2 resolution does not work for R4:

```
R5#ping 2001::3CC6:E16D:6E49:C346

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::3CC6:E16D:6E49:C346, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
!
!R5#ping 2001::4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001::4, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R5#show ipv6 neighbors

IPv6 Address                              Age Link-layer Addr State Interface
FE80::382E:ADD1:BD9A:934D                   8 0017.5aed.28b0  STALE Fa0/0
!
!R6#show ipv6 neighbors

IPv6 Address                              Age Link-layer Addr State Interface
FE80::242F:1A27:B845:9B68                   8 0017.e04a.9678  STALE Fa0/0
```

Activate NDP debugging to verify that NDP messages from R4 are being dropped because they are not secured:

```
R5#debug ipv6 nd
 R5#debug ipv6 nd secured

 SEND: Receive: ND_ROUTER_ADVERT
SEND:    src FE80::4
SEND:    dst FF02::1
SEND:    Received at: 0x530F4C43951C = 14:31:31 UTC Feb 27 2014
SEND:            option 1 len 8: ND_OPT_SOURCE_LINKADDR
SEND:            option 5 len 8: ND_OPT_MTU
SEND:            option 3 len 32: ND_OPT_PREFIX_INFORMATION SEND: !RA without CGA option
SEND: Unsecure message discarded
SEND: ! DROP: ND_ROUTER_ADVERT  src FE80::4 dst FF02::1 reason=2
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

Section 3 Introduction

> **This section utilizes Firewall Services, Security Services, and Web Security.**

Lab Topics

This section is divided into two areas, Intrusion Prevention and Content Security. The Cisco IPS can provide some content security functions, but the Cisco WSA is a dedicated Content Security solution offering a complete set of features. The CCIE Security Lab exam will focus on the Web Security aspect of the WSA where L4TM (Traffic Monitor) is not included. The following high-level topics are provided by Cisco as *blueprint* topics:

- IPS 4200 Series Sensor Appliance
  - Initialize the Sensor Appliance
  - Sensor Appliance management
  - Virtual Sensors on the Sensor Appliance
  - Implementing security policies
  - Promiscuous and inline monitoring on the Sensor Appliance
  - Tune signatures on the Sensor Appliance
  - Custom signatures on the Sensor Appliance
  - Actions on the Sensor Appliance
  - Signature engines on the Sensor Appliance
  - Use IDM/IME to the Sensor Appliance
  - Event action overrides/filters on the Sensor Appliance
  - Event monitoring on the Sensor Appliance
  - VACL/SPAN & RSPAN on Cisco switches

- WSA
  - Implementing WCCP

https://t.me/learningnets

- Active Dir Integration
- Custom Categories
- HTTPS Config
- Services Configuration (Web Reputation)
- Configuring Proxy By-pass Lists
- Web proxy modes
- App visibility and control

As with previous topics, we can expand upon this blueprint with the following checklist:

Technology	Mastered?
Understanding Cisco IPS System Architecture (System Design, MainApp, SensorApp, EventStore)	
Understanding Cisco IPS User Roles (Administrator, Operator, Viewer, Service)	
Understanding Cisco IPS Command Modes (Privileged, Global, Service, Multi-Instance)	
Understanding Cisco IPS Interfaces (Command and Control, Sensing, Alternate TCP Reset)	
Understanding Promiscuous (IDS) vs. Inline (IPS) Monitoring	
Initialization Basic Sensor (IP Address, Mask, Default Route, etc.)	
Troubleshooting Basic Connectivity Issues	
Managing Sensor ACLs	

https://t.me/learningnets

Technology	Mastered?
Allowing Services Ping and Telnet from/to Cisco IPS	
Enabling Physical Interfaces	
Promiscuous Mode	
Inline Interface Mode	
Inline VLAN Pair Mode	
VLAN Group Mode	
Inline Bypass Mode	
Interface Notifications	
Understanding the Analysis Engine	
Creating Multiple Security Policies and Applying Them to Individual Virtual Sensors	
Understanding and Configuring Virtual Sensors (vs0, vs1)	
Assigning Interfaces to the Virtual Sensor	
Understanding and Configuring Event Action Rules (rules0, rules1)	
Understanding and Configuring Signatures (sig0, sig1)	

https://t.me/learningnets

Technology	Mastered?
Adding Signatures to Multiple Virtual Sensors	
Understanding and Configuring Anomaly Detection (ad0, ad1)	
Using the Cisco IDM (IPS Device Manager)	
Using Cisco IDM Event Monitoring	
Displaying Events Triggered Using the Cisco IPS Console	
Troubleshooting Events Not Triggering	
Displaying and Capturing Live Traffic on the Cisco IPS Console (Packet Display and Packet Capture)	
SPAN and RSPAN	
Rate Limiting	
Configuring Event Action Variables	
Target Value Ratings	
Event Action Overrides	
Event Action Filters	
Configuring General Settings	

https://t.me/learningnets

Technology	Mastered?
General Signature Parameters	
Alert Frequency	
Alert Severity	
Event Counter	
Signature Fidelity Rating	
Signature Status	
Assigning Actions to Signatures	
AIC Signatures	
IP Fragment Reassembly	
TCP Stream Reassembly	
IP Logging	
Configuring SNMP	
Signature Tuning (Severity Levels, Throttle Parameters, Event Actions)	
Creating Custom Signatures (Using the CLI and Cisco IDM)	
Understanding Various Types of Signature Engines	

https://t.me/learningnets

Technology	Mastered?
Understanding Various Types of Signature Variables	
Understanding Various Types of Event Actions	
Creating a Custom String TCP Signature	
Creating a Custom Flood Engine Signature	
Creating a Custom AIC MIME-Type Engine Signature	
Creating a Custom Service HTTP Signature	
Creating a Custom Service FTP Signature	
Creating a Custom ATOMIC.ARP Engine Signature	
Creating a Custom ATOMIC.IP Engine Signature	
Creating a Custom TCP Sweep Signature	
Creating a Custom ICMP Sweep Signature	
Creating a Custom Trojan Engine Signature	
Enabling Shunning and Blocking (Enabling Blocking Properties)	
Enabling the TCP Reset Function	
Configure Cisco Ironport WSA	

https://t.me/learningnets

Technology	Mastered?
Configuring WCCP	
Active Dir Integration	
Custom Categories	
HTTPS Config	
Services Configuration (Web Reputation)	
Configuring Proxy By-pass Lists	
Web proxy modes	
Application visibility and control	

Configurations

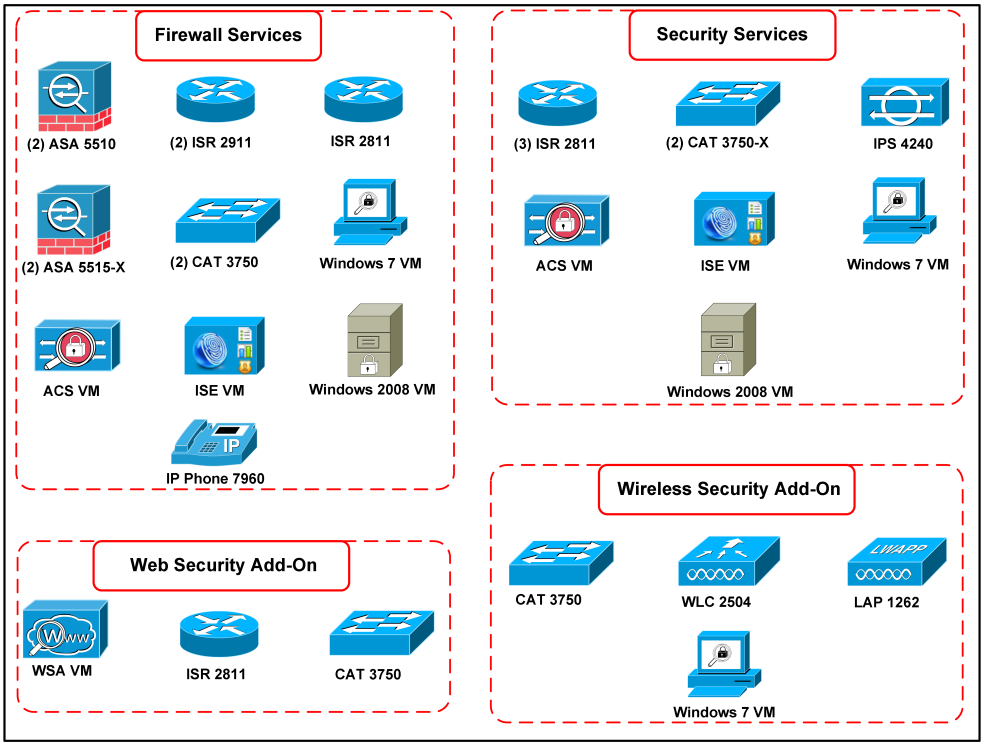
Initial configuration files are provided at the top of this page. Each lab has an initial configuration that is assumed. The initial configuration in some cases is the configuration that the previous lab ended with. If there is only one configuration file, the other devices did not change in the previous lab.

Load the Section 3 Initial Configuration Files to initialize your rack. Pay attention to the instructions in each task, because you may be required to load a different set of initial configuration files.

Section 3 General Lab Topology

The INE CCIE Security Racks are divided into modules as seen in the figure below. This section, Section 3, utilizes Firewall Services, Security Services, and Web Security.

https://t.me/learningnets



In this section, you will focus on how to provide various security services using IPS and WSA solutions. As you work on each task, focus on the small elements of building each technology. Gaining an in-depth knowledge of the workings of small tasks will lead to better troubleshooting of large scenarios in the CCIE Lab exam, as well in as production environments. Pay close attention to the topology, which may change.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124

https://t.me/learningnets

- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Setup Wizard

Task

- Initialize the WSA using the following values:
  - Hostname **wsa.inelab.local**.
  - NTP server 172.16.6.1 and DNS server 172.16.20.100.
  - IP address 172.16.6.100/24 with a default gateway of 172.16.6.1.
  - The **admin** user should have a password of **default1A**.
  - Disable Senderbase Network Participation.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.

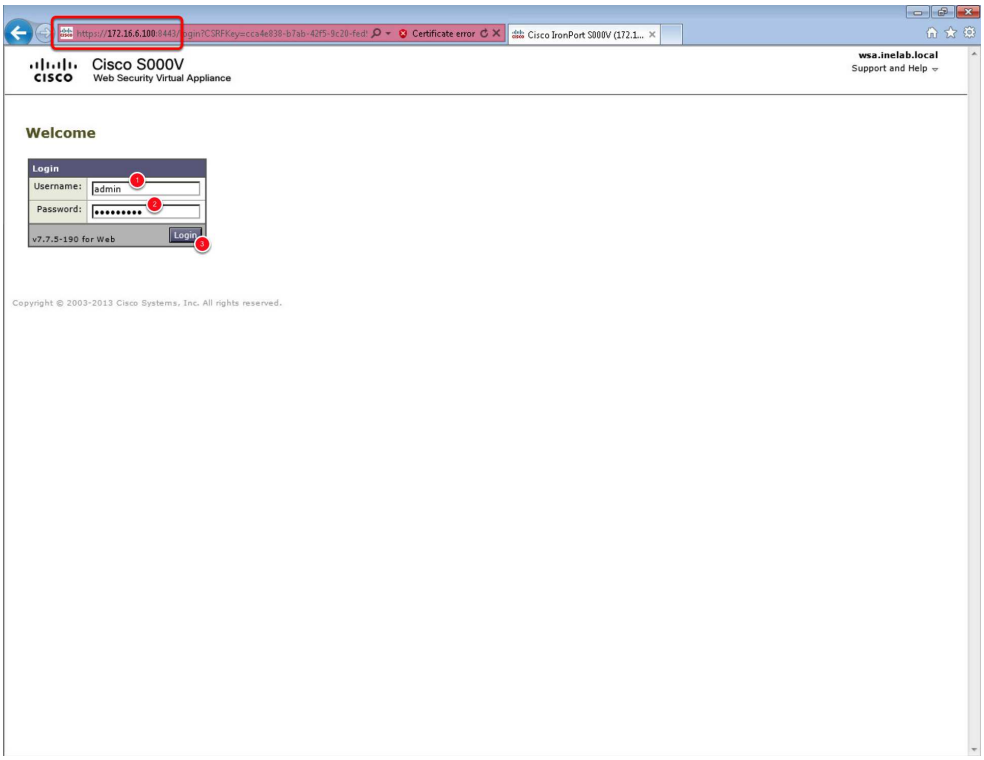
Overview

The WSA is pre-configured with the above requirements for all labs to allow you to quickly pick up any task and finish it successfully without going through the initial system setup wizard each time. However, as in the CCIE lab, you may find the WSA unconfigured, at the factory default state; it's good to know how to initialize it. This initial setup is a requirement and cannot be skipped; if configuration errors are made in this process, they can be changed afterward.

https://t.me/learningnets

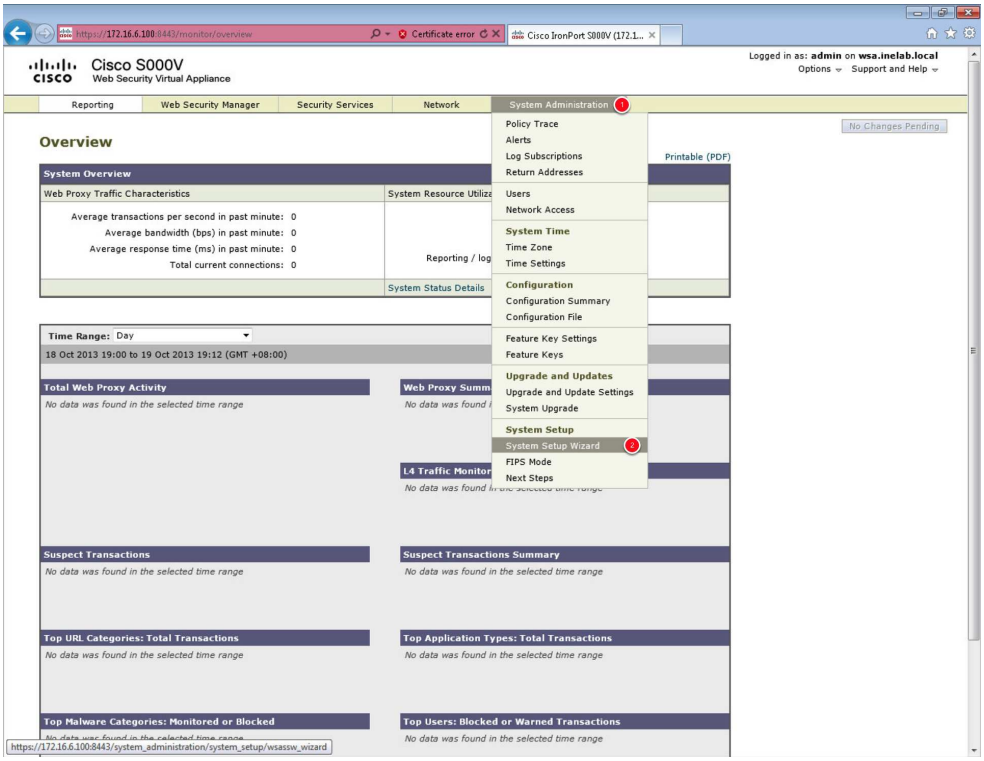
Configuration

Navigate to the WSA GUI using HTTP or HTTPS on the well-known ports of 8080 and 8443. The two URLs are **http://172.16.6.100:8080** and **https://172.16.6.100:8443**, with the HTTP request being redirected to HTTPS by default. Authenticate using **admin/default1A** as the **username/password** pair:

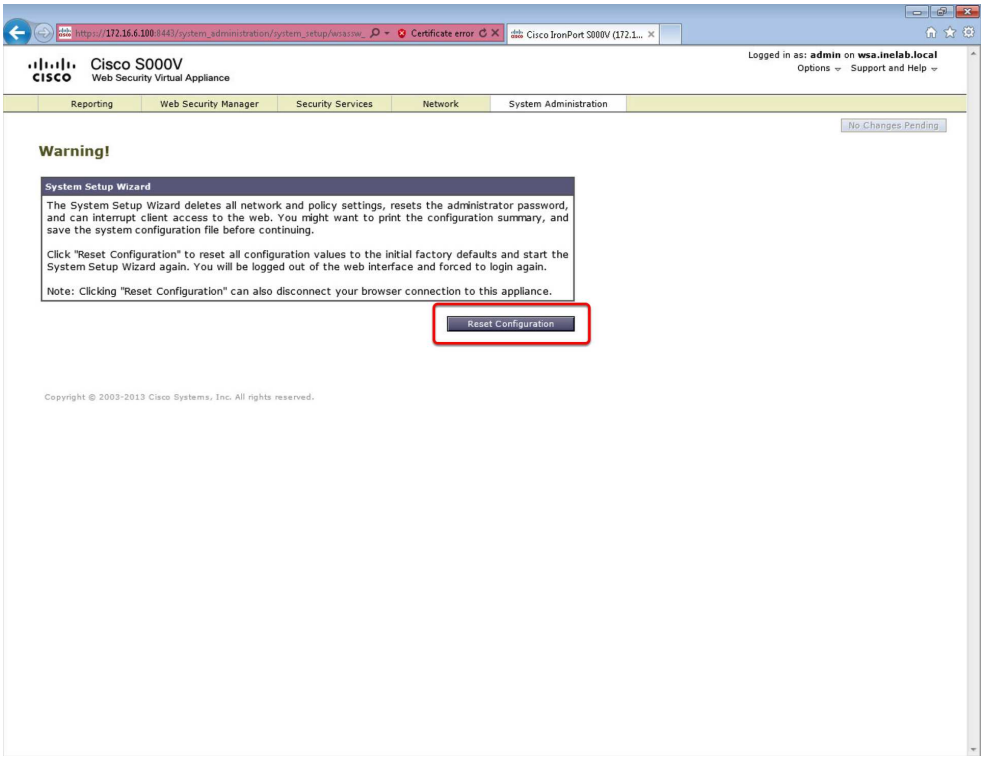


Browse through the GUI to start the initial setup wizard:

https://t.me/learningnets



Reset the WSA to factory defaults; this can also be achieved through the CLI by using the command `resetconfig` ; it is also useful if you want to quit any configuration changes and start over again:



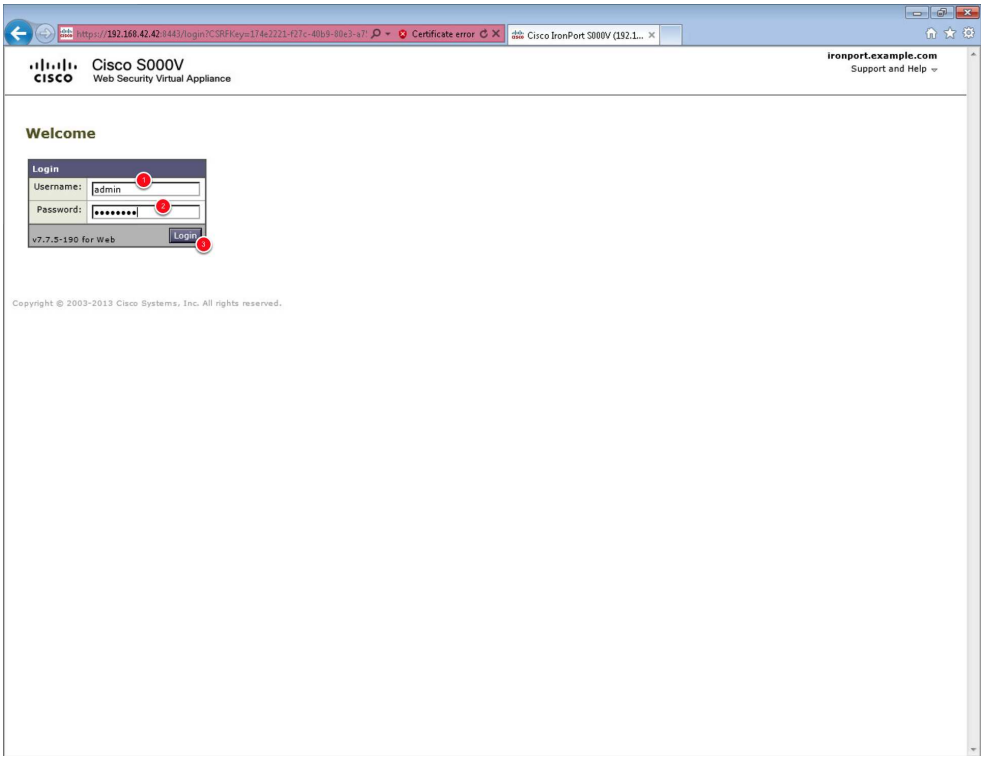
https://t.me/learningnets

When the WSA is at the factory default state, it has a preconfigured IP address of 192.168.42.42 on the M1 (management) interface. The quickest method to access it on this IP address is to assign **TEST-PC-B** to VLAN 28 and configure it with an IP address from 192.168.42.0/24 subnet:

```
SW3:

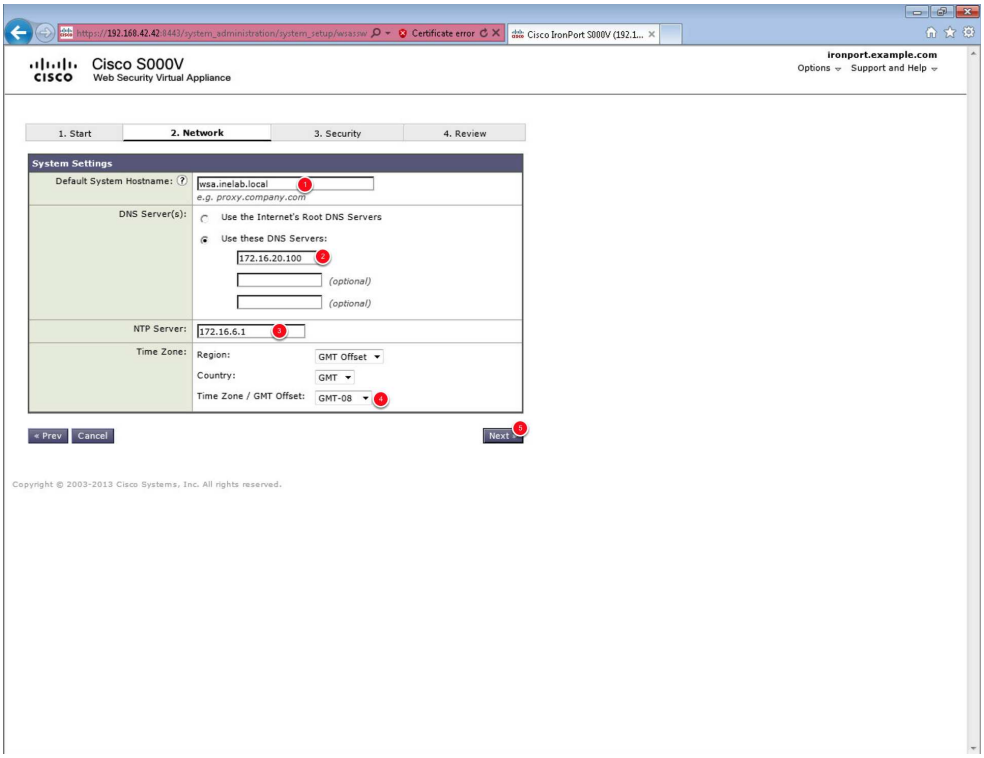

interface GigabitEthernet1/0/5
 switchport access vlan 28
```

Navigate to the WSA GUI using HTTP or HTTPS on the well-known ports of 8080 and 8443. The two URLs are **http://192.168.42.42:8080** and **https://192.168.42.42:8443**, with the HTTP request being redirected to HTTPS by default. Authenticate using **admin/ironport** as the **username/password** pair, the default values:

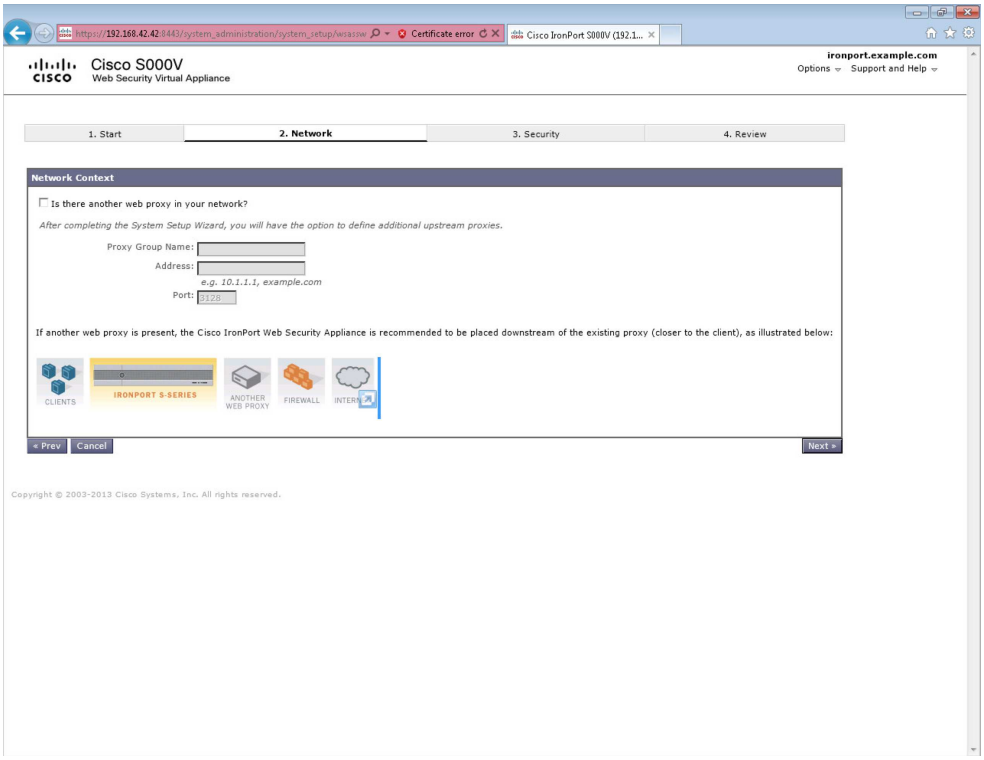


Configure the WSA hostname (as it's seen in the CLI and email notifications), DNS, and NTP server, and specify GMT offset:

https://t.me/learningnets

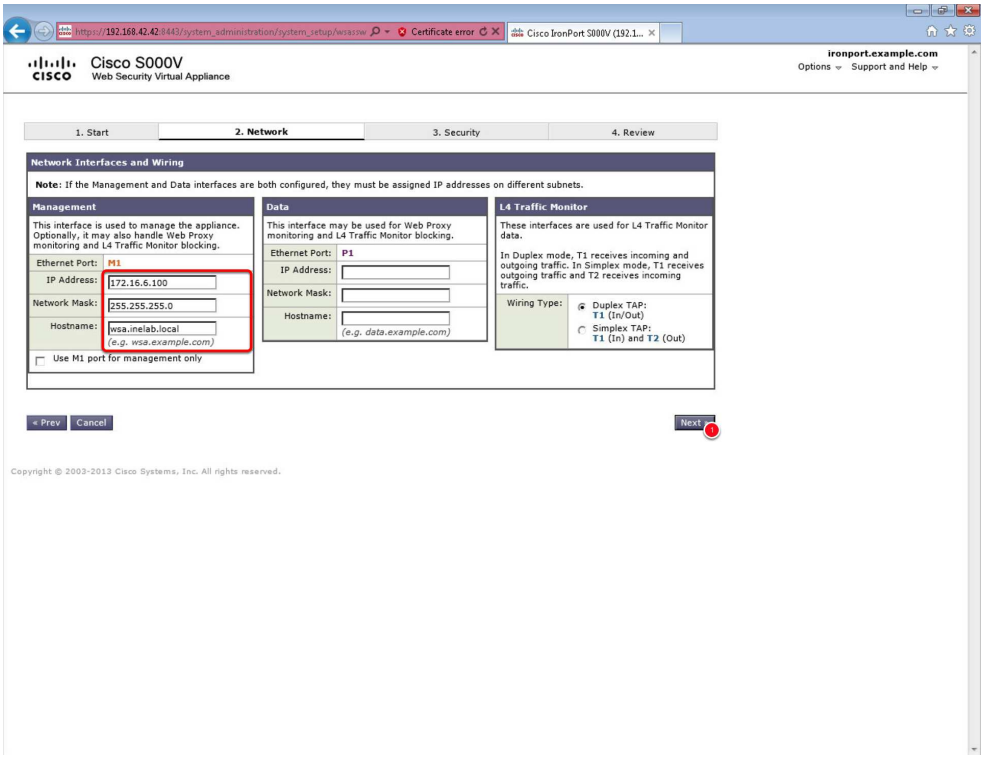


Configure nothing on the next page because there is no upstream proxy:

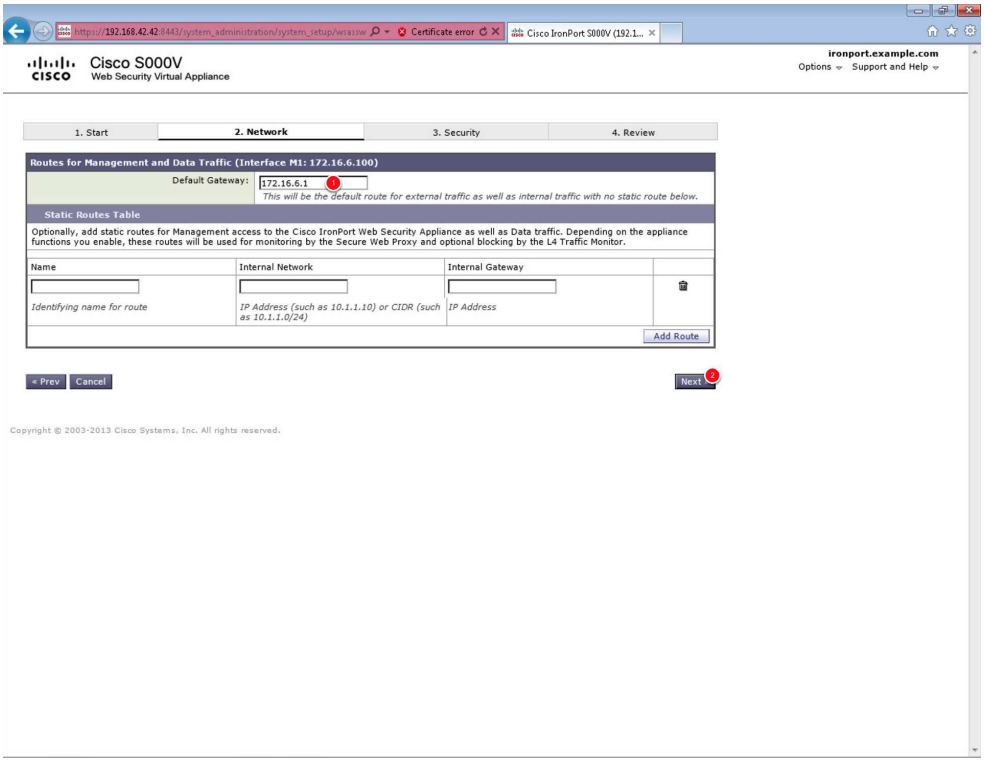


Configure the IP addressing and WSA hostname for the management interface (it will be used, for example, when WSA is joined to the Active Directory):

https://t.me/learningnets

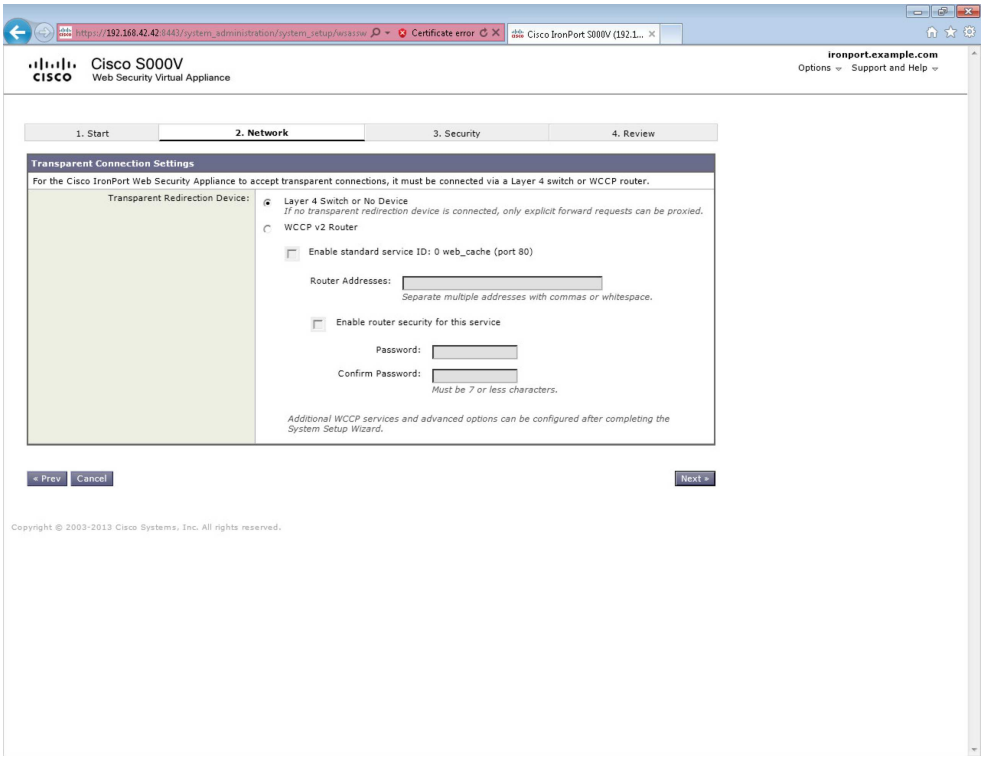


Configure the default gateway; there is no need for any specific routes:

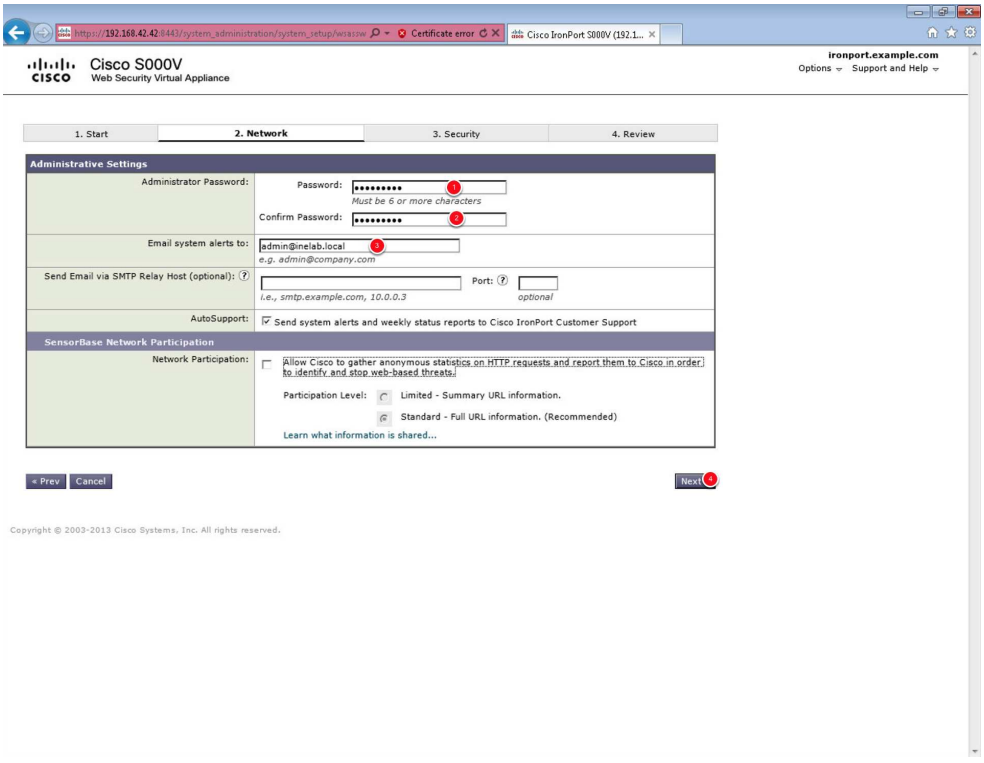


Leave the default values for transparent mode based on task requirements; these can be changed later:

https://t.me/learningnets

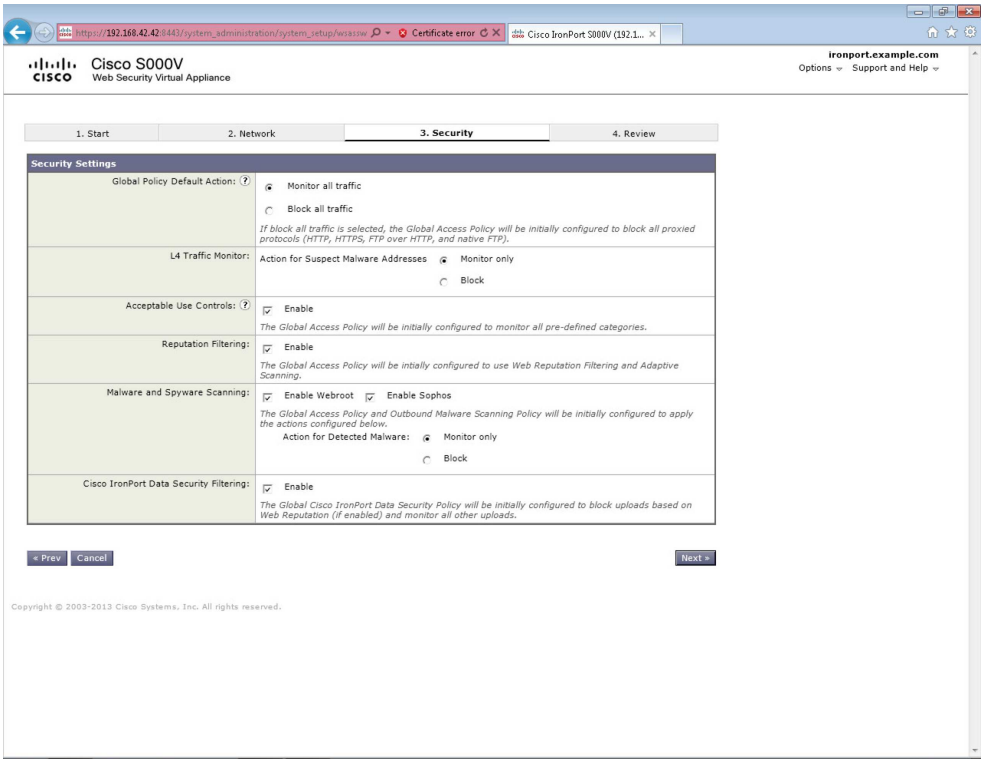


Change the default password for the **admin** user to **default1A**, specify the system alert email recipient, and disable Senderbase network participation:

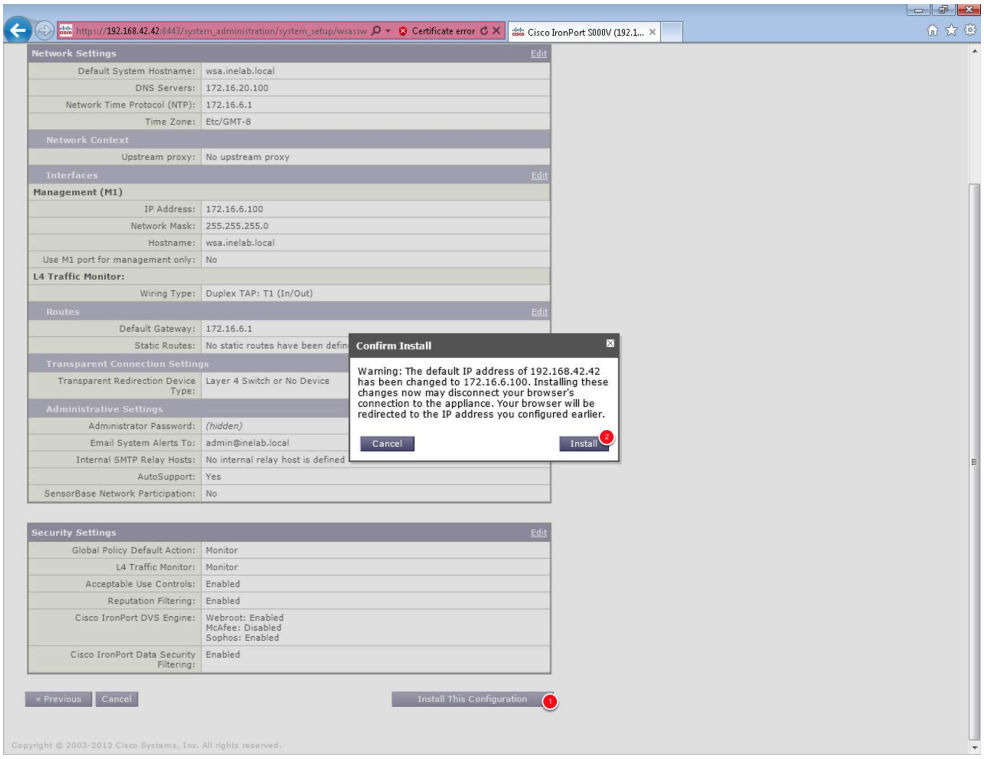


Leave the default values in the next screen based on task requirements; these can be changed later:

https://t.me/learningnets



Install the configuration:



Configure **TEST-PC-B** back to VLAN 37 and set it for DHCP-based IP addressing:

https://t.me/learningnets

SW3:

```
interface GigabitEthernet1/0/5
 switchport access vlan 37
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Management Access

Task

- Configure the WSA so that HTTPS redirection for management purposes is not active, and allow it to be managed from only 136.1.37.0/24 subnet.
- Using RADIUS integration with the ACS2 server:
  - Authorize the user **SUPERUSER** for the administrator role.
  - Authorize the user **GUESTUSER** for the guest role.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the WSA to the initial configuration (not to factory defaults), load it again by navigating to **System Administration > Configuration File > Load Configuration**, select the **WSA-DEFAULT.xml** file, click **Load**, and commit changes afterward.

Overview

WSA supports an RBAC (Role-Based Access Control) setup for management purposes and uses four built-in roles:

- **Administrator**, with full administrative rights.
- **Operator**, a restrictive administrator account with no rights for creating management

https://t.me/learningnets

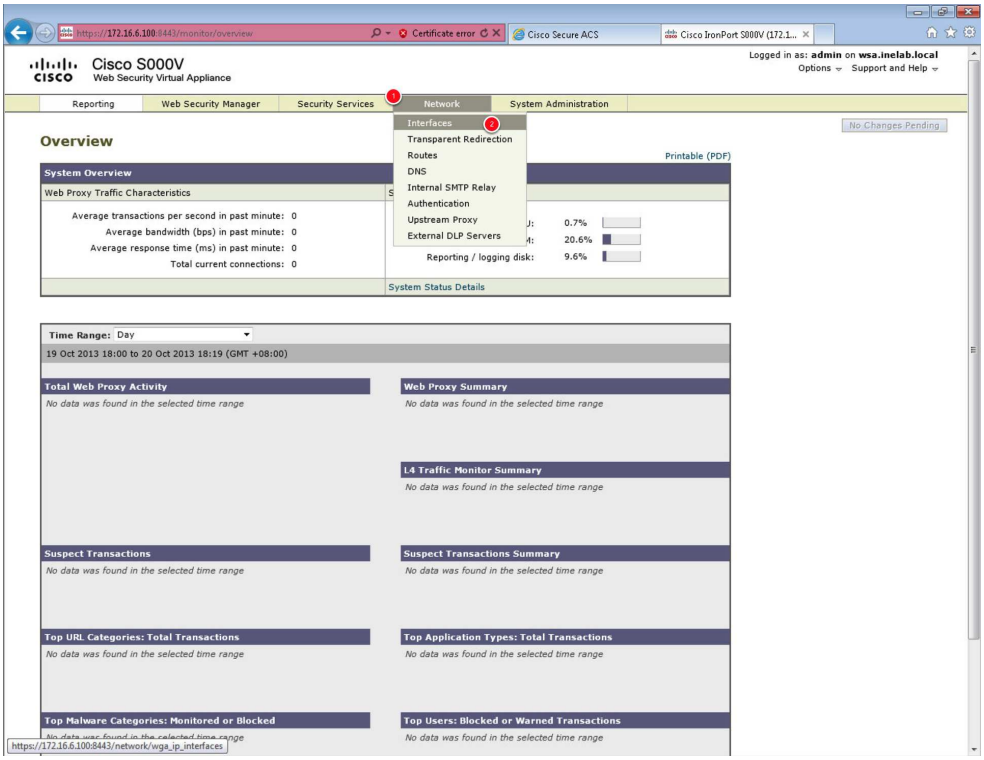
user accounts, reseting the configuration to factory defaults, or performing upgrades.

- **Read-Only Operator**, which can read configuration but make no changes.
- **Guest**, which can only view system status information.

RBAC is supported for both local user authentication and remote user authentication through RADIUS by the use of IETF Radius attribute 25 (Class Attribute). The RADIUS server sends a value for attribute 25 as part of the authorization process, which is mapped to a user role on WSA. When RADIUS integration is configured, WSA (unlike IOS routers) does not use a fallback mechanism to the local user database if the RADIUS server is unreachable; instead, authentication is simultaneously allowed both for local and remote RADIUS accounts. Basically when a user logs on to the WSA, if the user account is not found in the local database, the RADIUS server is queried, and therefore it's mandatory that usernames from local and RADIUS databases be different.

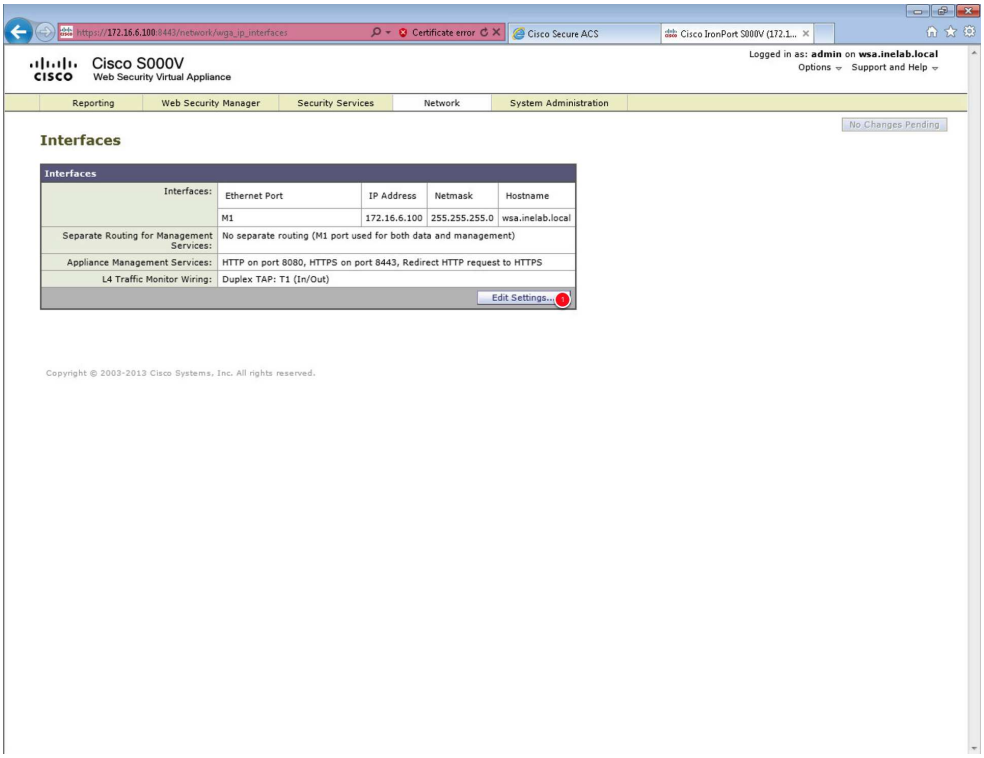
Configuration

Authenticate on the WSA GUI and browse to the network configuration page:

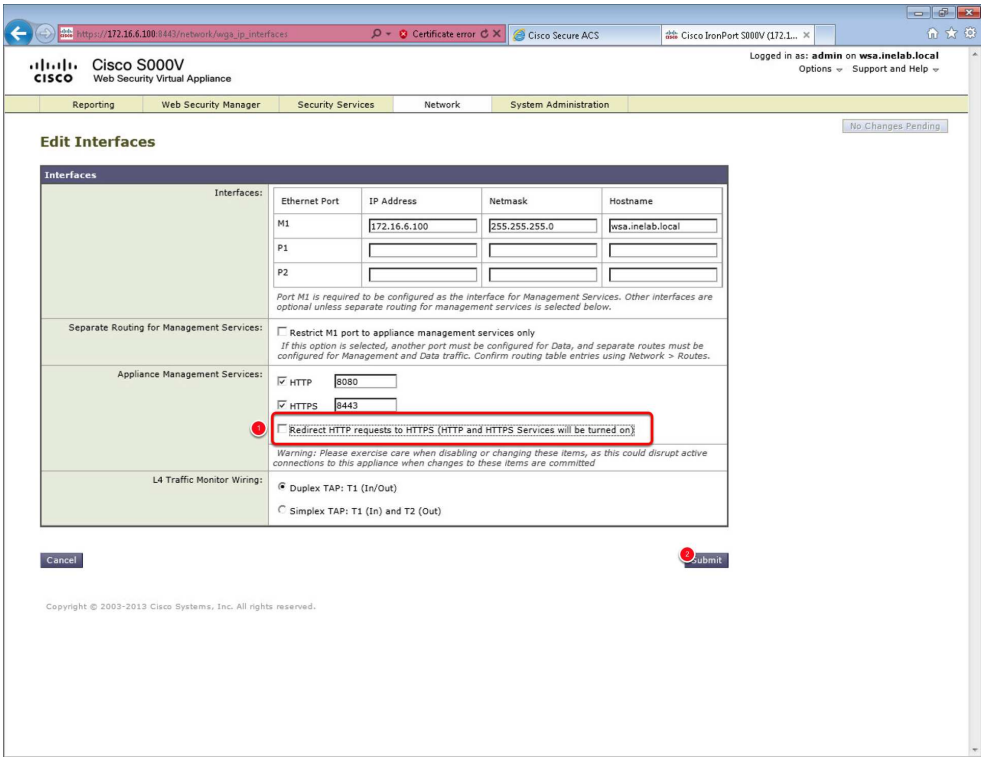


To change the network settings, click **Edit Settings**:

https://t.me/learningnets

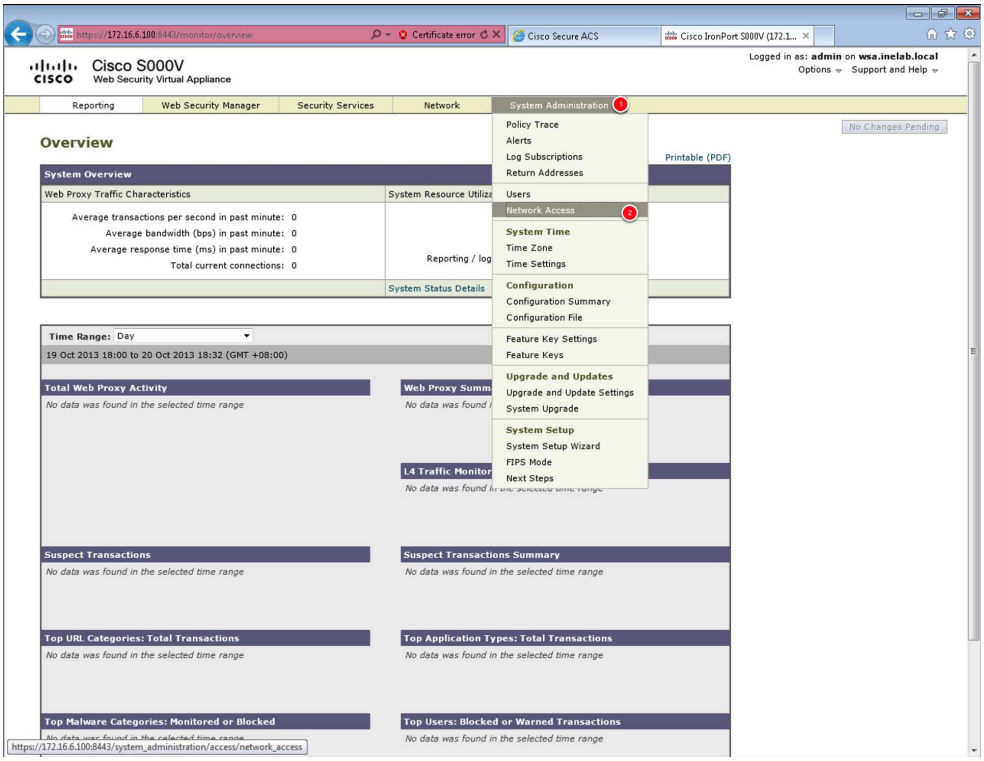


Disable the HTTP to HTTPS redirection:

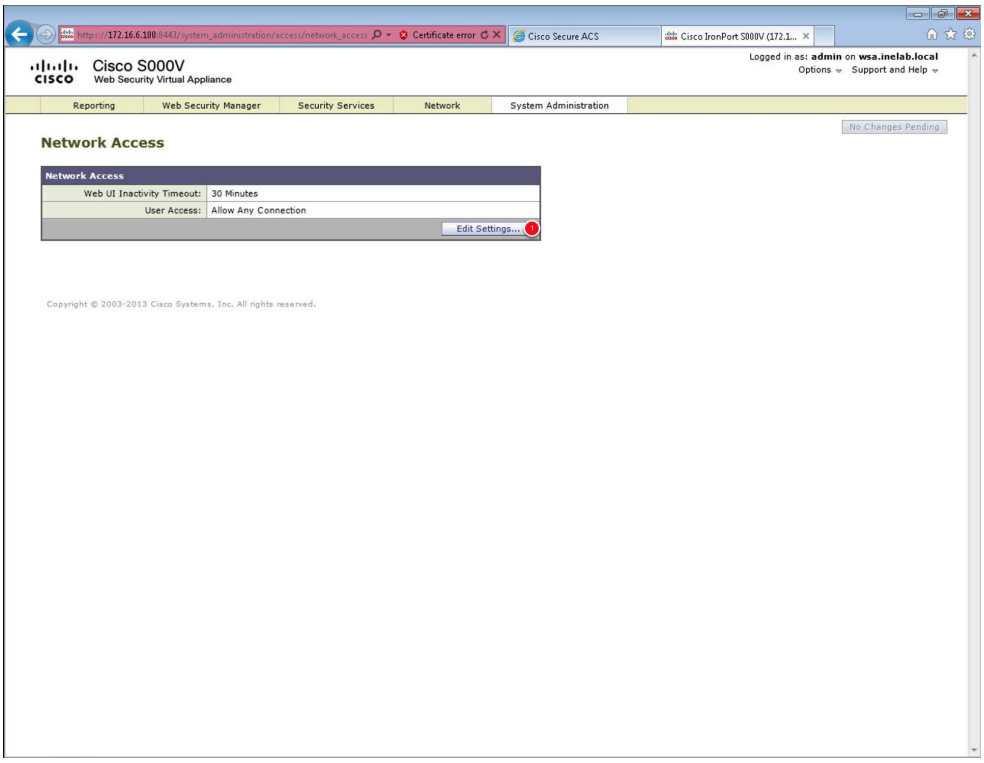


To restrict subnets from where management connections are allowed, browse to the network access page:

https://t.me/learningnets

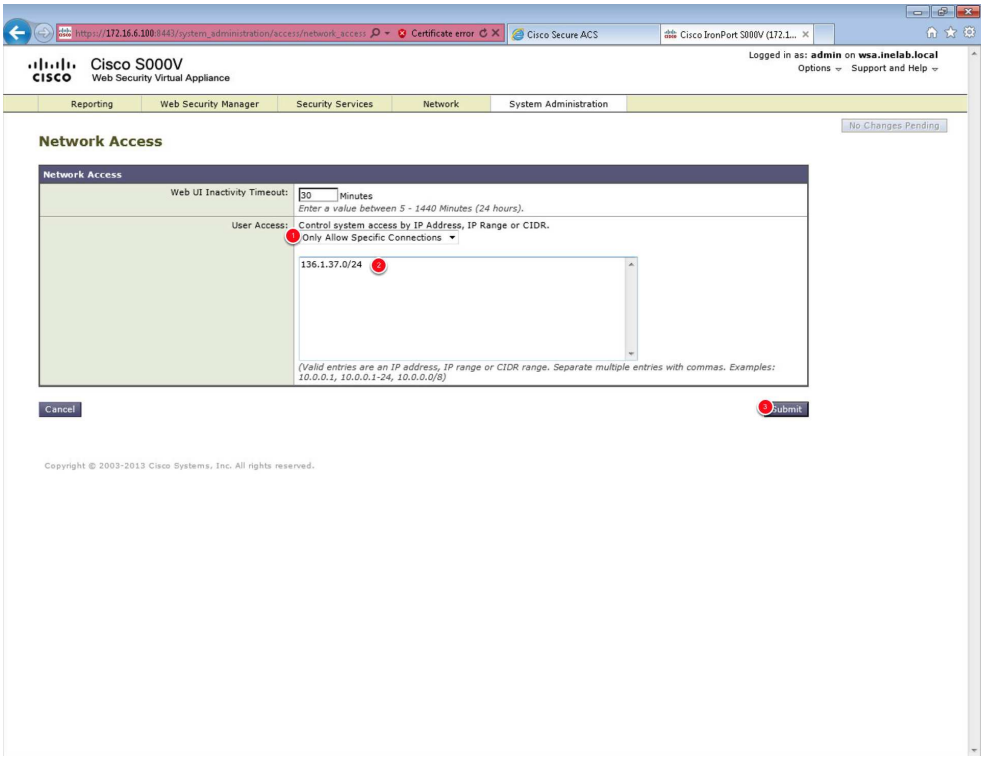


To make configuration changes, click **Edit Settings**:

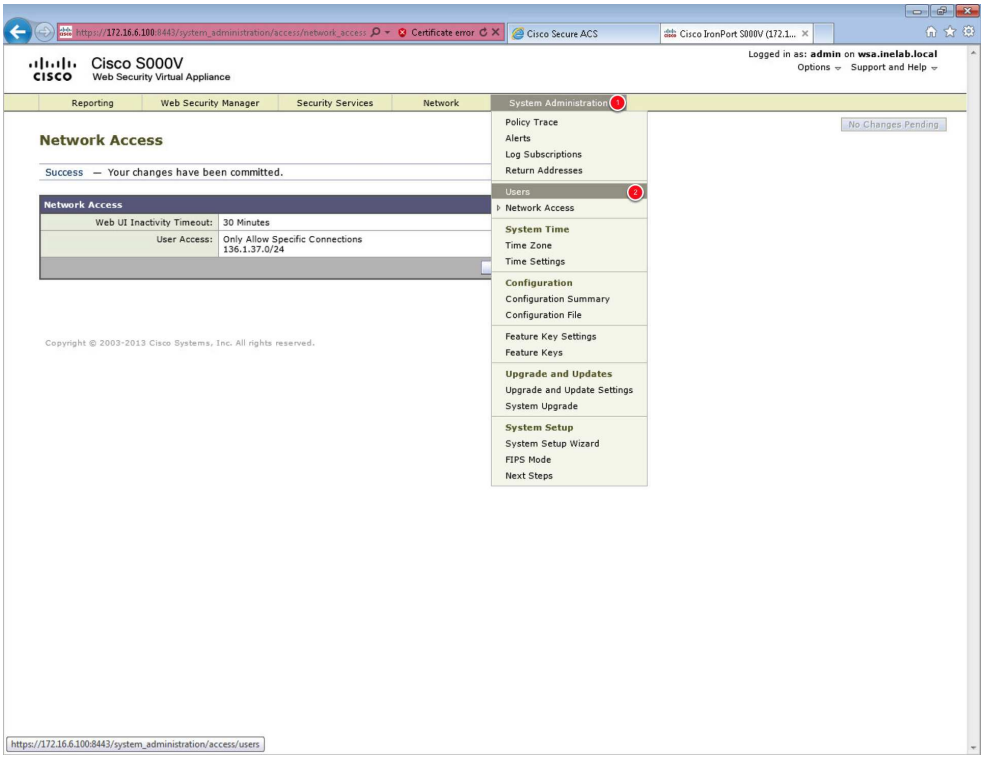


Enable the functionality, and then specify the subnet:

https://t.me/learningnets

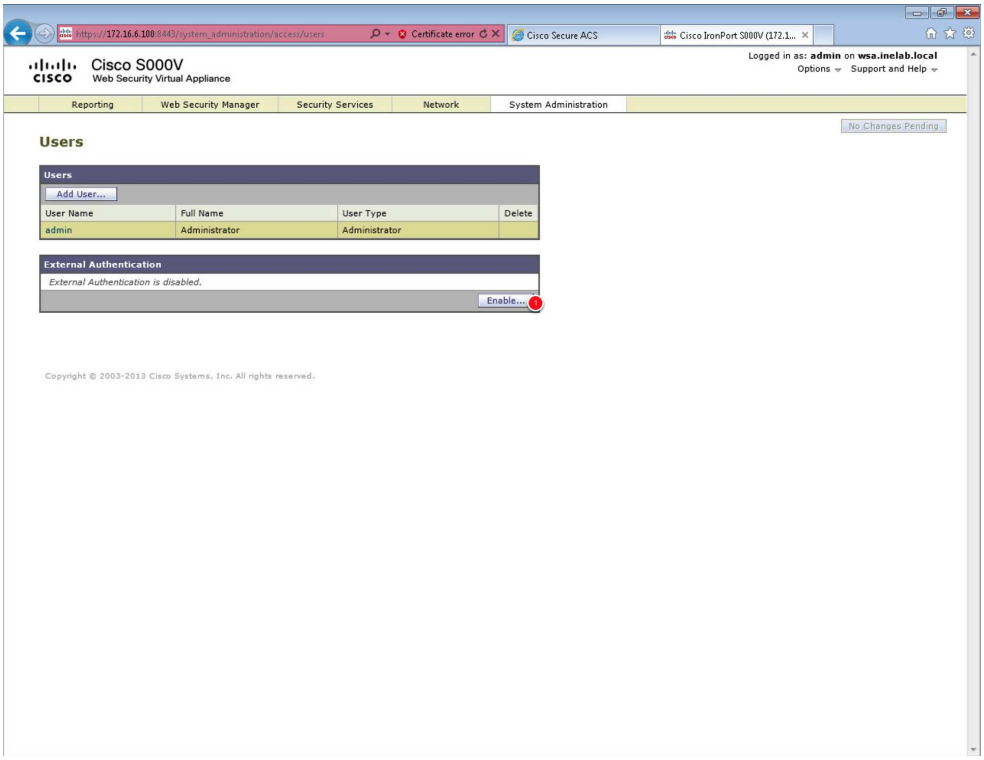


For RADIUS configuration settings, browse to users page:

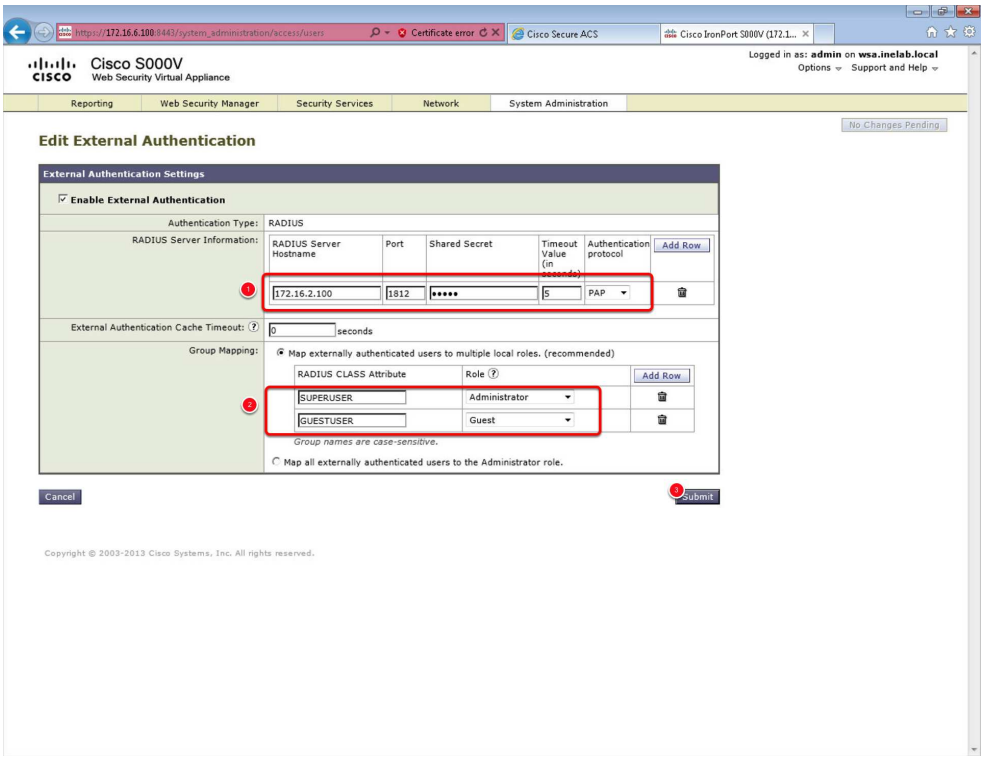


Enable the RADIUS integration:

https://t.me/learningnets

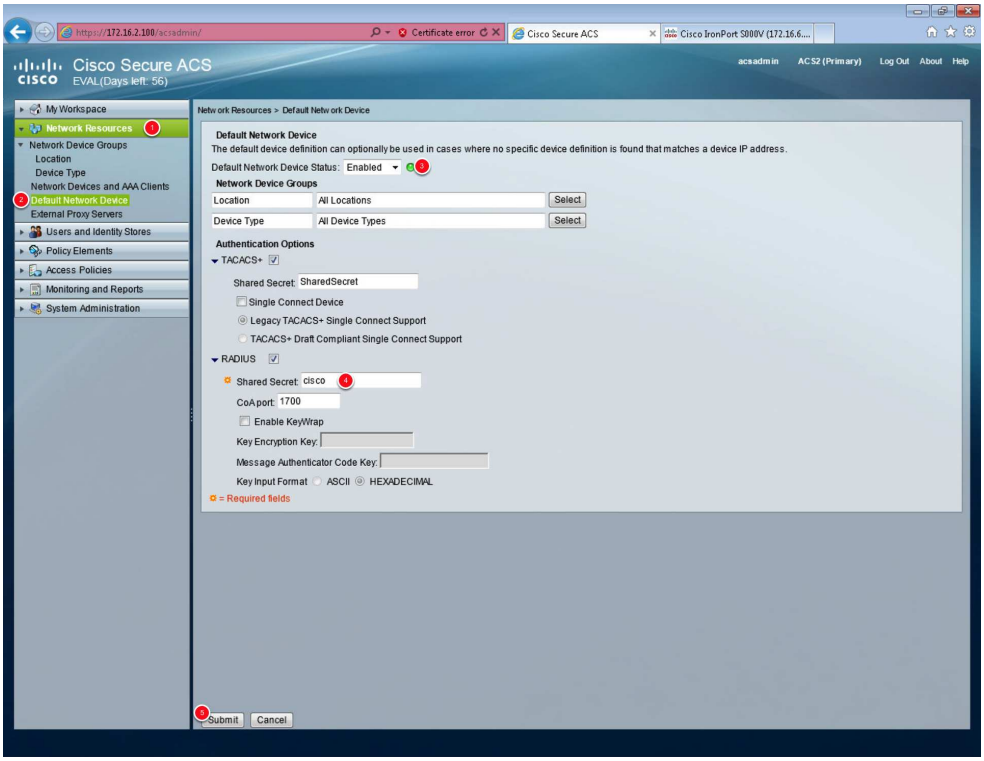


Configure the RADIUS server setting and roles through the use of the RADIUS Class Attribute:

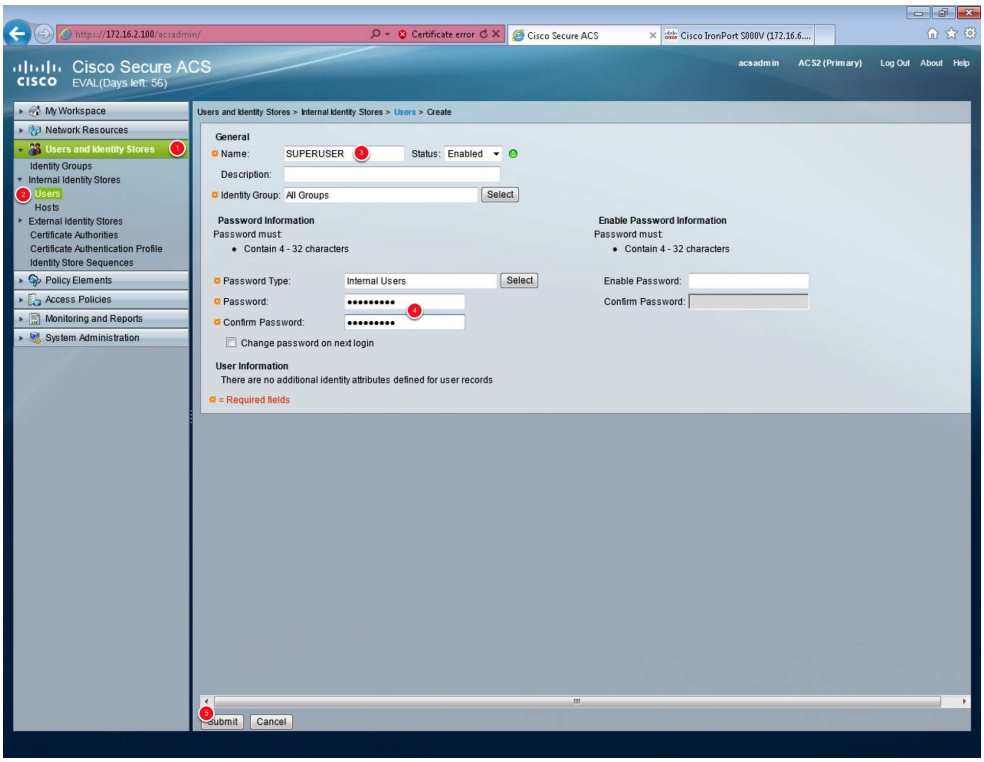


Add WSA as a NAD into ACS, or configure the default network device:

https://t.me/learningnets

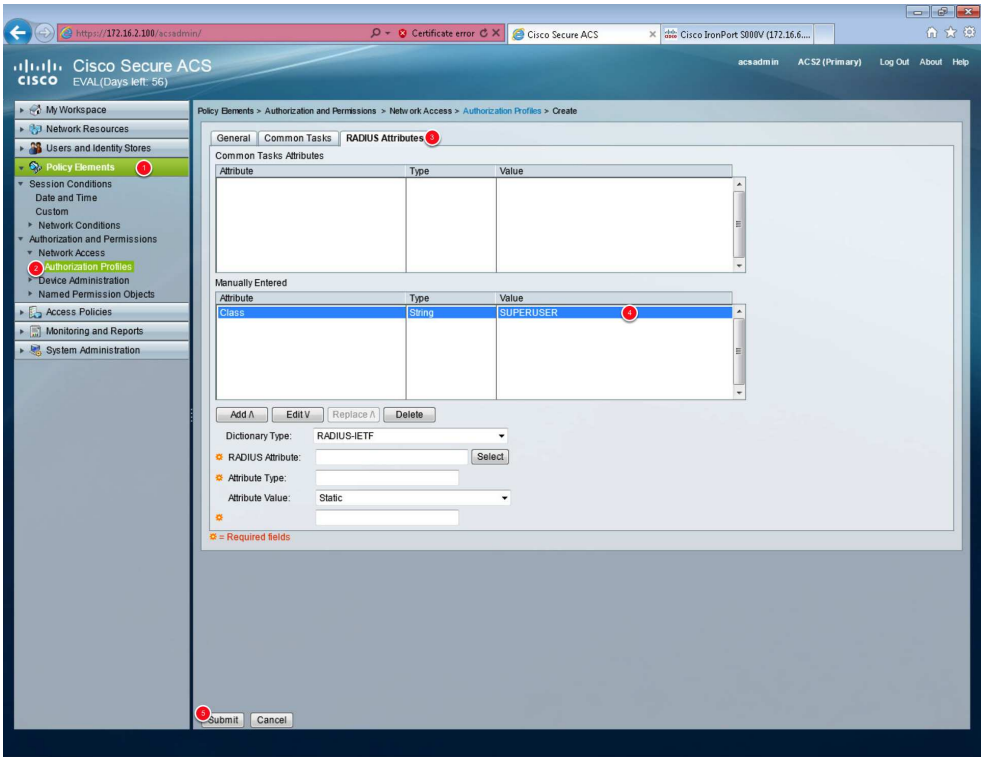


Configure the two users into ACS; for example, **SUPERUSER**:

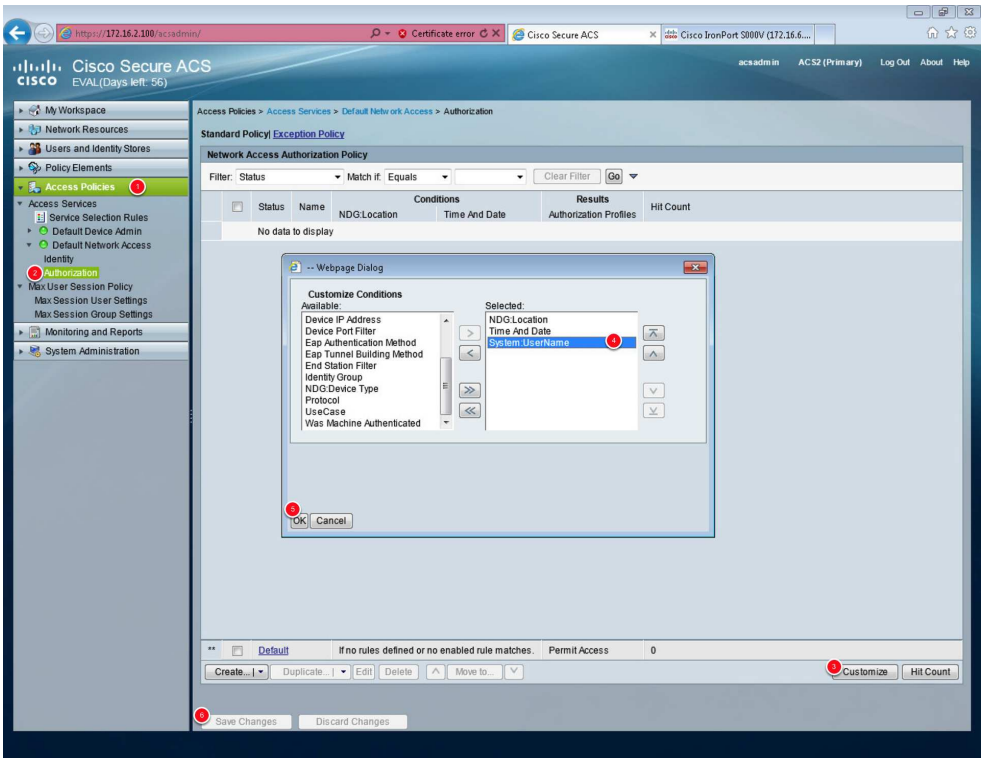


Configure two authorization profiles, one for **SUPERUSER** and one for **GUESTUSER**; assign RADIUS Attribute 25 with values configured on the WSA:

https://t.me/learningnets

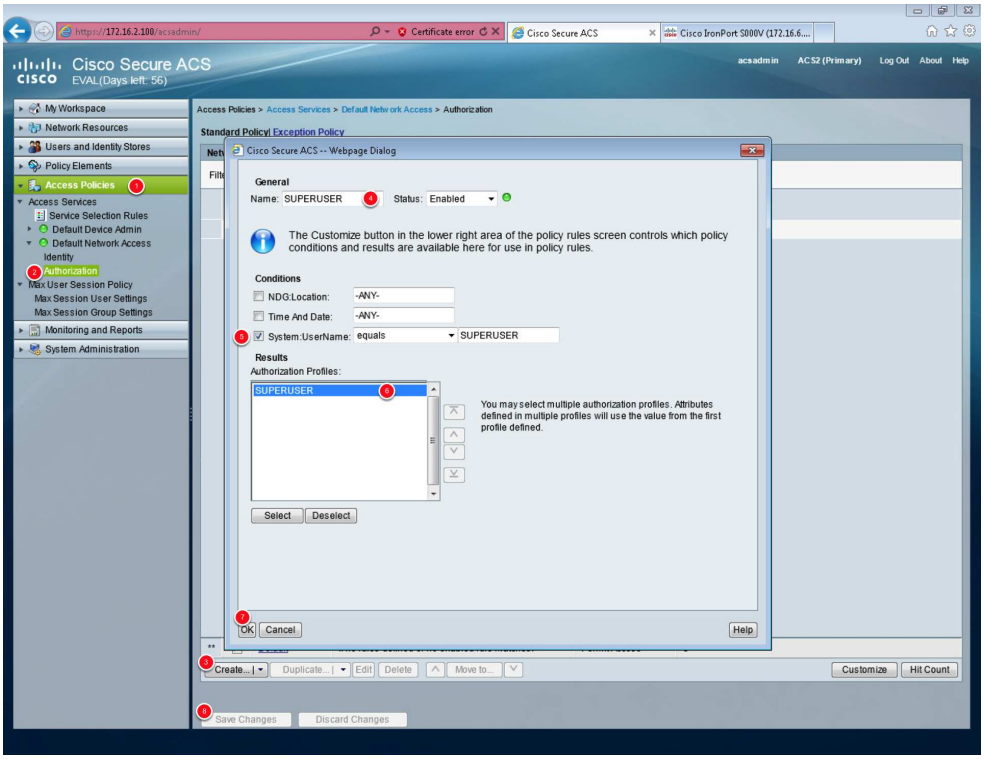


Add the username as a variable for RADIUS authorization policies:



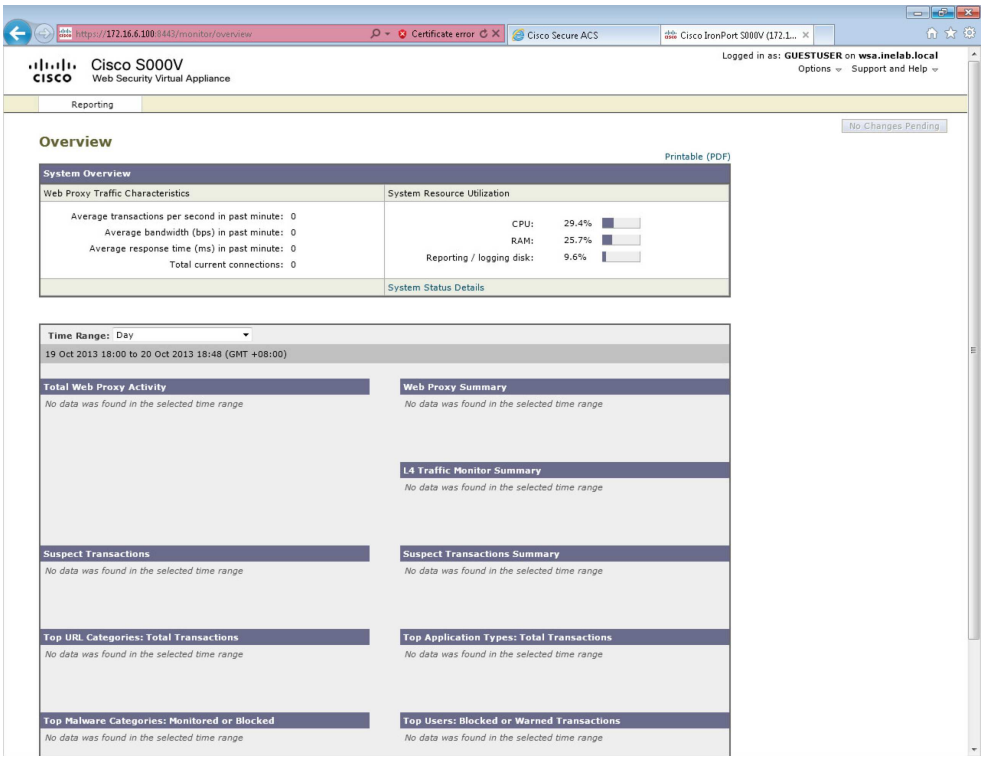
Configure two authorization policies, one for **SUPERUSER** and one for **GUESTUSER**:

https://t.me/learningnets



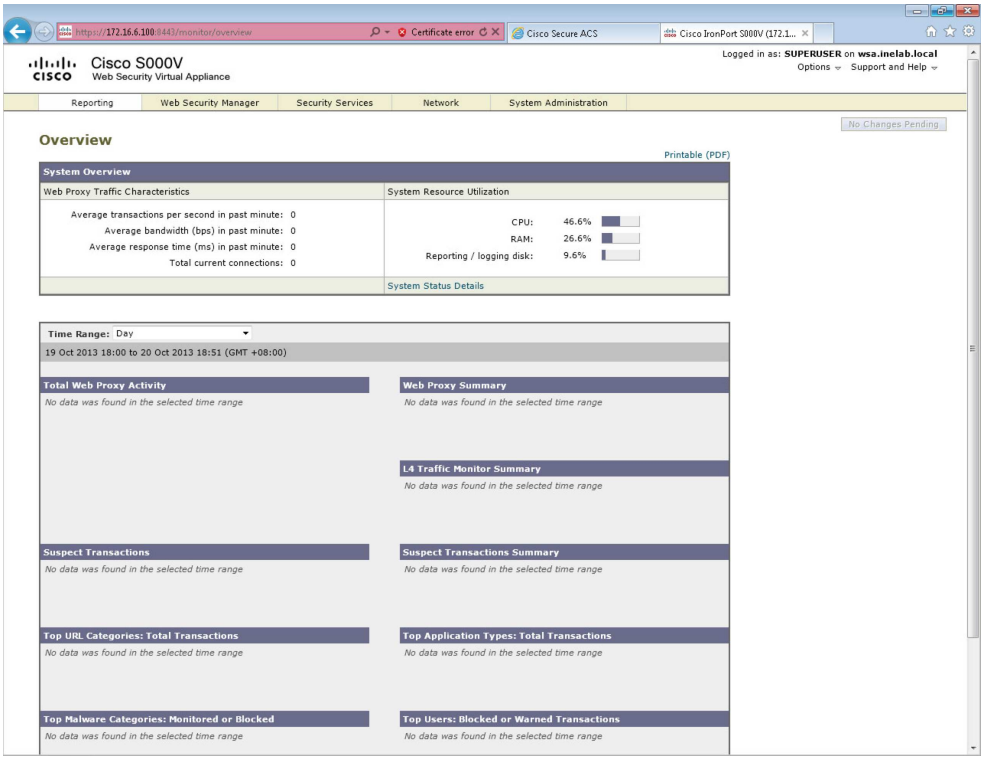
Verification

Authenticate with **GUESTUSER** and verify that it has restricted access to the GUI:



Authenticate with **SUPERUSER** and verify that it has unrestricted access to the GUI:

https://t.me/learningnets



Verify the WSA external authentication logs:

```
wsa.inelab.local> tail

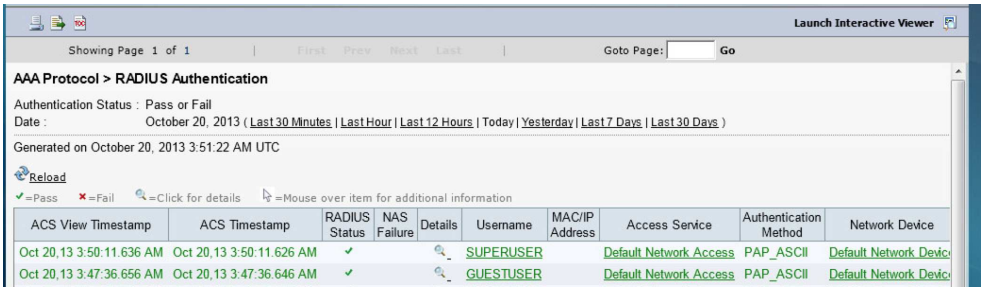
  Currently configured logs:
  1. "accesslogs" Type: "Access Logs" Retrieval: FTP Poll
  2. "authlogs" Type: "Authentication Framework Logs" Retrieval: FTP Poll
  3. "avc_logs" Type: "AVC Engine Logs" Retrieval: FTP Poll
  4. "bypasslogs" Type: "Proxy Bypass Logs" Retrieval: FTP Poll
  5. "cli_logs" Type: "CLI Audit Logs" Retrieval: FTP Poll
  6. "configdefragd_logs" Type: "Configuration Logs" Retrieval: FTP Poll
  7. "dca_logs" Type: "DCA Engine Logs" Retrieval: FTP Poll
  8. "external_auth_logs" Type: "External Authentication Logs" Retrieval: FTP  Poll
  9. "feedback_logs" Type: "Feedback Logs" Retrieval: FTP Poll
  10. "fips_logs" Type: "FIPS Logs" Retrieval: FTP Poll
  11. "ftpd_logs" Type: "FTP Server Logs" Retrieval: FTP Poll
  12. "gui_logs" Type: "GUI Logs" Retrieval: FTP Poll
  13. "haystackd_logs" Type: "Haystack Logs" Retrieval: FTP Poll
  14. "idsdataloss_logs" Type: "Data Security Logs" Retrieval: FTP Poll
  15. "logderrorlogs" Type: "Logging Logs" Retrieval: FTP Poll
  16. "mcafee_logs" Type: "McAfee Logs" Retrieval: FTP Poll
  17. "musd_logs" Type: "AnyConnect Secure Mobility Daemon Logs" Retrieval: FTP
  Poll
  18. "ocspd_logs" Type: "OCSP Logs" Retrieval: FTP Poll
  19. "pacd_logs" Type: "PAC File Hosting Daemon Logs" Retrieval: FTP Poll
  20. "proxylogs" Type: "Default Proxy Logs" Retrieval: FTP Poll
  21. "reportd_logs" Type: "Reporting Logs" Retrieval: FTP Poll
```

https://t.me/learningnets

```
22. "reportqueryd_logs" Type: "Reporting Query Logs" Retrieval: FTP Poll
23. "saas_auth_log" Type: "SaaS Auth Logs" Retrieval: FTP Poll
24. "shd_logs" Type: "SHD Logs" Retrieval: FTP Poll
25. "snmp_logs" Type: "SNMP Logs" Retrieval: FTP Poll
26. "sntpd_logs" Type: "NTP Logs" Retrieval: FTP Poll
27. "sophos_logs" Type: "Sophos Logs" Retrieval: FTP Poll
28. "status" Type: "Status Logs" Retrieval: FTP Poll
29. "system_logs" Type: "System Logs" Retrieval: FTP Poll
30. "trafmon_errlogs" Type: "Traffic Monitor Error Logs" Retrieval: FTP Poll
31. "trafmonlogs" Type: "Traffic Monitor Logs" Retrieval: FTP Poll
32. "uds_logs" Type: "UDS Logs" Retrieval: FTP Poll
33. "updater_logs" Type: "Updater Logs" Retrieval: FTP Poll
34. "wbnp_logs" Type: "WBNP Logs" Retrieval: FTP Poll
35. "webcat_logs" Type: "Web Categorization Logs" Retrieval: FTP Poll
36. "webrootlogs" Type: "Webroot Logs" Retrieval: FTP Poll
37. "welcomeack_logs" Type: "Welcome Page Acknowledgement Logs" Retrieval:
FTP Poll
Enter the number of the log you wish to tail. []> 8

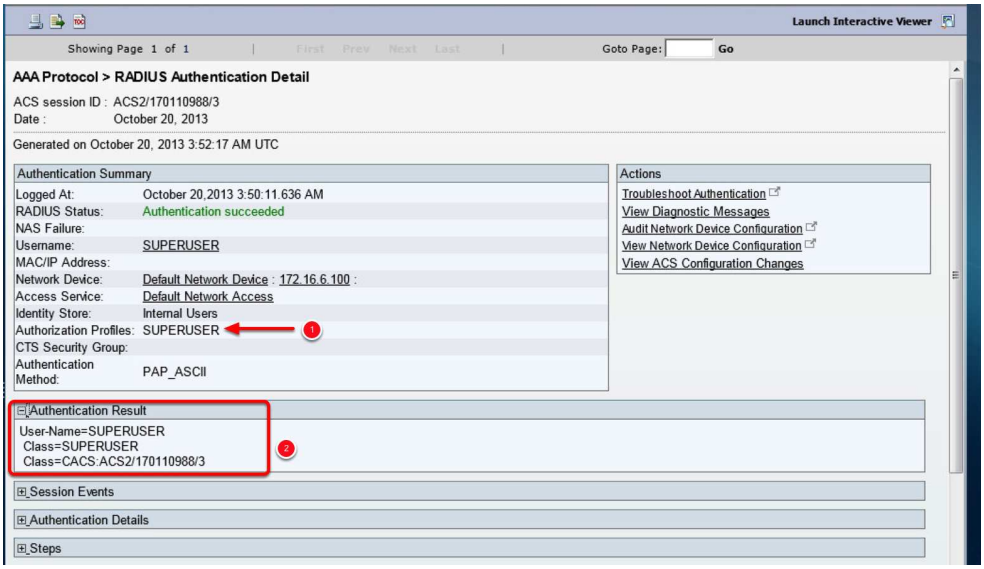
Press Ctrl-C to stop.
Sun Oct 20 17:58:15 2013 Info: User admin from 136.1.37.51 was authenticated successfully.
Sun Oct 20 18:27:18 2013 Info: User admin from cuau0 was authenticated successfully.
Sun Oct 20 18:48:06 2013 Info: logout:136.1.37.51 user:admin session:BWdnPLSfQTWkyrDInGhT
Sun Oct 20 18:48:25 2013 Info: Trying RADIUS server 172.16.2.100
Sun Oct 20 18:48:25 2013 Info: User 'GUESTUSER' is given 'guest' role
Sun Oct 20 18:48:25 2013 Info: User GUESTUSER from 136.1.37.51 was authenticated successfully.
Sun Oct 20 18:52:03 2013 Info: Trying RADIUS server 172.16.2.100
Sun Oct 20 18:52:03 2013 Info: User 'SUPERUSER' is given 'admin' role
Sun Oct 20 18:52:03 2013 Info: User SUPERUSER from 136.1.37.51 was authenticated successfully.
```

Verify the RADIUS authentication logs from ACS:



Expand the authentication log for **SUPERUSER** and check the matched RADIUS authorization policy:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Explicit Forward Mode

Task

- Configure Mozilla Firefox on **TEST-PC-B** to use WSA as its HTTP proxy server on port 3333.
- Define one identity and one access-policy for **TEST-PC-B**, but accommodate for the fact that it's being DHCP assigned and:
  - Allow access to **Gambling** and **Games** URL categories only on the weekends.
  - Access to **www.google.com** and **www.ine.com** should be allowed, bypassing any inspections.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the WSA to the initial configuration (not to factory defaults), load it again by navigating to **System Administration > Configuration File > Load Configuration**, select the **WSA-DEFAULT.xml** file, click **Load**, and commit changes afterward.

Overview

A proxy server, namely the WSA, can be deployed in two ways, and it can serve

https://t.me/learningnets

client requests for both at the same time (this was not the case in previous AsyncOS versions):

- Explicit forward mode:
  - The client has the proxy configured in its browser settings, so the client directs traffic to the proxy server.
  - There is no need to configure the network infrastructure to redirect traffic to the WSA.
  - DNS resolution is performed by the proxy server, not by the client.

- Transparent mode
  - The client is not aware of the proxy; it's not configured in the browser.
  - The network infrastructure redirects the client's request to the WSA through the use of WCCP.
  - DNS resolution is performed by the client.

By default, WSA's HTTP proxy listens on ports 80 and 3128, because 3128 is used by most other proxies such as Squid and Microsoft TMG. Useful CLI commands for proxy configuration are **proxyconfig** and **advancedproxyconfig**. Additionally, the **webcache** command can be used to configure settings for caching.

Identities represent a classifier for any request received by the WSA (regardless of the deployment method, forward or transparent mode). There can be several identities configured, and these will be inspected top down with a stop on first match, just like entries in an access-list are processed. User authentication is not a transaction classifier but an optional requirement for the transaction to be allowed and is configured within the identity settings. There is one global identity that acts like a gateway of last resort if none of the user-configured identities are matched. WSA can classify transactions based on one or more of the following:

- Client subnet
- Protocol (HTTP/HTTPS/FTP)—mandatory
- Proxy ports (generally used in forward mode deployments)
- URL categories
- User-Agent (the client application making the request, such as Internet Explorer, Mozilla Firefox, or Yahoo Messenger)

Access policies define how HTTP traffic is inspected (including HTTPS decrypted traffic); whether the user's request is allowed is based on the configured settings and inspection results. There can be several access policies configured, and these will be inspected top down with a stop on first match, just like entries in an access-list are processed. There is one default global access policy that acts as a gateway

https://t.me/learningnets

of last resort if none of the user-configured access policies are matched; the global policy cannot be disabled, but it can be configured to deny all requests if it's being matched. For each access-policy, the following inspections can be performed and configured:

- Allowed protocols and user-agents
- URL Filtering (both predefined and custom categories)
- AVC (Application Visibility and Control), defining which services are allowed for certain applications (for example, allow Yahoo Messenger chat service but drop Yahoo Messenger voice service)
- Content download (for example, allow Word files but drop PDF files)
- Web Reputation Filtering, which automatically allows or denies requests based on the scoring (-10 to + 10) for the URL of the requested resource
- Antimalware, which bidirectionally inspects the session for malware and allows or denies it

In a regular deployment, you create identities to classify a request and map these identities to access policies. Technically, configuring an identity is not strictly required unless client authentication is mandatory, because the same classifiers from the identity configuration are available as well in the access policy settings. So an incoming HTTP request will be processed as follows:

- Classify the request by parsing the configured identities and stop on first match; if no identity is configured or matched, it will match on the default global identity.
- Verify whether the matched identity requires authentication; authentication must be successful to pass further.
- Inspect the configured access policies and stop on first match (based on the identities used by each access policy); if no access policies are configured or matched, it will match on the default global access policy.

WSA supports URL Filtering based on built-in categories, as well as custom (user-configured) URL categories. Available actions for matching traffic are:

- **Monitor**, which will allow the user access to the request web page. However, this is subject to further inspections such as antimalware or reputation score, so the final verdict may be to restrict access.
- **Allow**, which will allow the user access to the requested web page without any further inspections; it's basically a whitelisting.
- **Block**, which will restrict the user from accessing the requested web page and present a notification message.
- **Redirect**, which offers the possibility of redirecting the user to a certain web page.

https://t.me/learningnets

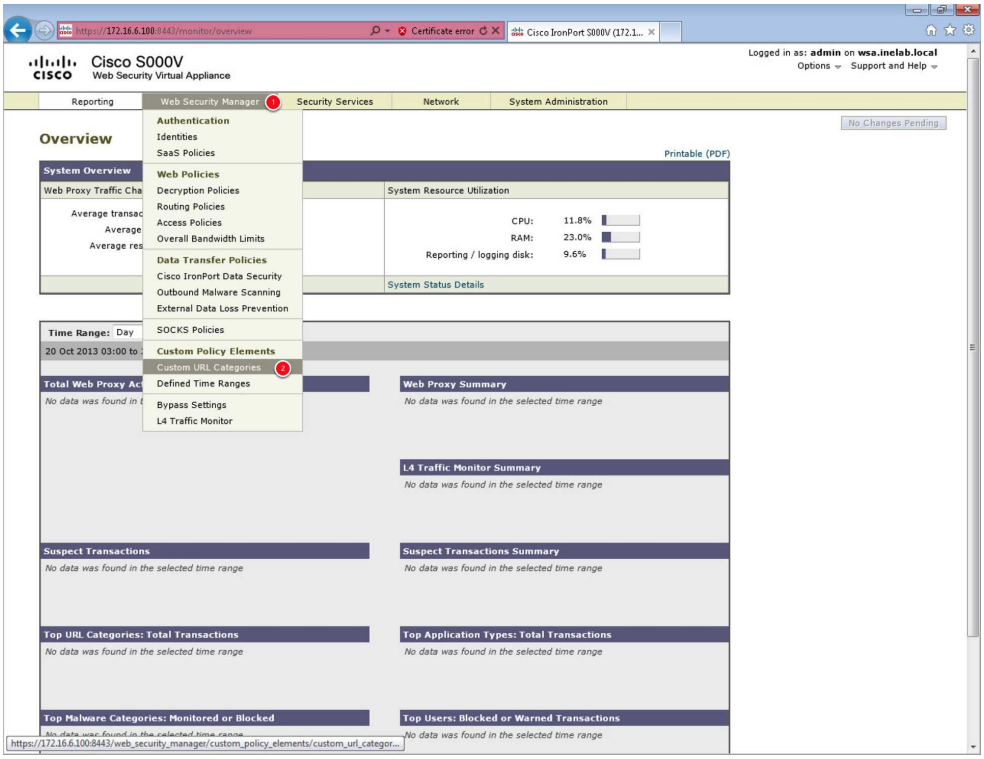
- **Warn**, which will present the user with an AUP (Acceptable use Policy) that must be accepted before access is allowed.

WSA passes a transaction through multiple steps before allowing or denying the connection. For HTTP sessions, the most important are, in order:

- Custom URL Category action
- Pre-defined URL Category action
- Uncategorized URL action
- Reputation Score action
- Antimalware action
- Application Visibility and Control action

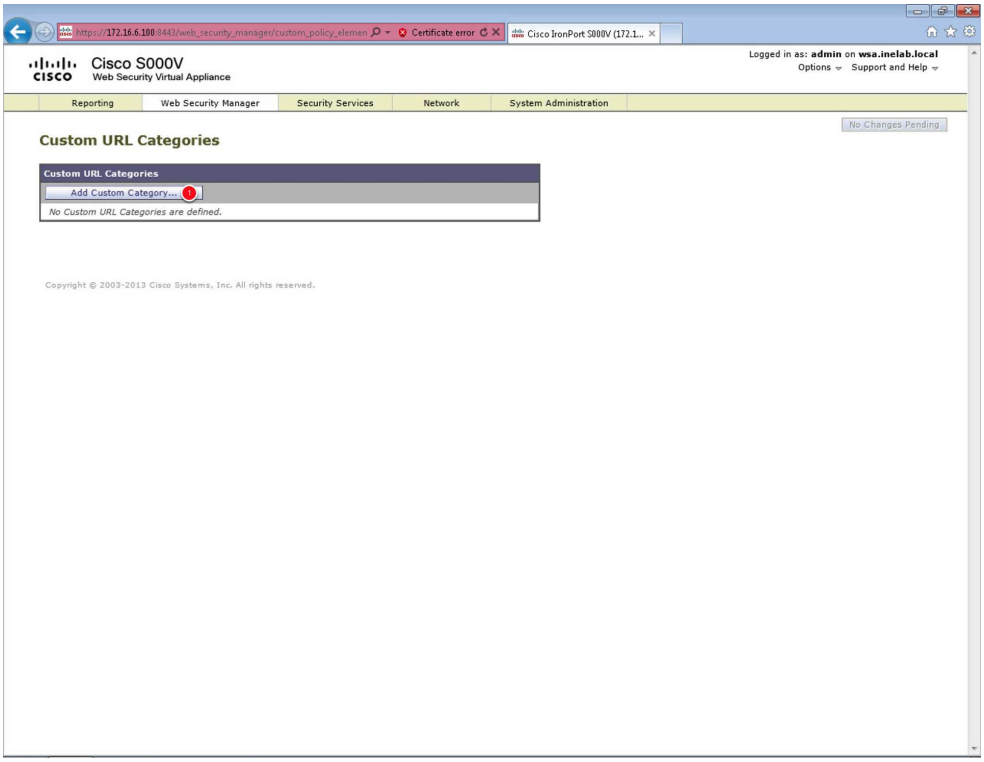
Configuration

Browse through the GUI to create a Custom URL Category:

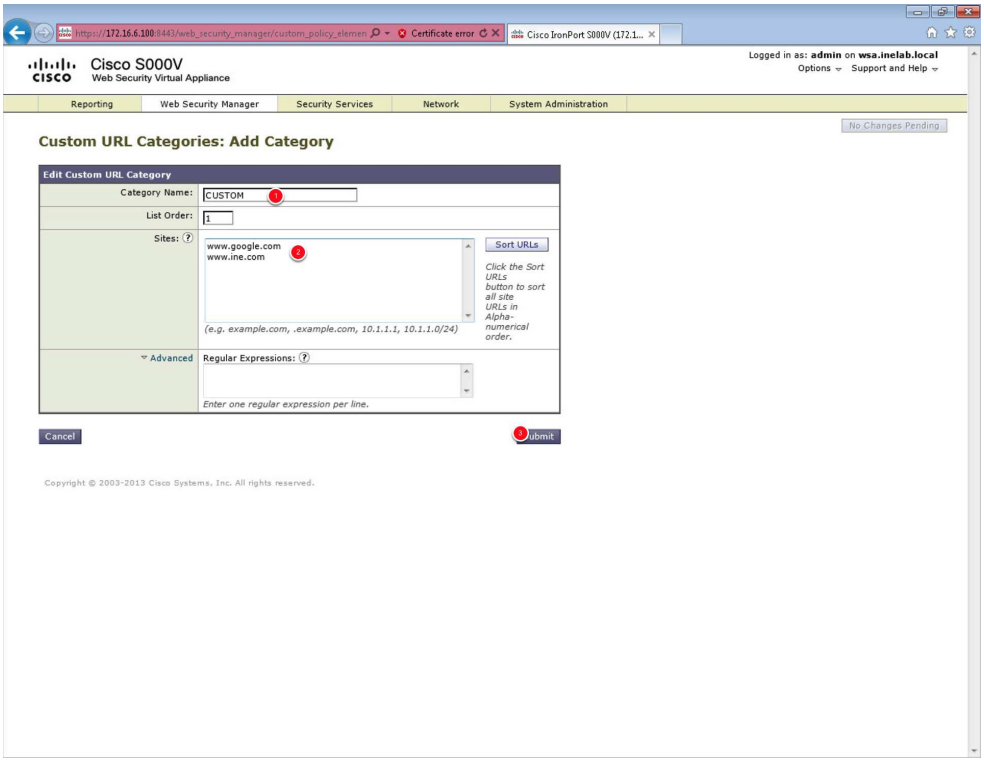


Click the **Add Custom Category** button:

https://t.me/learningnets

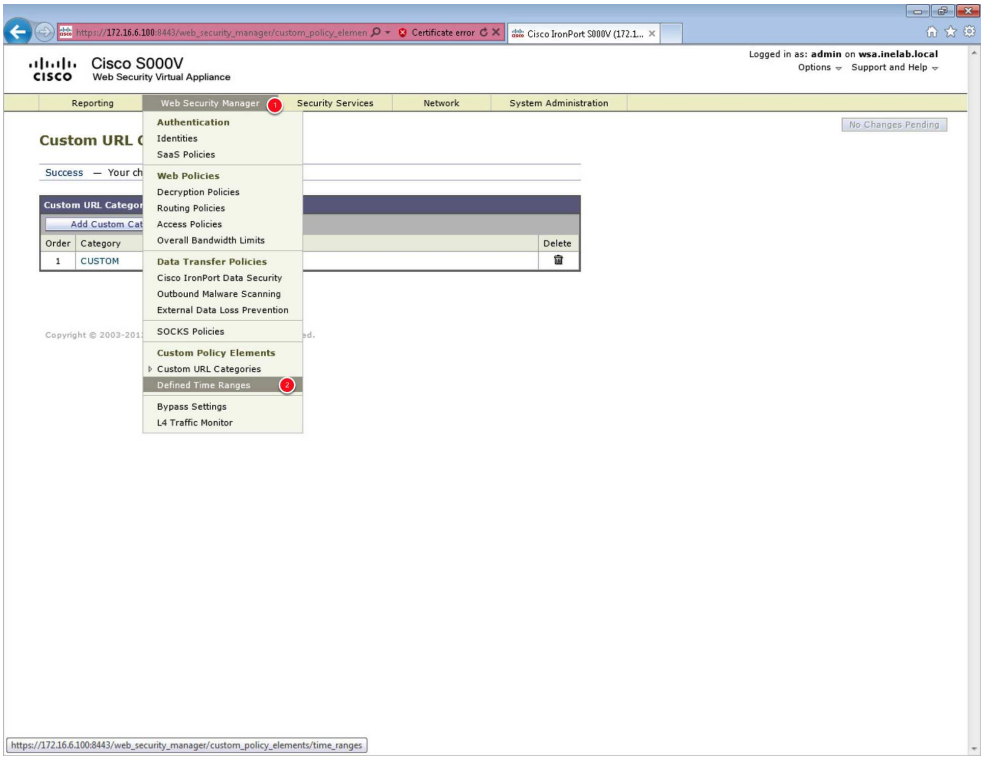


Add the two URLs to the URL category:

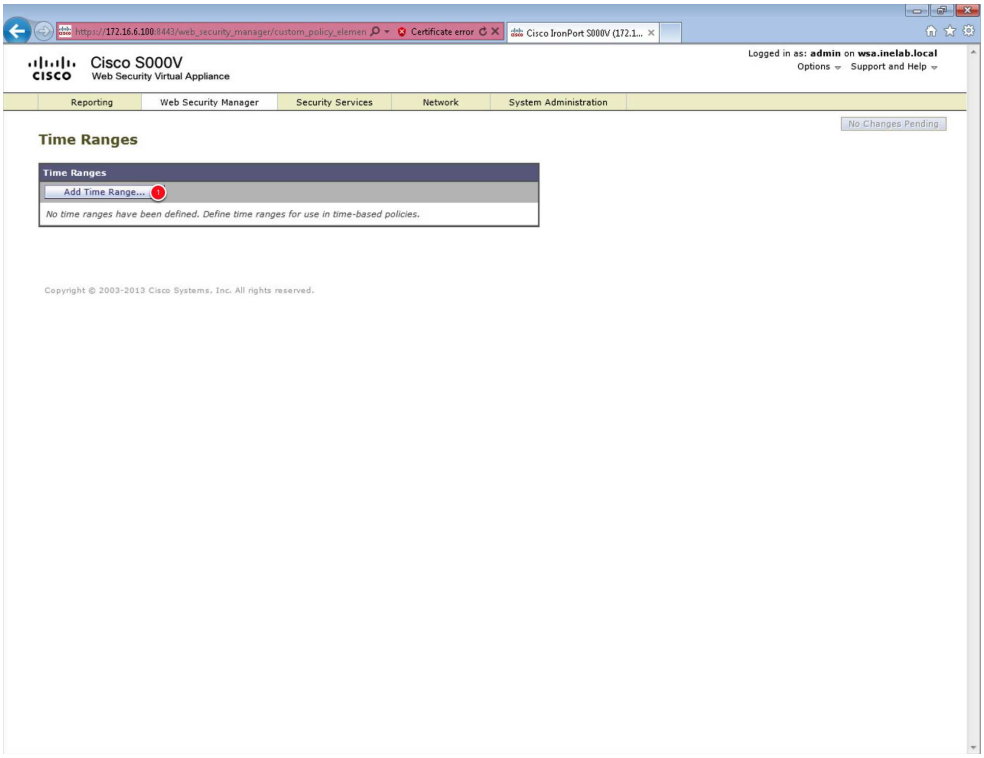


Browse through the GUI to create a time range for the weekend:

https://t.me/learningnets

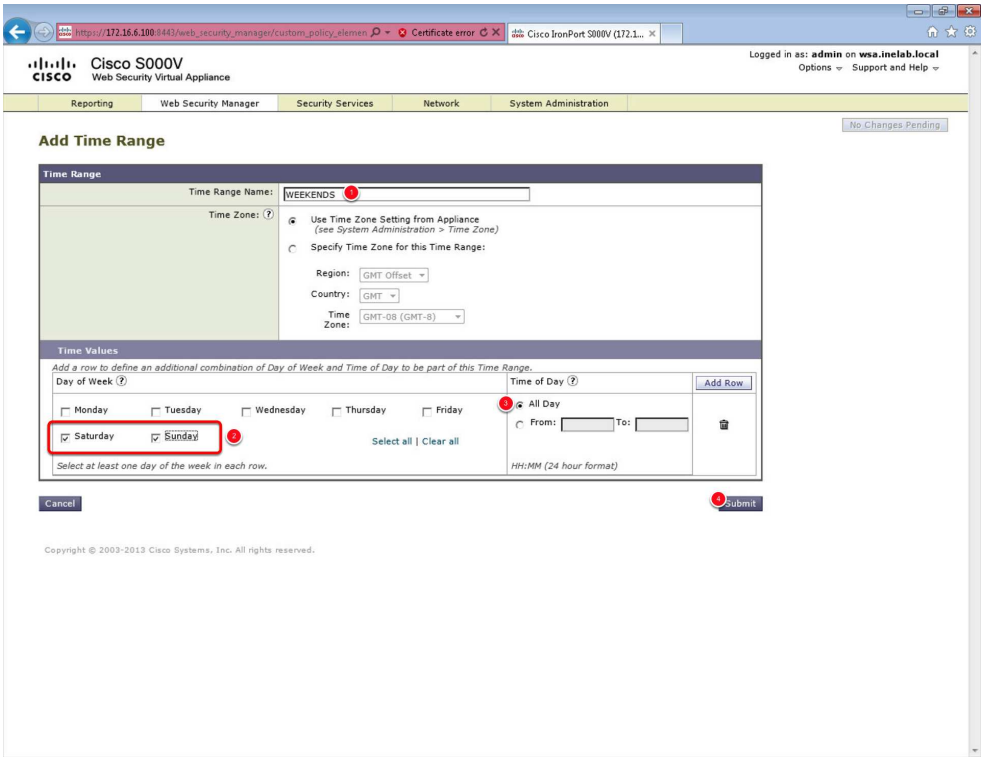


Click the **Add Time Range** button:

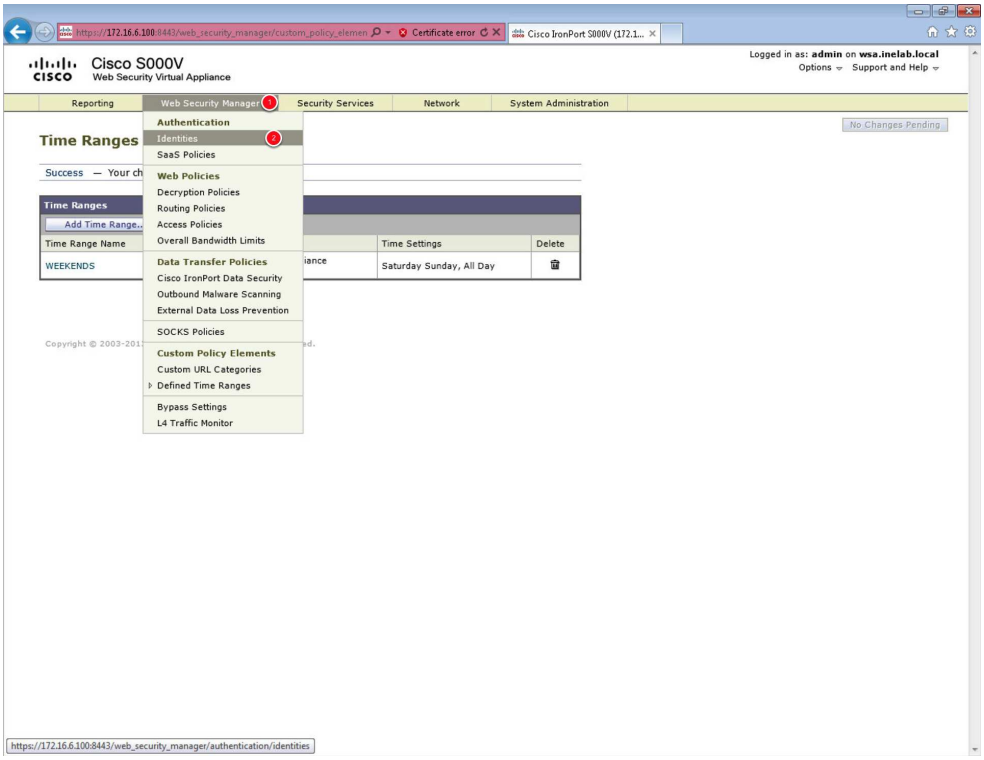


Define Saturday and Sunday as valid for the time range:

https://t.me/learningnets

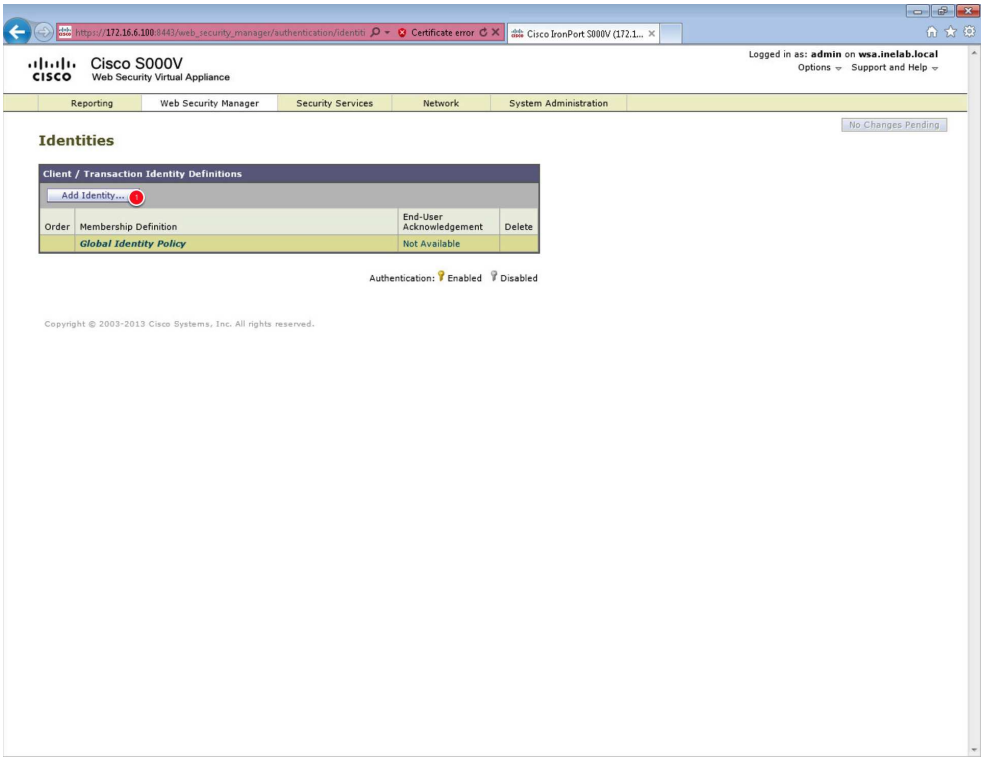


Browse through the GUI to configure an identity for **TEST-PC-B**:

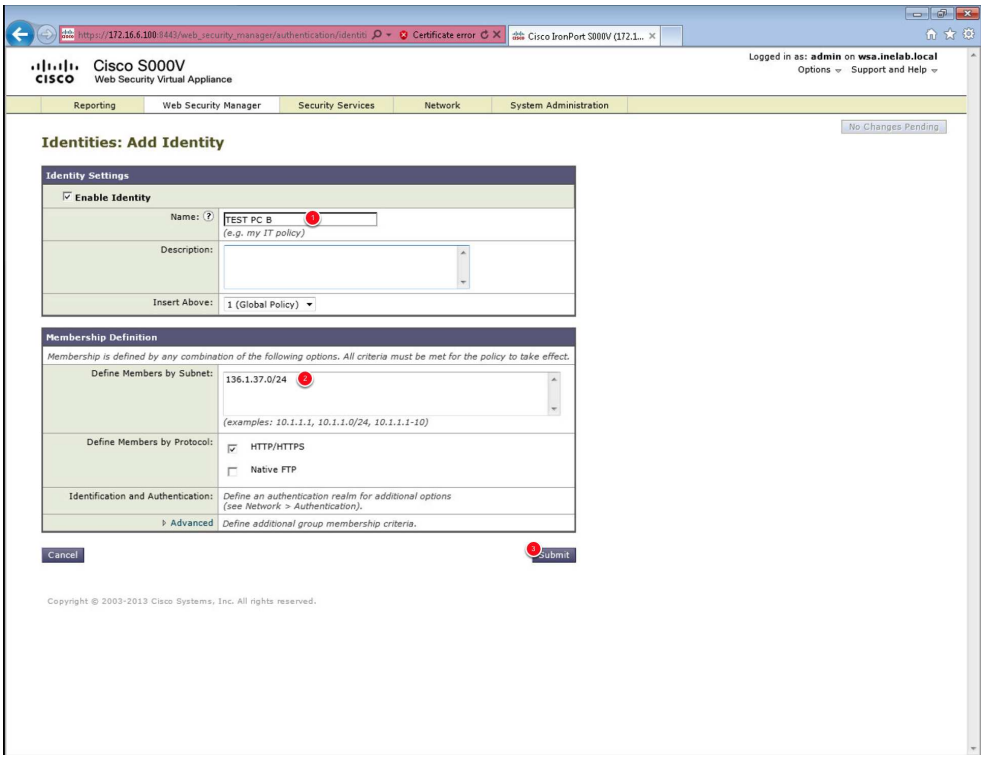


Click the **Add Identity** button:

https://t.me/learningnets

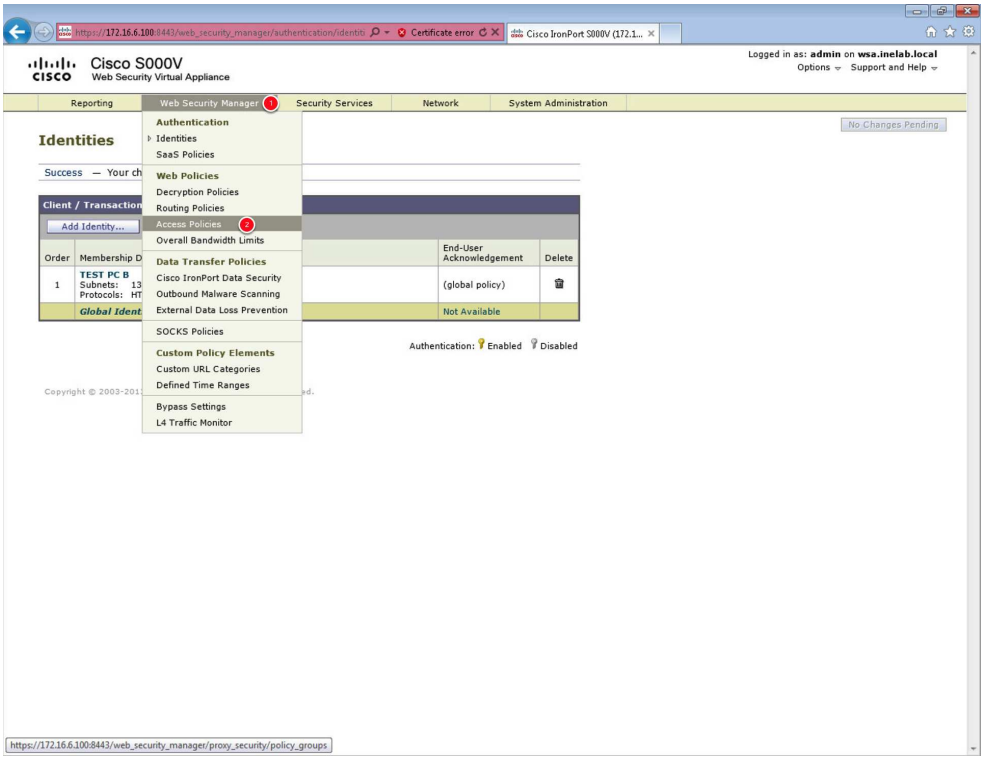


Because **TEST-PC-B** has a DHCP-assigned IP address, configure the complete subnet as the identity match:

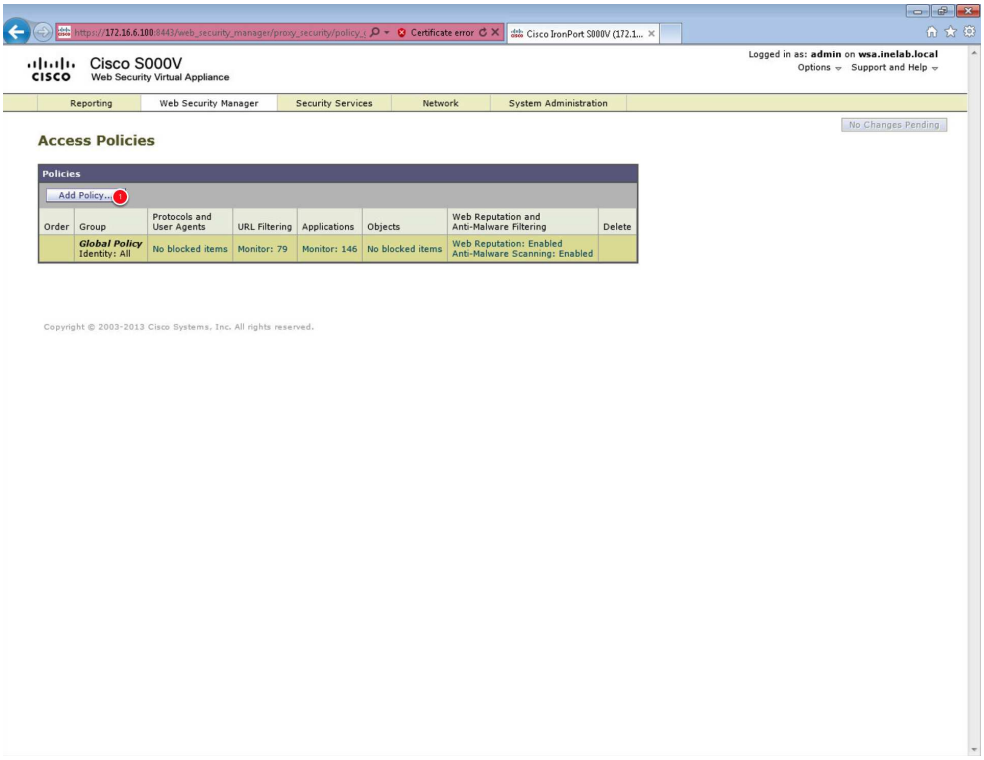


Browse through the GUI to configure an access policy:

https://t.me/learningnets

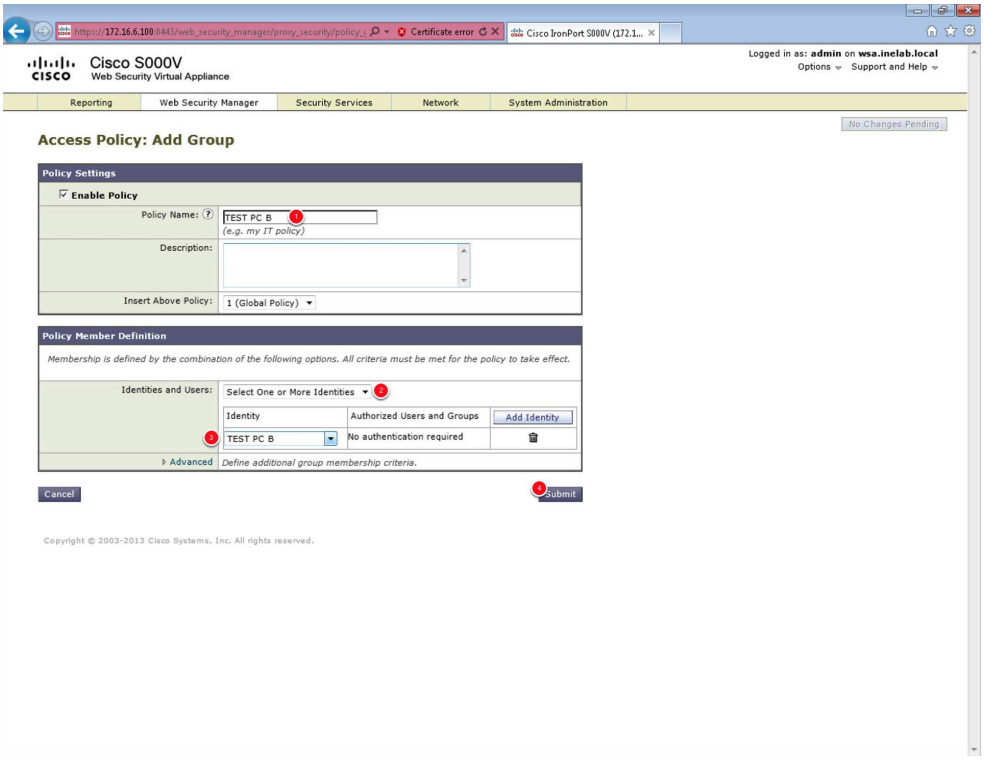


Click the **Add Policy** button:

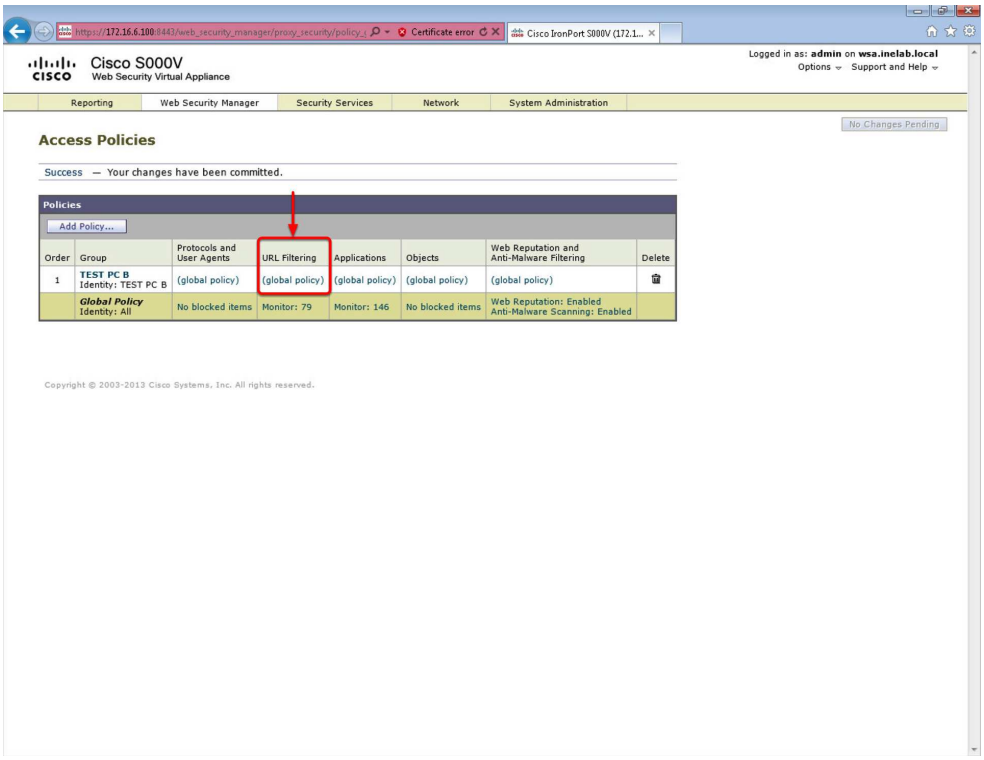


Select the previously defined identity as policy membership:

https://t.me/learningnets

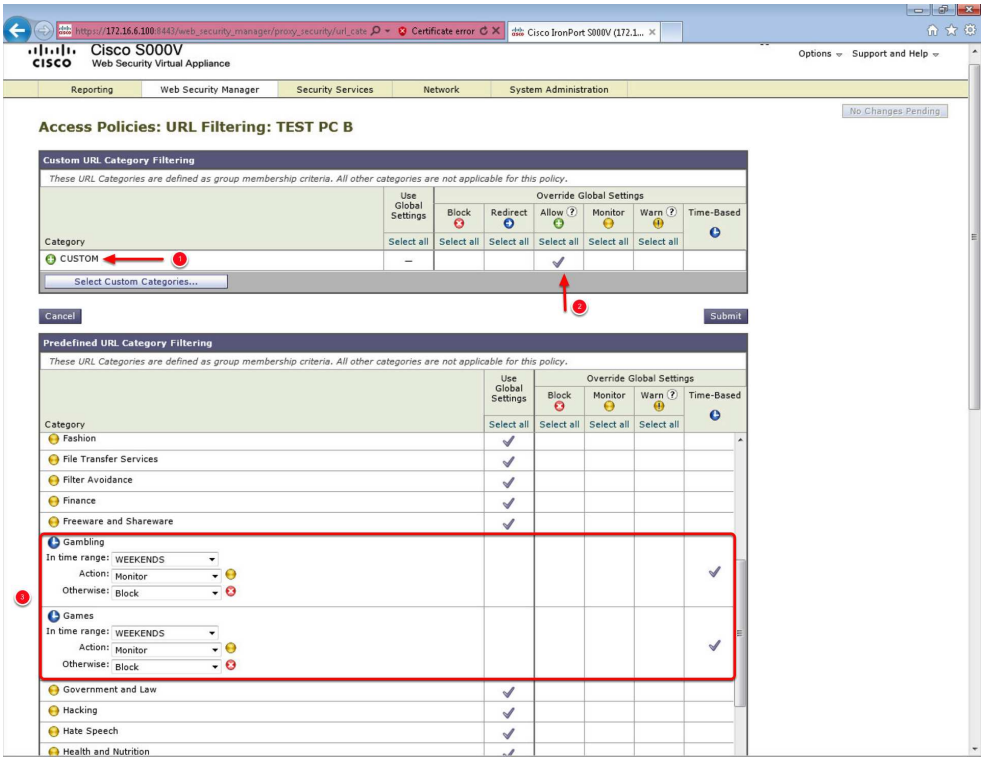


Edit the URL Filtering settings for the configured access policy:

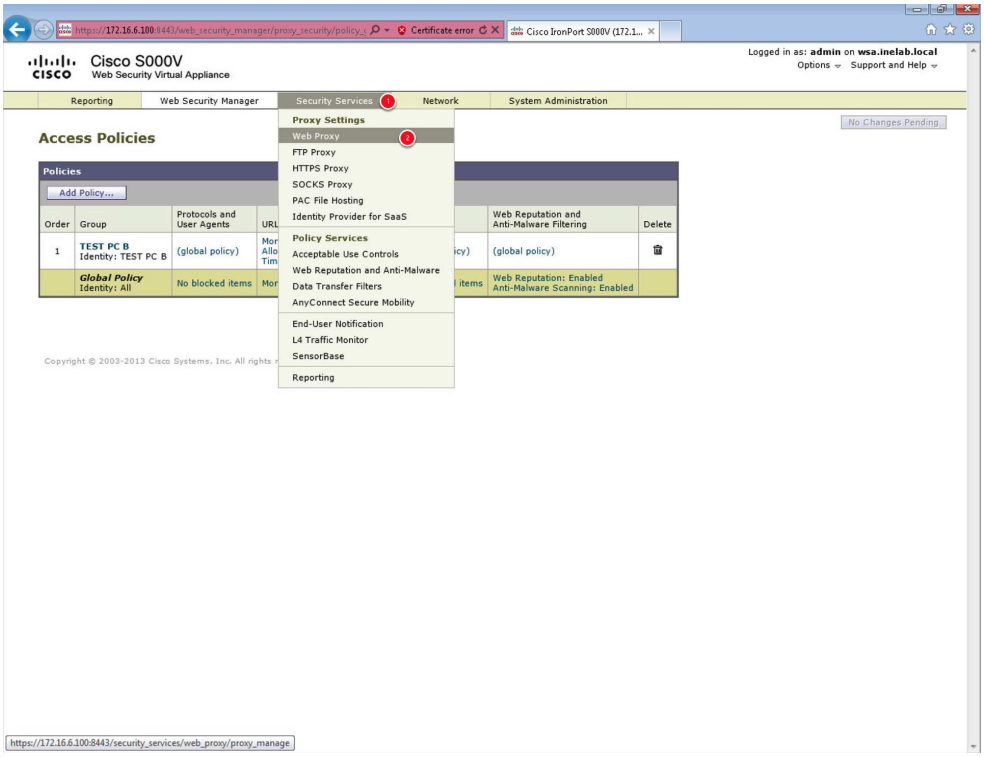


Set the Custom URL Category to **allow** so that all inspections are bypassed and the pre-defined URL categories to **monitor** within the allowed time range:

https://t.me/learningnets

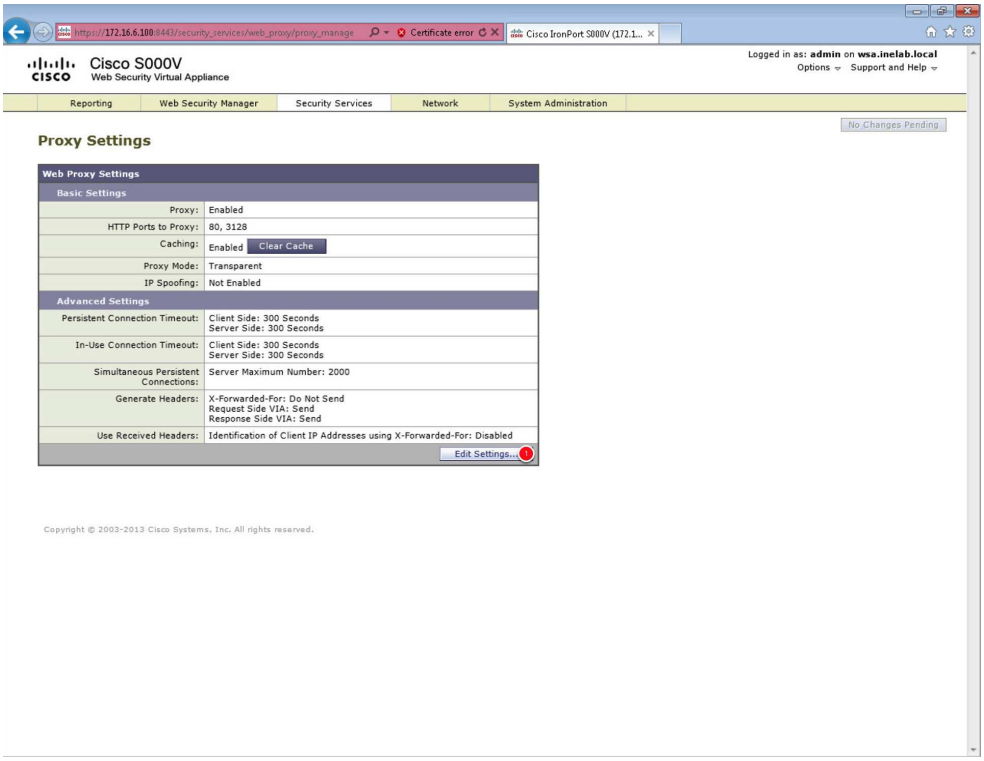


Browse through the GUI to edit the HTTP proxy settings:

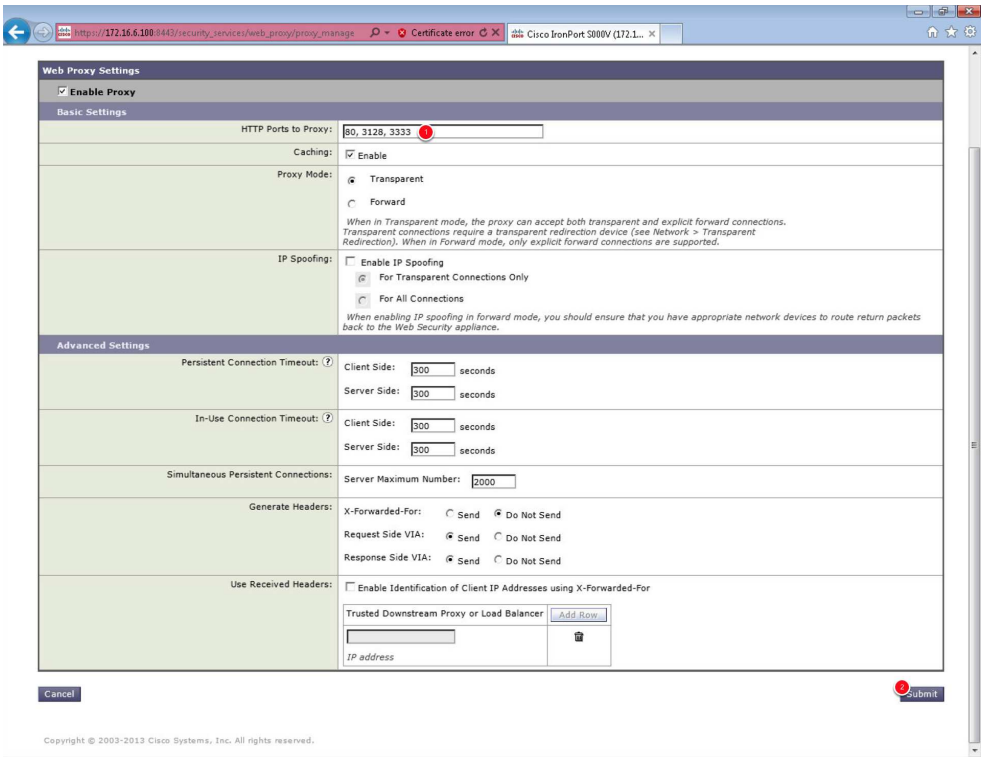


Click the **Edit Settings** button:

https://t.me/learningnets



Add port 3333 to the proxy service, so that WSA can serve client requests received on port 3333:



Verification

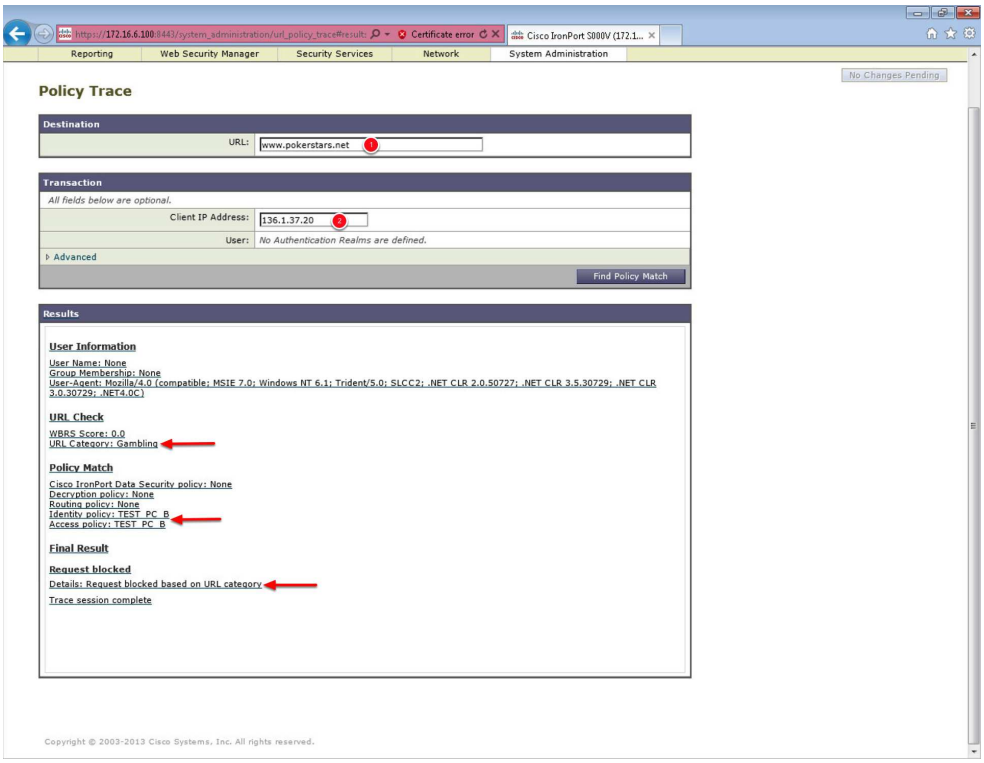
Check the time on the WSA:

https://t.me/learningnets

```
wsa.inelab.local> date

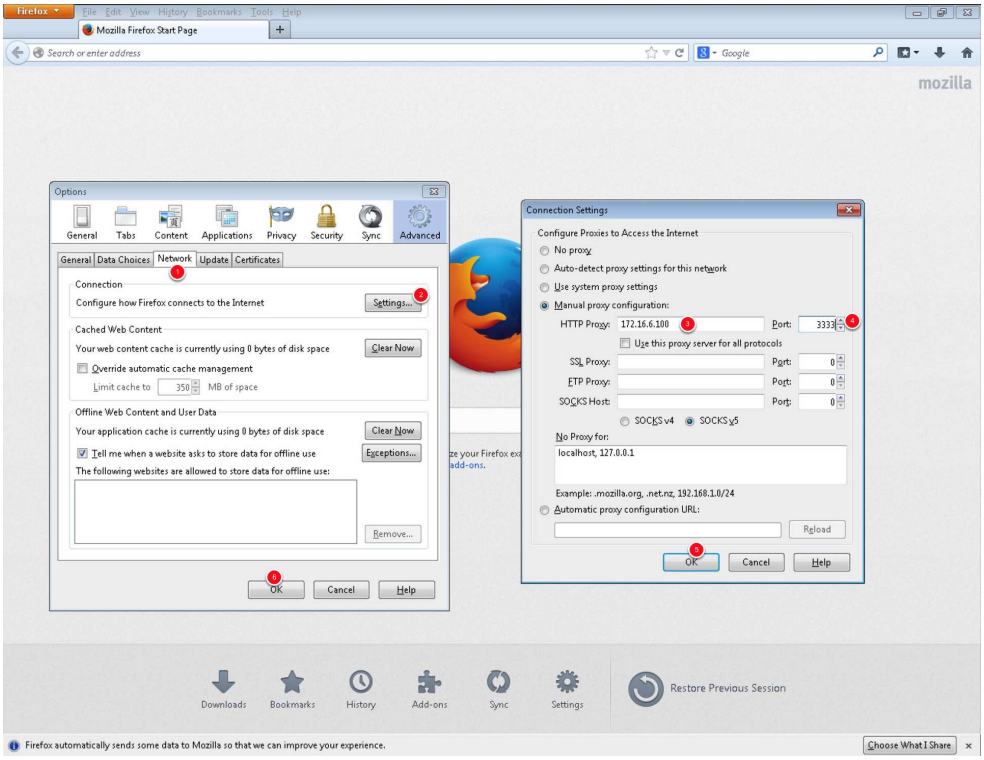
Mon Oct 21 04:05:35 2013 GMT-8
```

Based on the current time, access to the pre-defined URL categories should be blocked; test using the **Policy Trace** tool from the **System Administration** menu:

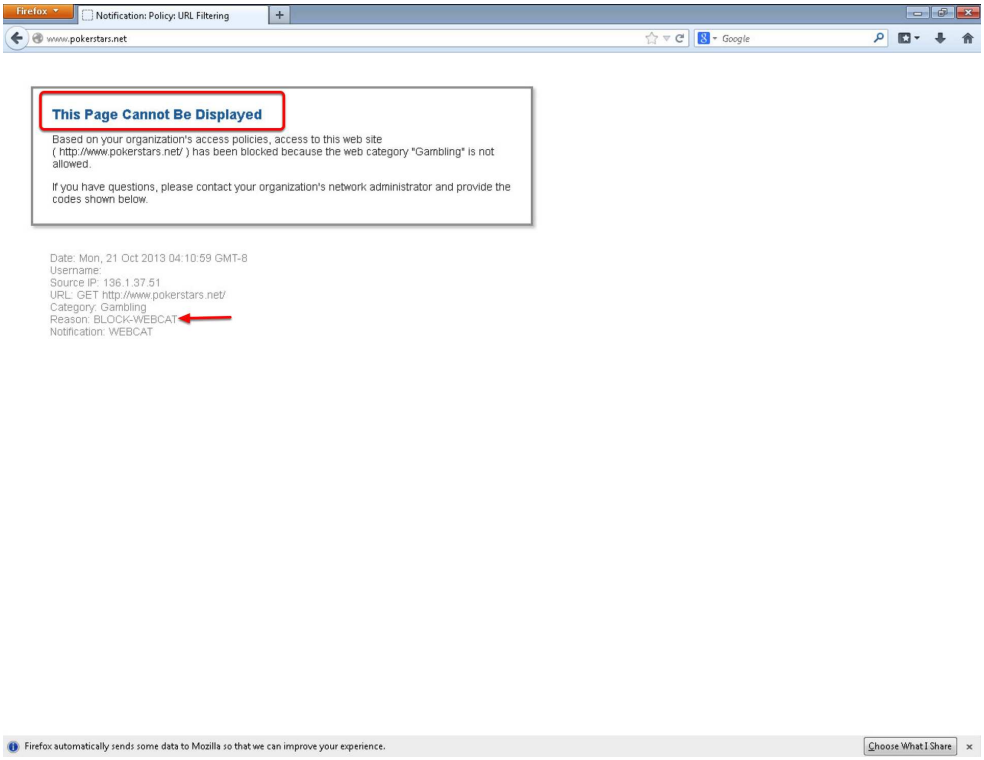


Configure Mozilla Firefox for proxy functionality:

https://t.me/learningnets



Using Firefox, browse to **www.pokerstars.net**; notice the reason for the request being blocked, namely **BLOCK-WEBCAT**, which means the request was denied because of URL Category settings:



Modify the time on the WSA so that the request matches the configured time range:

```
wsa.inelab.local> settime
```

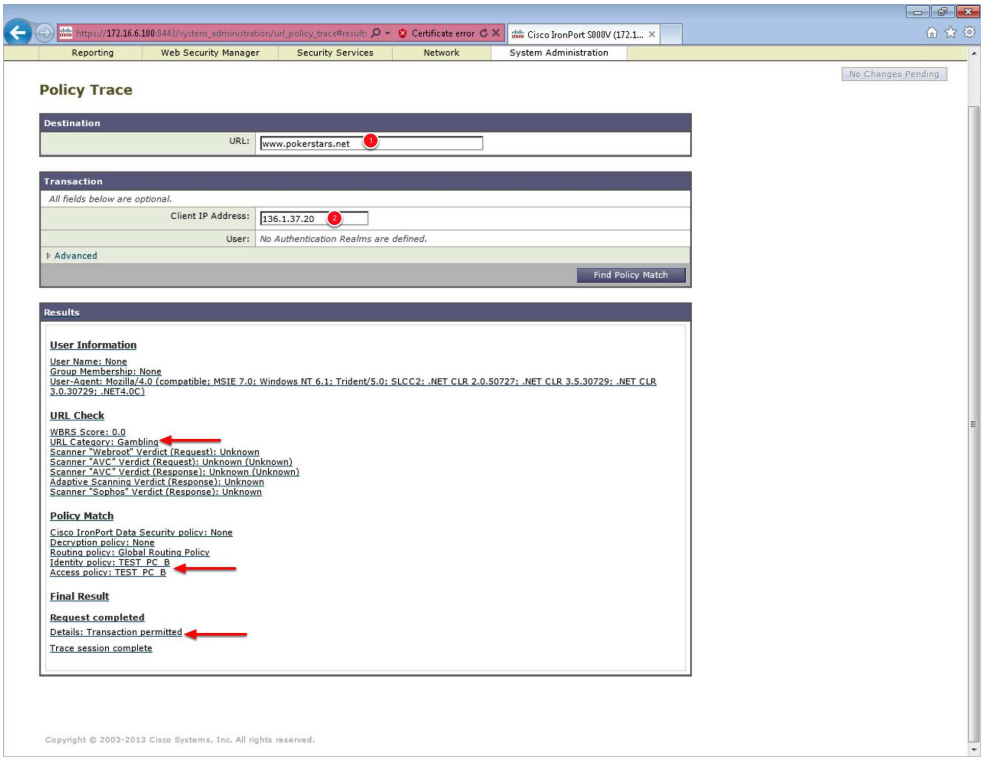
https://t.me/learningnets

```
WARNING: Changes to system time will take place immediately and do not
require the user to run the commit command.

This machine is currently running NTP.  In order to manually set the time,
NTP must be disabled.Do you want to stop NTP and manually set the time? [N]> Y

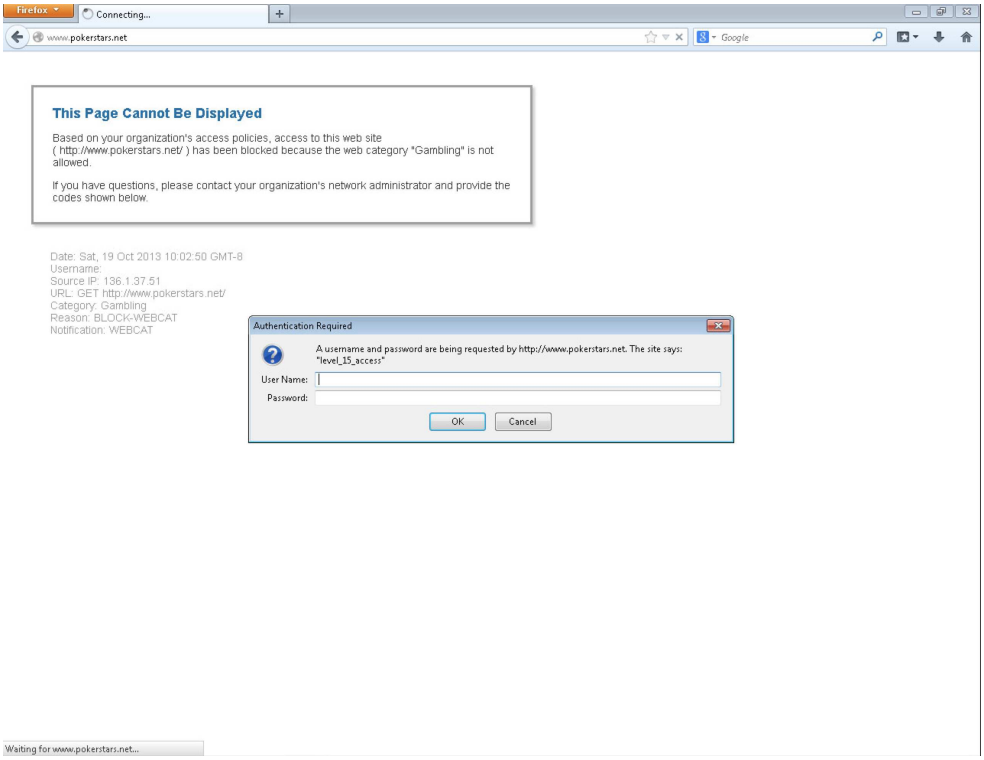
Current time Mon Oct 21 04:15:07 2013 GMT-8.
Please enter the time in MM/DD/YYYY HH:MM:SS format.[]> 10/20/2013 10:00:00
Time set to Sun Oct 20 10:00:00 2013 GMT-8.
```

Based on the current time, access to the pre-defined URL categories should be allowed; test using the **Policy Trace** tool from the **System Administration** menu:

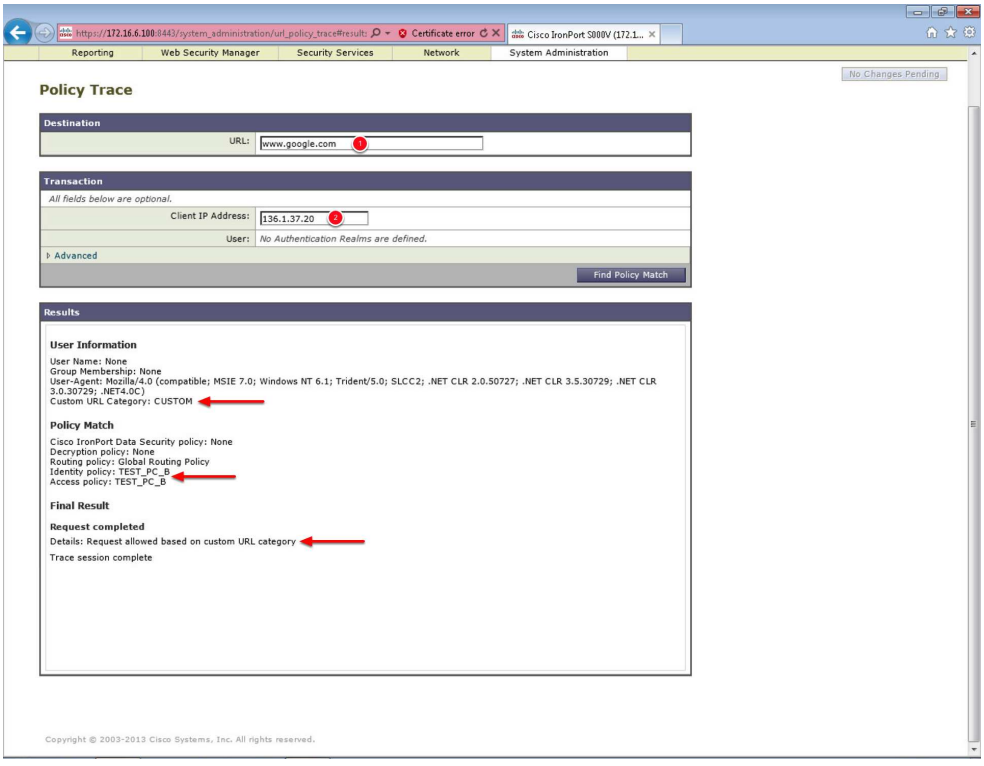


Using Firefox, browse again to **www.pokerstars.net**; this time the request should be allowed. The web server from R4 should ask for authentication:

https://t.me/learningnets



Using the **Policy Trace** tool from the **System Administration** menu, test access to one of the whitelisetd domains, such as **www.google.com**:



Verify the WSA **access logs** to confirm that the initial request to **www.pokerstars.net** was denied, and after the time change it was allowed, as well as for **www.google.com**; **TCP_MISS** means that content was not retrieved from the cache:

https://t.me/learningnets

```
wsa.inelab.local> tail


Currently configured logs:1. "accesslogs" Type: "Access Logs" Retrieval: FTP Poll
2. "authlogs" Type: "Authentication Framework Logs" Retrieval: FTP Poll
3. "avc_logs" Type: "AVC Engine Logs" Retrieval: FTP Poll
4. "bypasslogs" Type: "Proxy Bypass Logs" Retrieval: FTP Poll
5. "cli_logs" Type: "CLI Audit Logs" Retrieval: FTP Poll
6. "configdefragd_logs" Type: "Configuration Logs" Retrieval: FTP Poll
7. "dca_logs" Type: "DCA Engine Logs" Retrieval: FTP Poll
8. "external_auth_logs" Type: "External Authentication Logs" Retrieval: FTP
Poll
9. "feedback_logs" Type: "Feedback Logs" Retrieval: FTP Poll
10. "fips_logs" Type: "FIPS Logs" Retrieval: FTP Poll
11. "ftpd_logs" Type: "FTP Server Logs" Retrieval: FTP Poll
12. "gui_logs" Type: "GUI Logs" Retrieval: FTP Poll
13. "haystackd_logs" Type: "Haystack Logs" Retrieval: FTP Poll
14. "idsdataloss_logs" Type: "Data Security Logs" Retrieval: FTP Poll
15. "logderrorlogs" Type: "Logging Logs" Retrieval: FTP Poll
16. "mcafee_logs" Type: "McAfee Logs" Retrieval: FTP Poll
17. "musd_logs" Type: "AnyConnect Secure Mobility Daemon Logs" Retrieval: FTP
Poll
18. "ocspd_logs" Type: "OCSP Logs" Retrieval: FTP Poll
19. "pacd_logs" Type: "PAC File Hosting Daemon Logs" Retrieval: FTP Poll
20. "proxylogs" Type: "Default Proxy Logs" Retrieval: FTP Poll
21. "reportd_logs" Type: "Reporting Logs" Retrieval: FTP Poll
22. "reportqueryd_logs" Type: "Reporting Query Logs" Retrieval: FTP Poll
23. "saas_auth_log" Type: "SaaS Auth Logs" Retrieval: FTP Poll
24. "shd_logs" Type: "SHD Logs" Retrieval: FTP Poll
25. "snmp_logs" Type: "SNMP Logs" Retrieval: FTP Poll
26. "sntpd_logs" Type: "NTP Logs" Retrieval: FTP Poll
27. "sophos_logs" Type: "Sophos Logs" Retrieval: FTP Poll
28. "status" Type: "Status Logs" Retrieval: FTP Poll
29. "system_logs" Type: "System Logs" Retrieval: FTP Poll
30. "trafmon_errlogs" Type: "Traffic Monitor Error Logs" Retrieval: FTP Poll
31. "trafmonlogs" Type: "Traffic Monitor Logs" Retrieval: FTP Poll
32. "uds_logs" Type: "UDS Logs" Retrieval: FTP Poll
33. "updater_logs" Type: "Updater Logs" Retrieval: FTP Poll
34. "wbnp_logs" Type: "WBNP Logs" Retrieval: FTP Poll
35. "webcat_logs" Type: "Web Categorization Logs" Retrieval: FTP Poll
36. "webrootlogs" Type: "Webroot Logs" Retrieval: FTP Poll
37. "welcomeack_logs" Type: "Welcome Page Acknowledgement Logs" Retrieval:
FTP Poll
Enter the number of the log you wish to tail.[]> 1

Press Ctrl-C to stop.
```

https://t.me/learningnets

```
#Date: 2013-10-20 00:12:28
#System: 172.16.6.100 - wsa.inelab.local
#Software: AsyncOS for Web 7.7.5-190
#Fields: %t %e %a %w/%h %s %2r %A %H/%d %c %D %Xr %?BLOCK_SUSPECT_USER_AGENT,MONITOR_SUSPECT_USER_AGENT?%<User-Agent
1382299859.540 0 136.1.37.51
TCP_DENIED/403 0 GET http://www.pokerstars.net/ - NONE/- - BLOCK_WEBCAT_12-TEST_PC_B-TEST_PC_B-NONE-NONE-NONE-NONE
 <IW_gamb,0.0,-,"-",-,-,-,-,"-",-,-,-,"-",-,-,"-","-",-,-,IW_gamb,-,"-","-","Unknown","Unknown","-","-",0.00,0,-,"-"
1382234447.159 19 136.1.37.51
TCP_MISS/401 278 GET http://www.pokerstars.net/ - DIRECT/www.pokerstars.net text/plain DEFAULT_CASE_12-TEST_PC_B-TES
 <IW_gamb,0.0,0,"-",0,0,0,1,"-",-,-,-,"-",0,262686,"-","-",-,-,IW_gamb,-,"Unknown","-","Unknown","Unknown","-","-",1
1382234584.644 9 136.1.37.51
TCP_MISS/401 278 GET http://www.google.com/ - DIRECT/www.google.com - ALLOW_CUSTOMCAT_12-TEST_PC_B-TEST_PC_B-NONE-NO
 <C_CUST,-,-,"-",-,-,-,-,"-",-,-,-,"-",-,-,"-","-",-,-,-,-,"-","-","-","-","-","-",247.11,0,-,"-","-"> -
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Active Directory Integration

Task

- Configure the WSA for Active Directory integration with **TEST-SRV-B**; use **administrator/Cisco1234** as credentials for domain joining.
- After the first succesfull authentication, user credentials should be cached.
- Without defining any identities:
  - Inspect and allow access to all URLs only to user **USER3** from **CCIELAB** group, but ensure that a URL with a WBRS of -4 is blocked.
  - All other users (including those with failed authentication) should be allowed access only to categorized URLs with a WBRS higher than or equal to zero, without performing any antimalware scans.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the WSA to the initial configuration (not to factory defaults), load it again by navigating to **System Administration > Configuration File > Load Configuration**, select the **WSA-DEFAULT.xml** file, click **Load**, and commit changes afterward.

Overview

https://t.me/learningnets

WSA supports authentication through LDAP or NTLM (Active Directory). The authentication methods are:

basic: Credentials are sent unencrypted, in clear-text.
NTLMSSP: Credentials are never sent over the wire, based on challenge-response mechanism.
TUI: Transparent user authentication is based on an agent reading the Active Directory audit logs and mapping users logged in to AD to their IP addresses.

For the purpose of the CCIE lab, the first two methods are of interest. Because WSA integration with Active Directory follows the same behavior as a regular PC, which actually joins the domain and creates an object in the Active Directory schema, WSA can only integrate with a single domain/forest at this moment. With LDAP, this is not the case; multiple concurrent LDAP servers are supported at the same time, because the authentication process is solely based on a query/response mechanism.

When the WSA is configured for LDAP and/or AD user authentication, the authentication itself is enforced and applied at the identity level. By default, if users who fails authentication are denied any access, unless guest support is enabled at the identity level, in which case an access policy can be configured for users failing authentication; otherwise, the default global access policy settings will apply.

Configuration

Set the time on WSA to match the time on **TEST-SRV-B**, the Domain Controller; otherwise, AD integration is not possible:

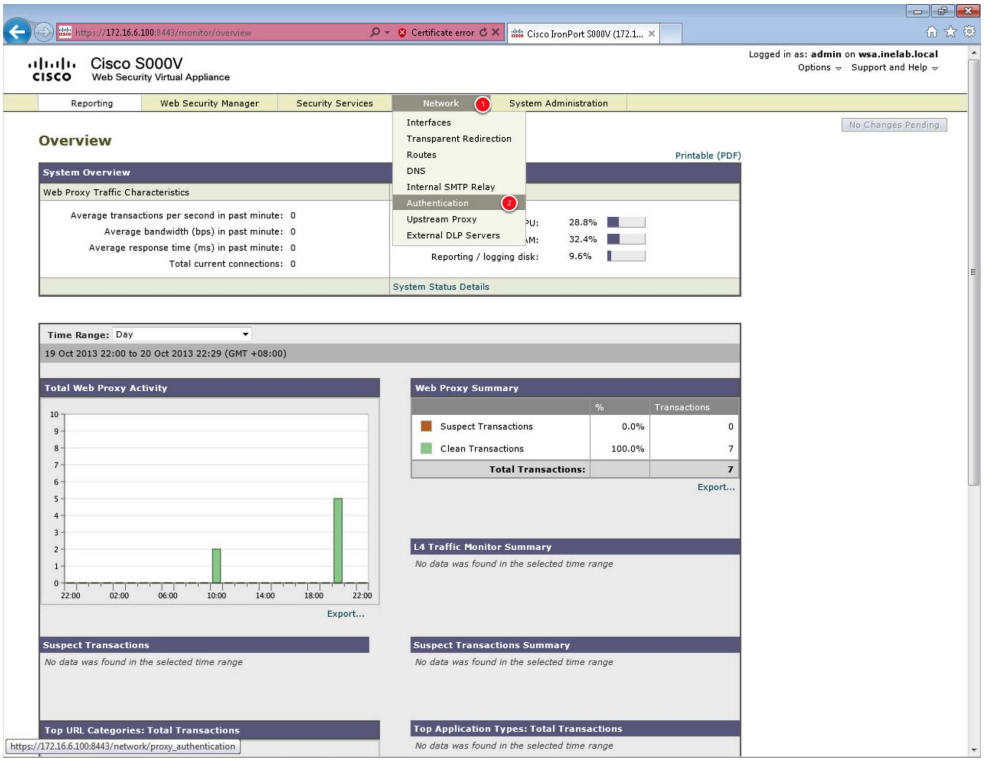
```
wsa.inelab.local> settime


WARNING: Changes to system time will take place immediately and do not
require the user to run the commit command.

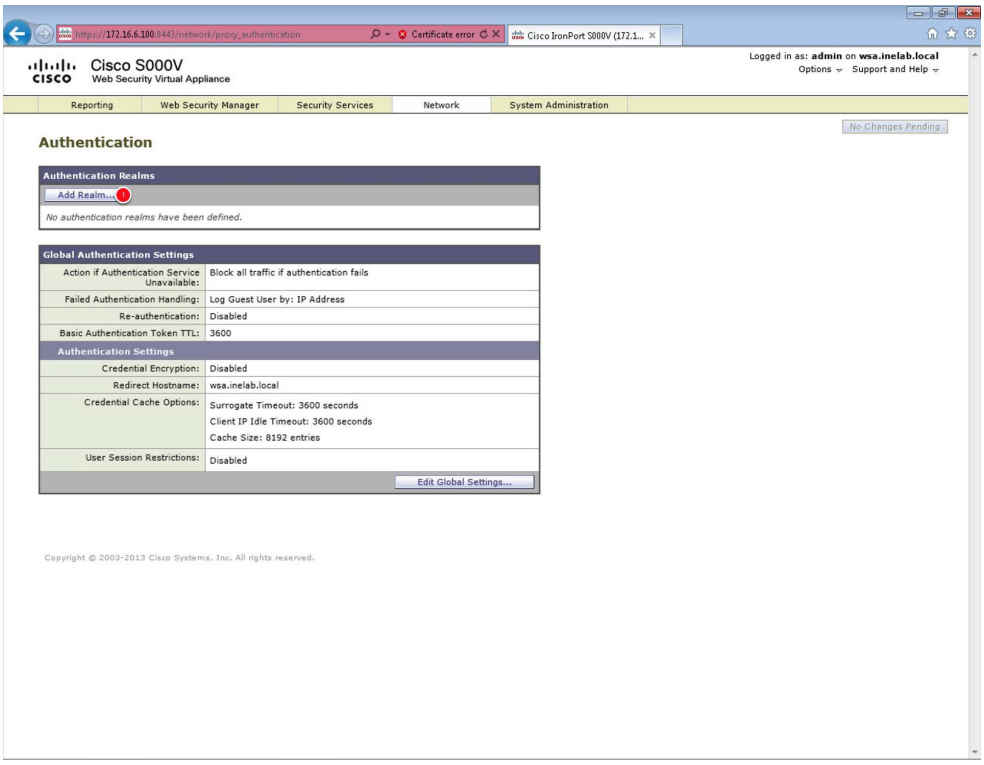
Current time Mon Oct 21 02:01:32 2013 GMT-8.
Please enter the time in MM/DD/YYYY HH:MM:SS format.[]> 10/21/2013 17:02:00
Time set to Mon Oct 21 17:02:00 2013 GMT-8.
```

Navigate through the GUI to configure integration with an LDAP/AD server for user authentication:

https://t.me/learningnets

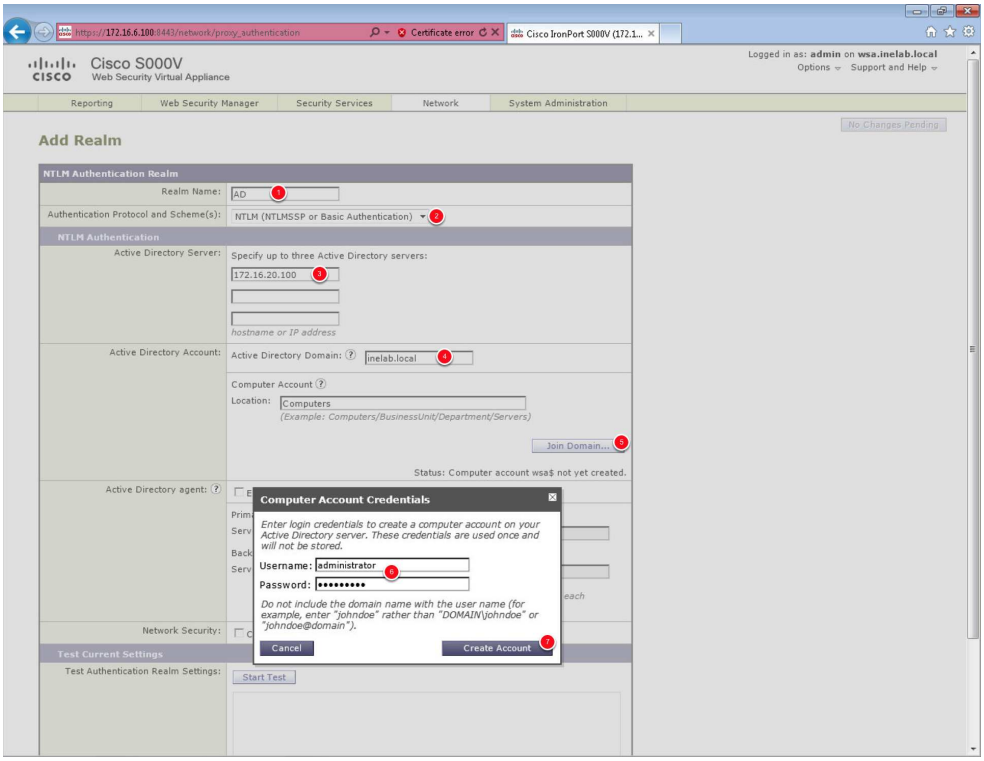


Click the **Add Realm** button:

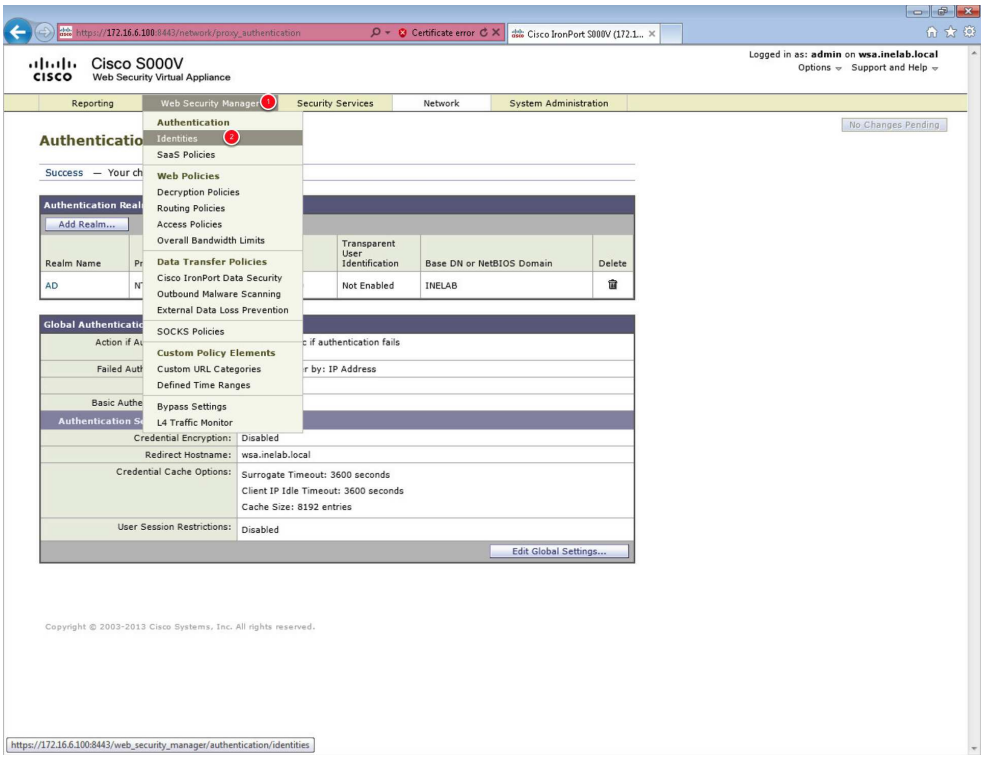


Select NTLM as the authentication scheme, configure the domain name and IP address of domain controller, and join WSA to AD:

https://t.me/learningnets

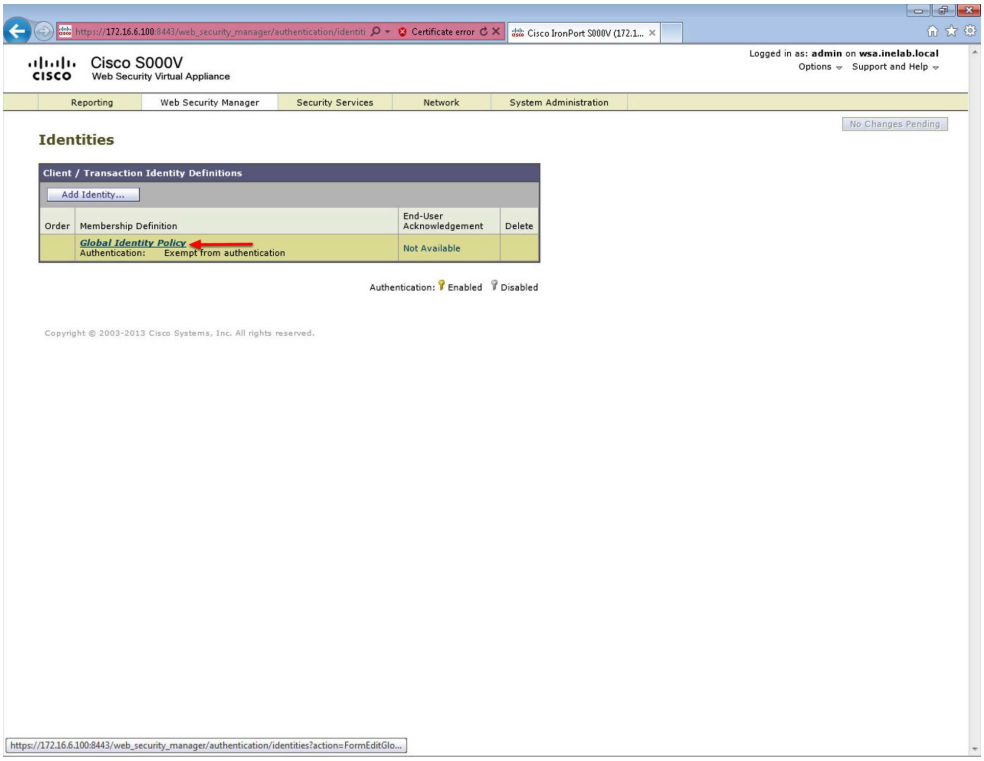


Navigate through the GUI to enable authentication with guest privileges for the global identity policy:

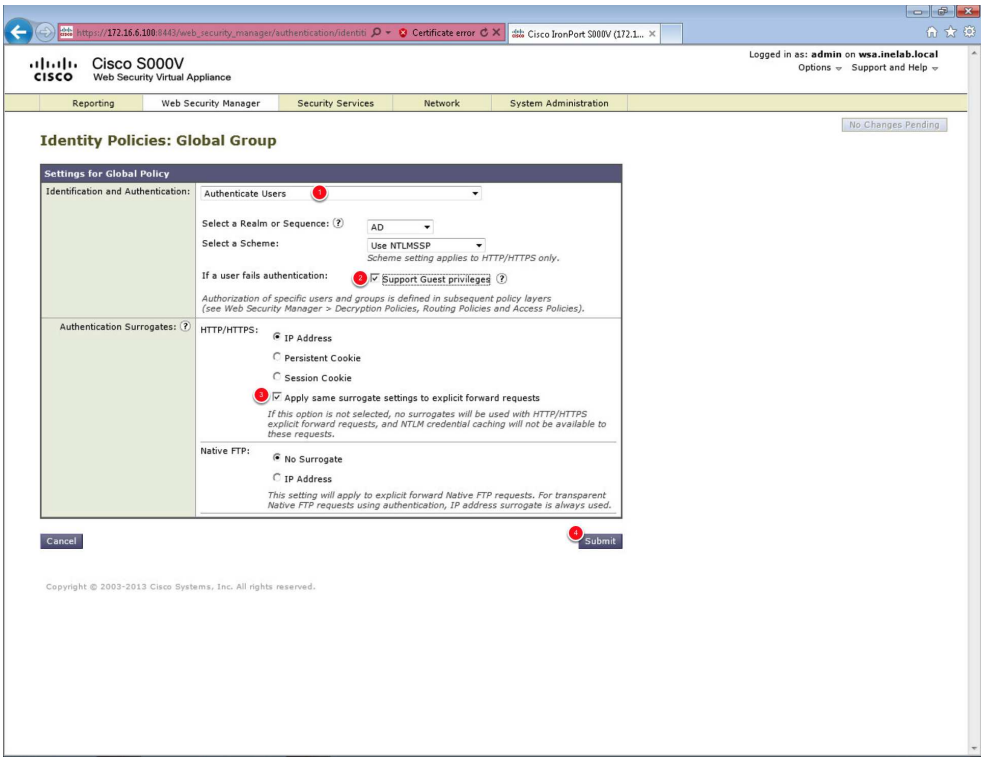


Click the **Global Identity Policy** hyperlink:

https://t.me/learningnets

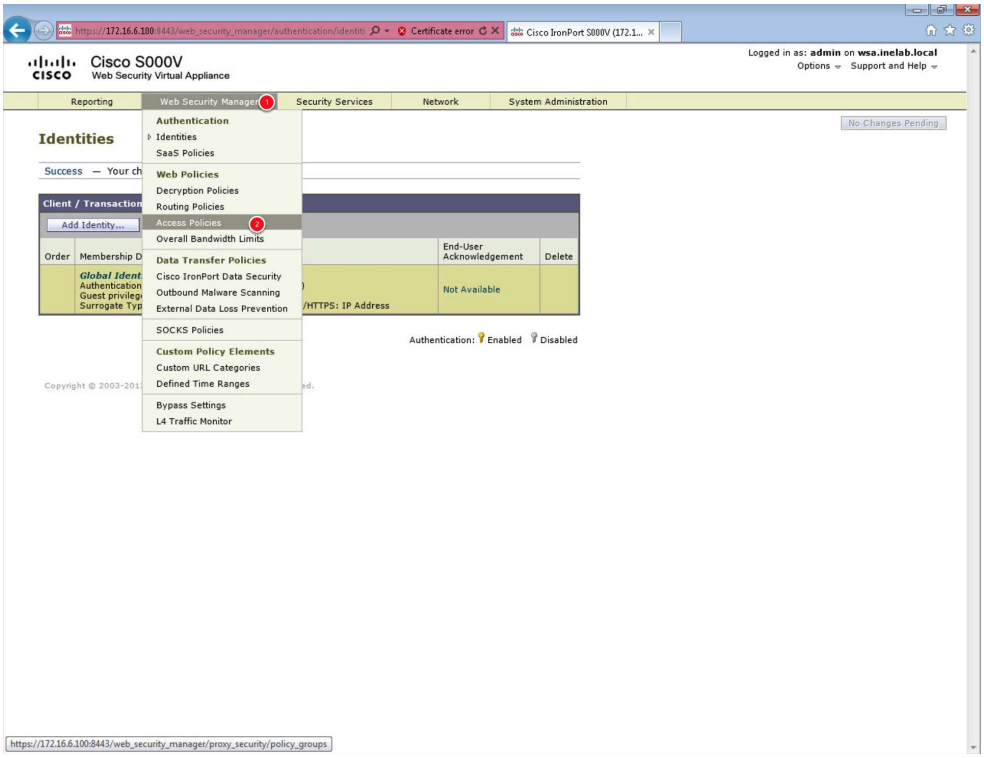


Enable authentication and guest support, and allow for NTLM credential caching:

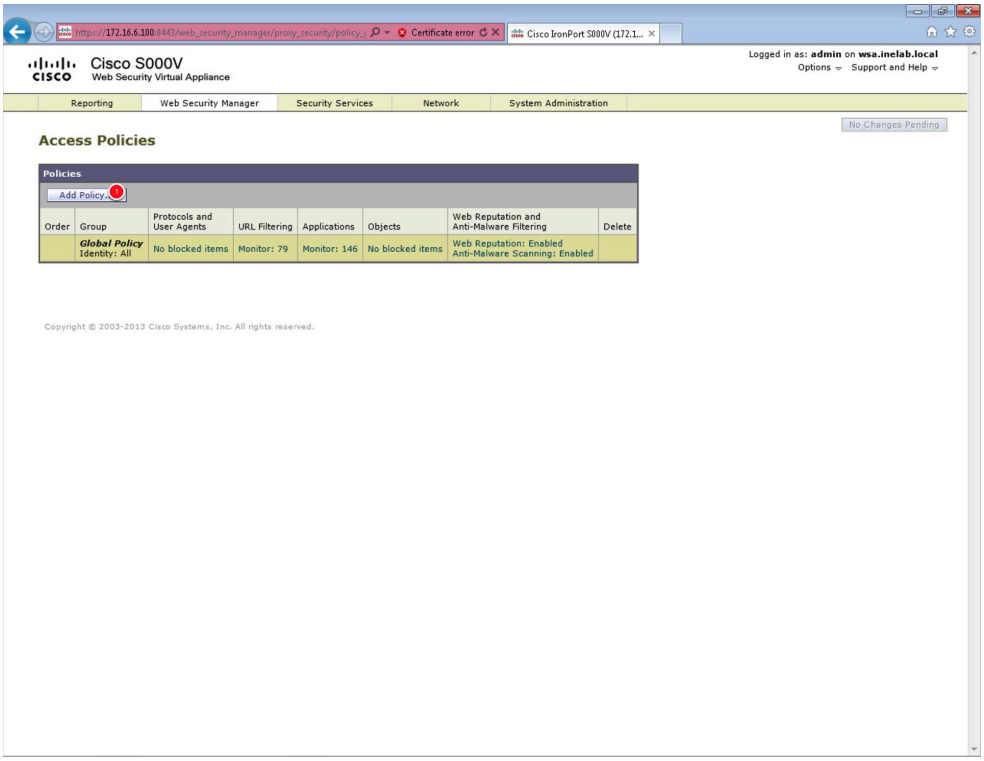


Navigate through the GUI to configure an access policy for the authenticated users:

https://t.me/learningnets

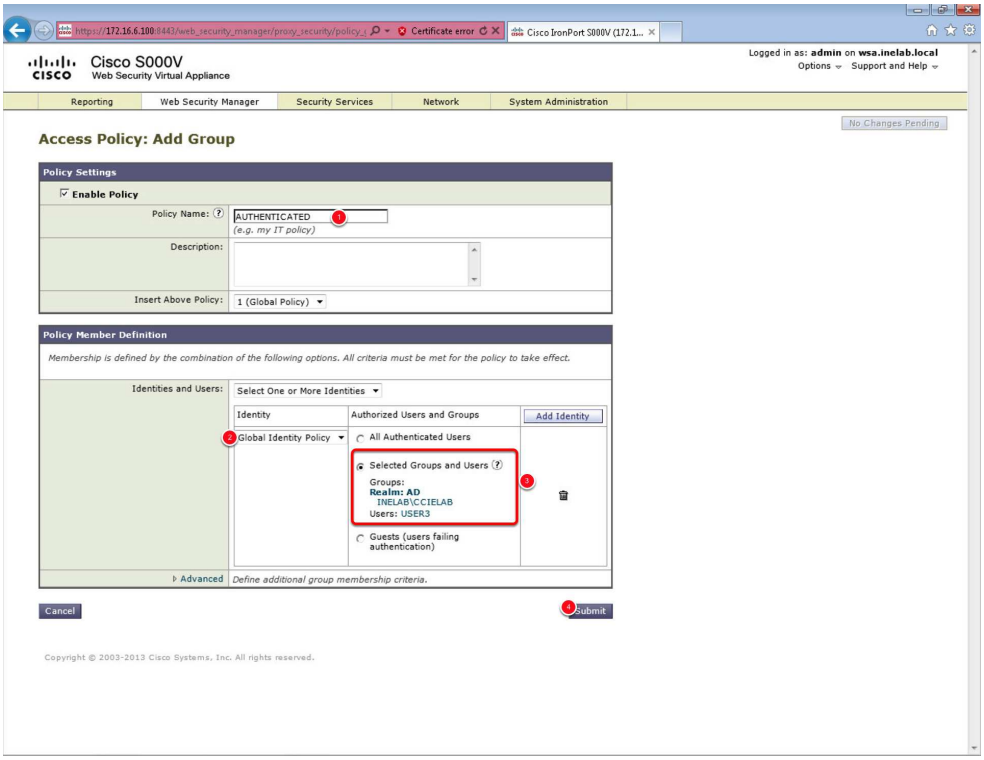


Click the **Add Policy** button:

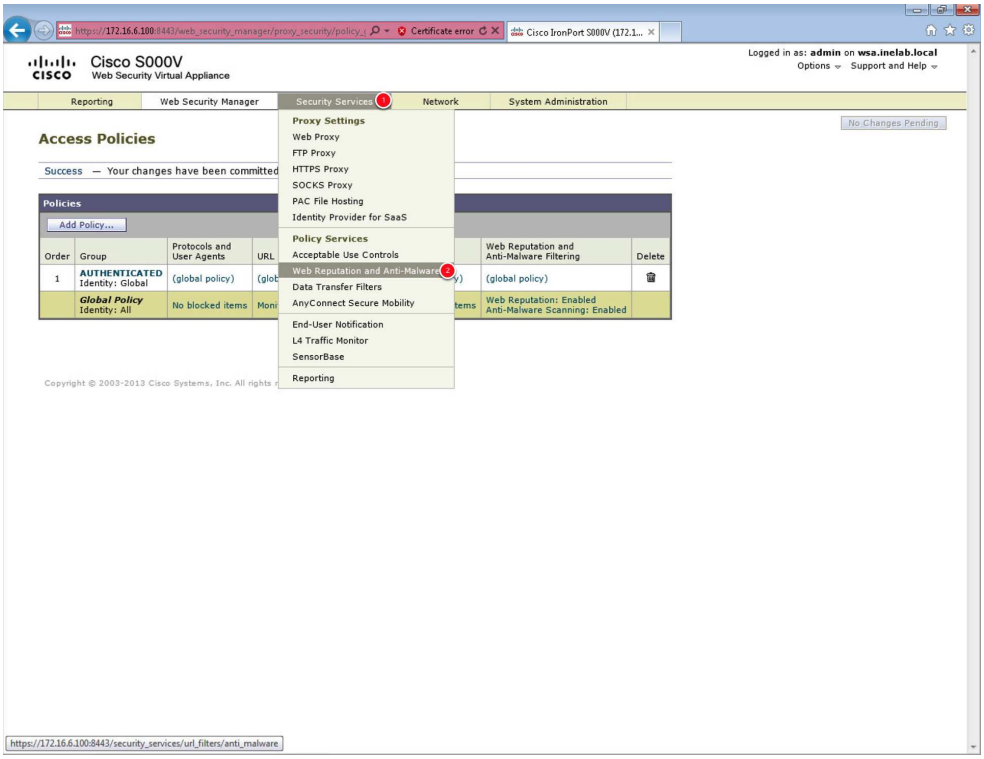


Select the global identity and allow access only to user **USER3** if it's a member of the **CCIELAB** group:

https://t.me/learningnets

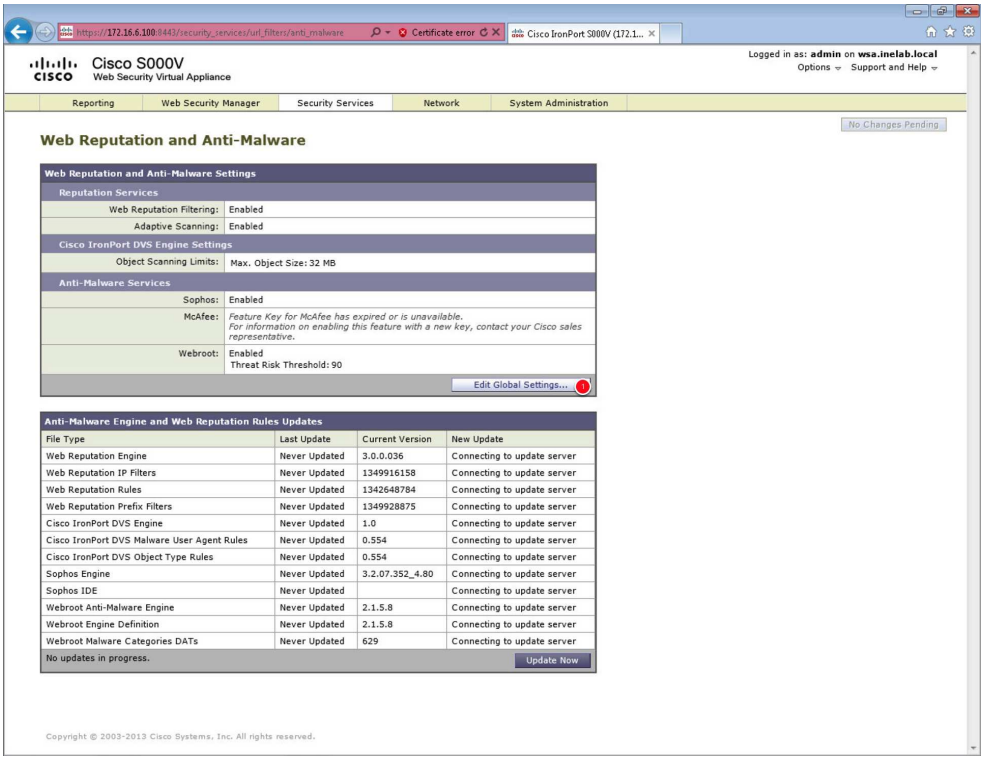


To modify the default WBRS settings, adaptive scanning must be disabled first:

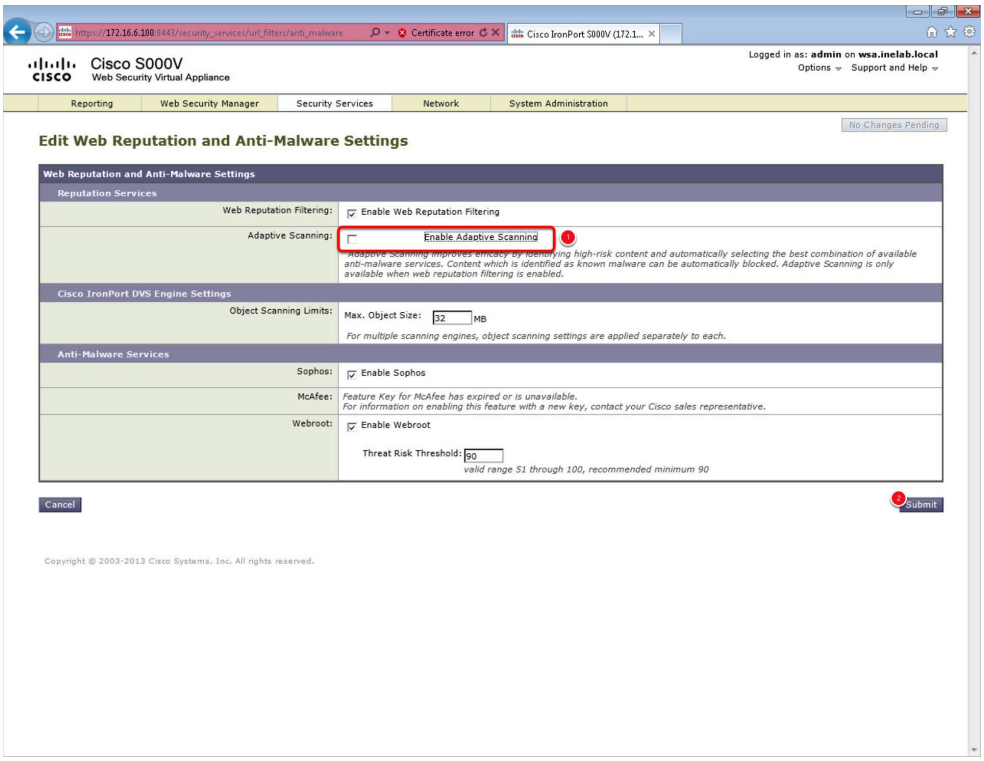


Click the **Edit Global Settings** button:

https://t.me/learningnets

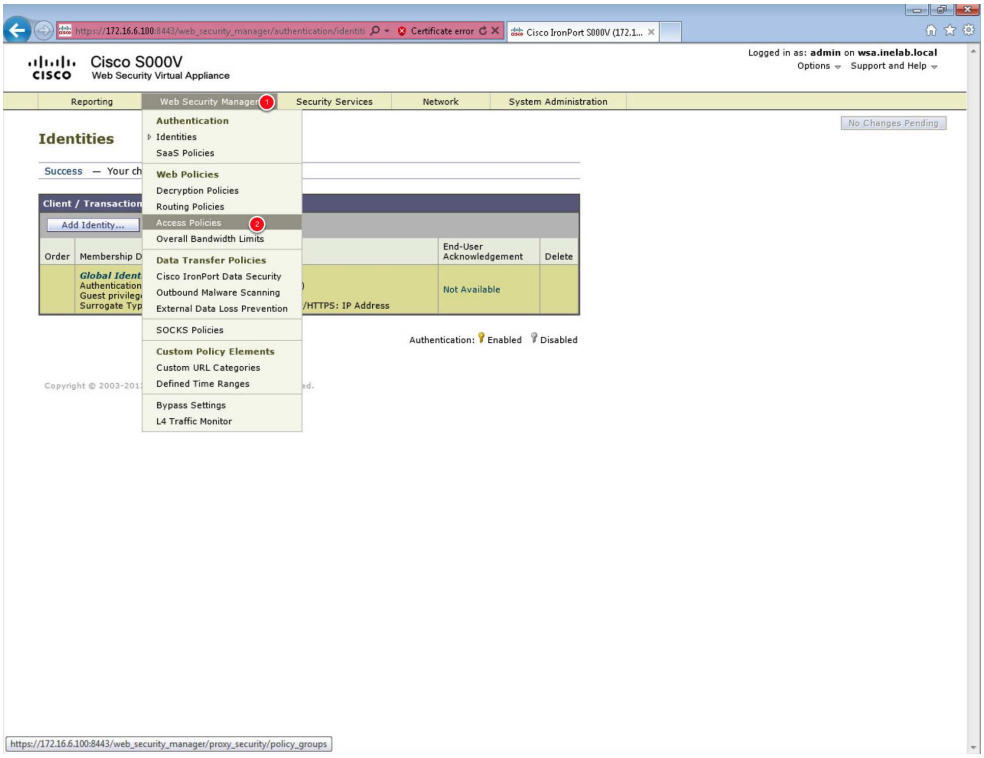


Disable adaptive scanning:

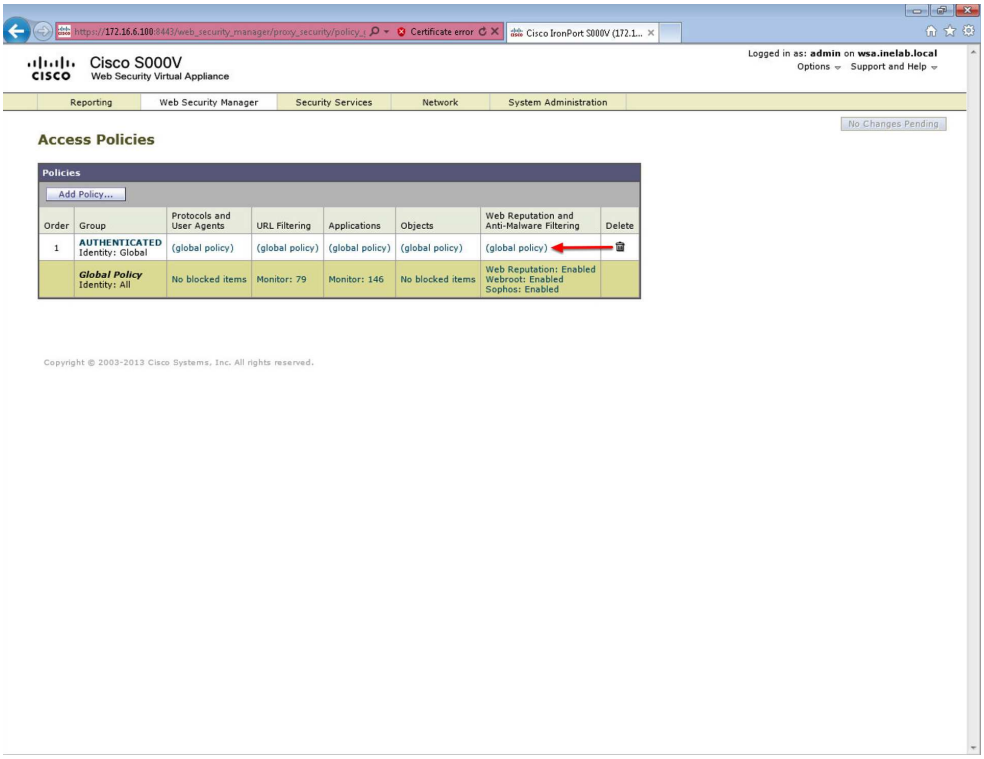


Navigate through the GUI to modify access policy settings for the authenticated users:

https://t.me/learningnets

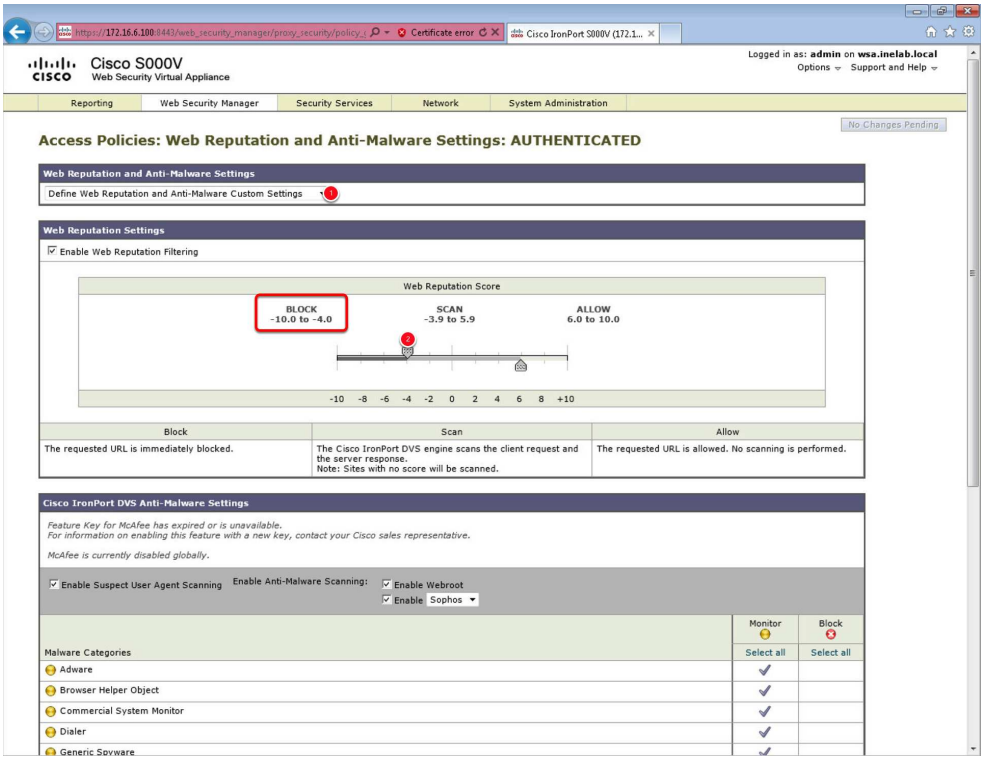


Click to edit the web reputation settings for the authenticated users access policy:

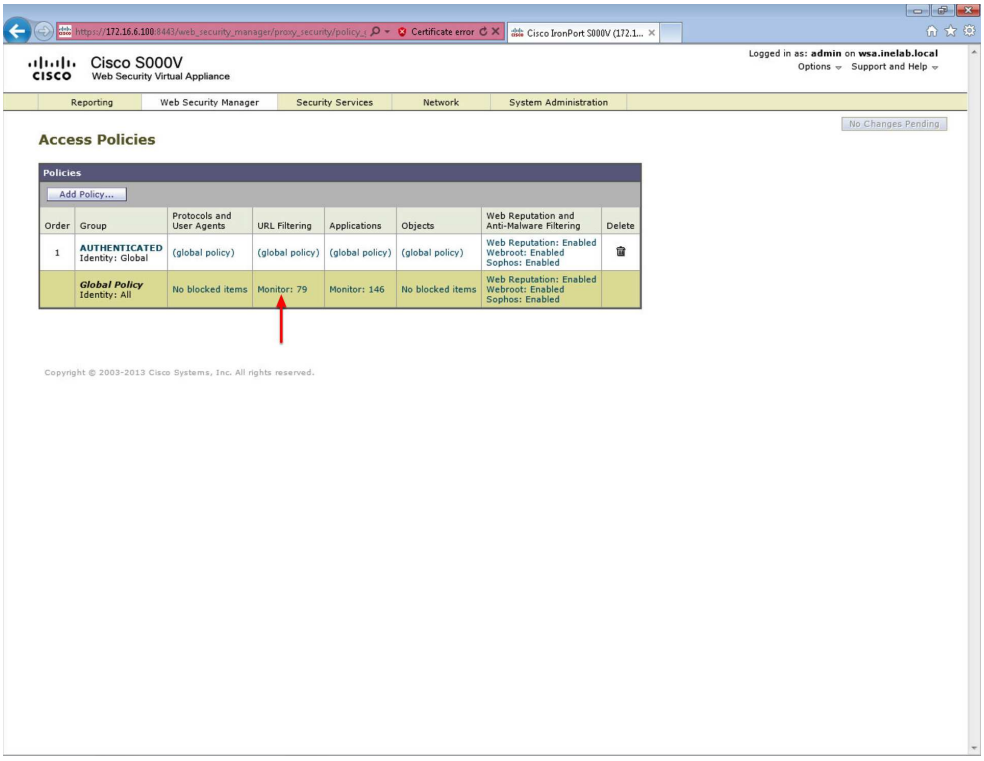


Configure WBRS so that access to a URL with a score of -4 is blocked:

https://t.me/learningnets

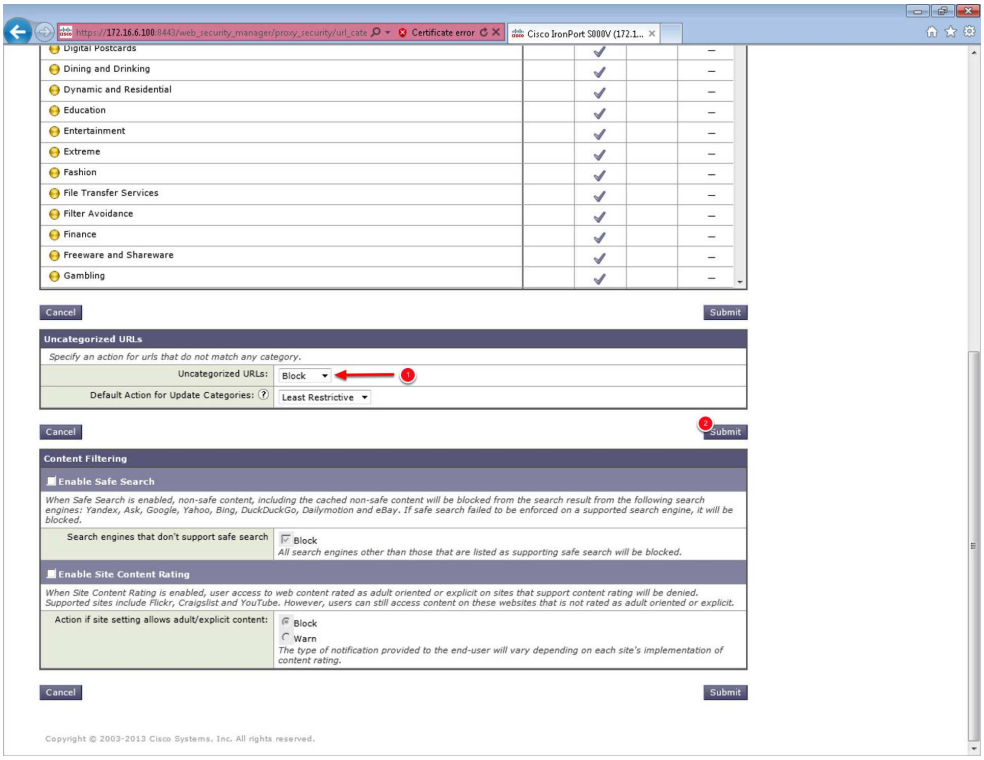


Navigate through the GUI to change the settings for uncategorized URLs, for the global access policy:

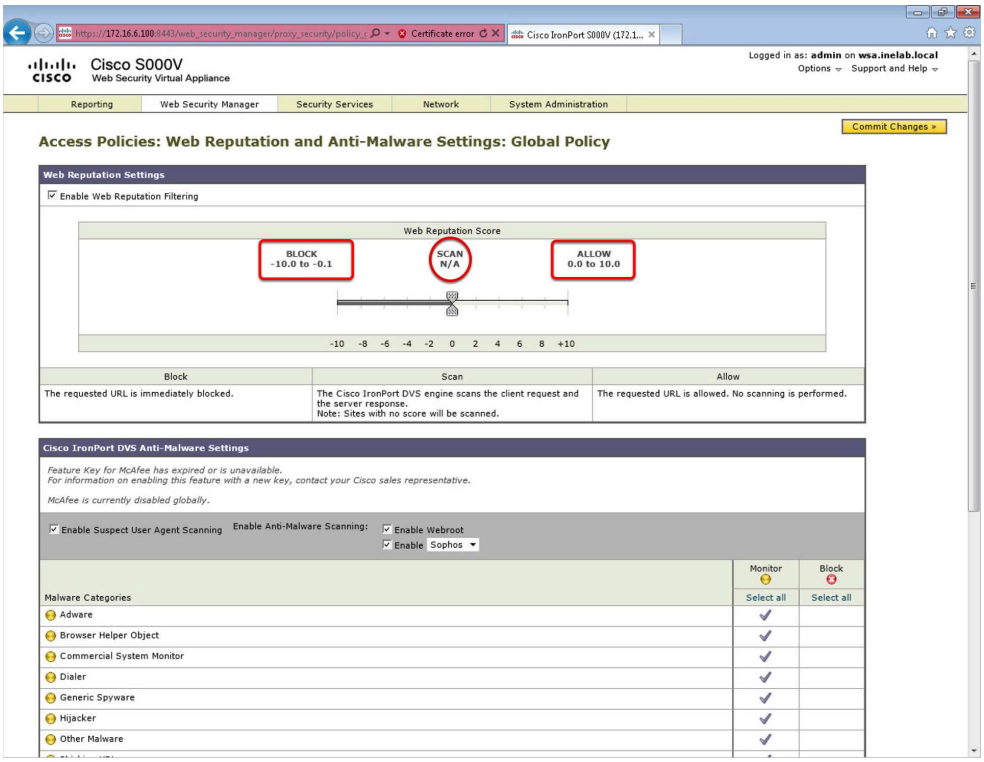


Change the action for uncategorized URLs to **block**:

https://t.me/learningnets



Change the WBRS actions for global access policy to match the task requirements; block all URLs with scores lower than zero, and allow URLs with a score higher than zero without any inspection:

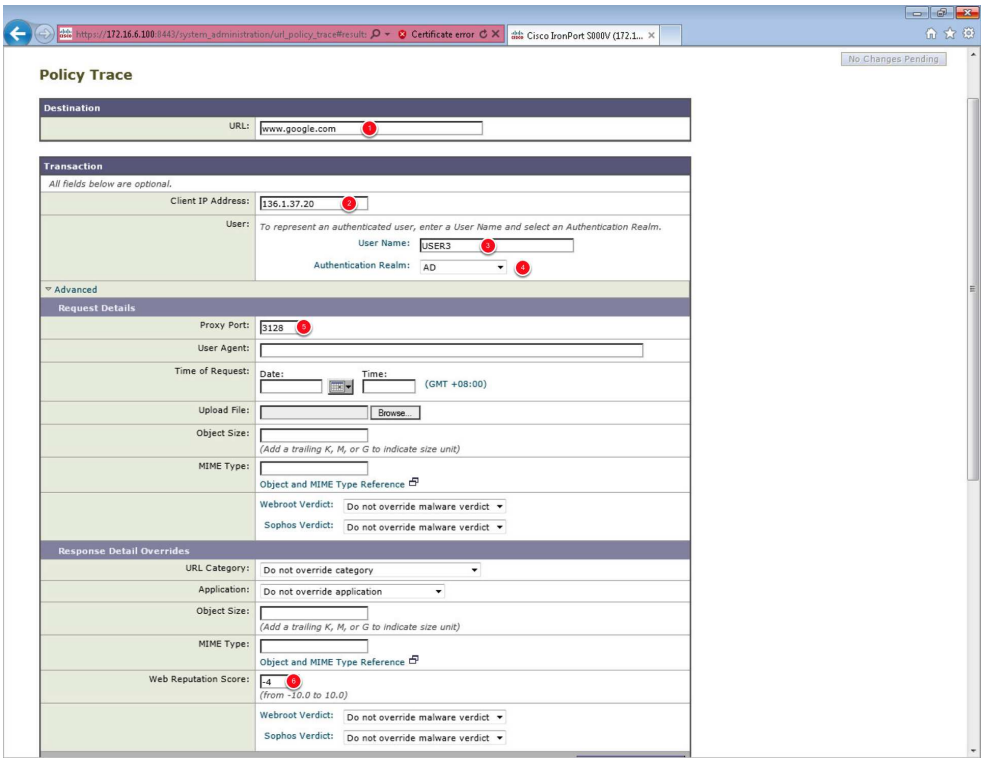


Verification

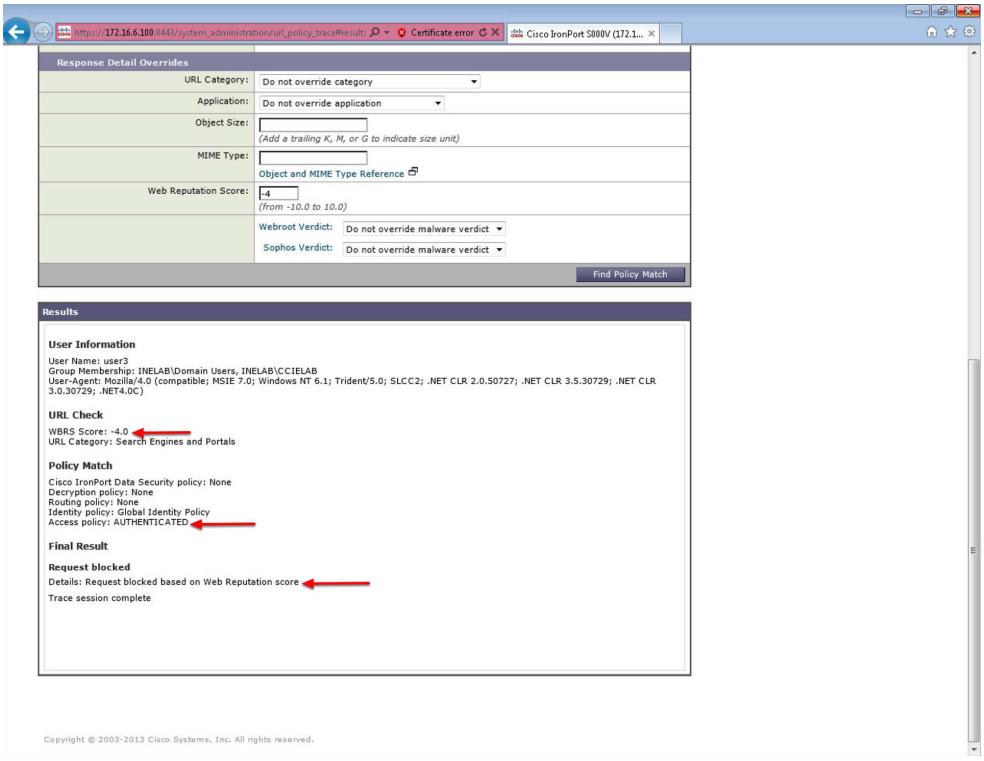
Using the Policy trace tool, simulate a request to **www.google.com**; override the

https://t.me/learningnets

WBRS to -4, and simulate authentication with USER3:

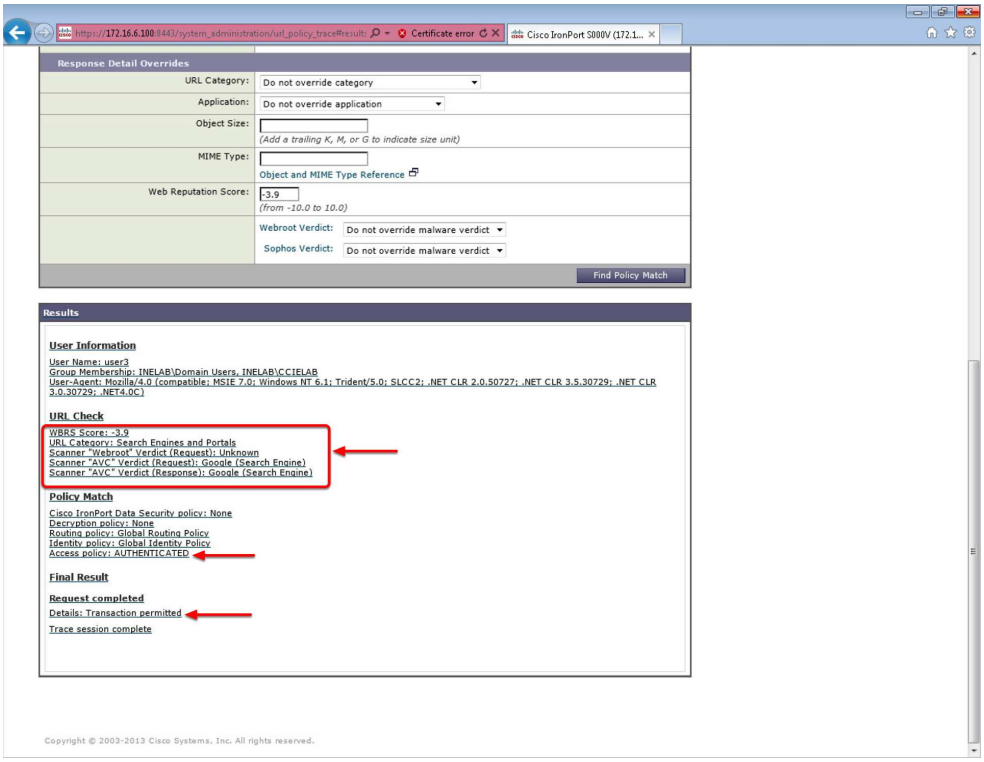


Based on the configured settings, the request should be denied because of the WBRS score; verify that correct access policy is matched:

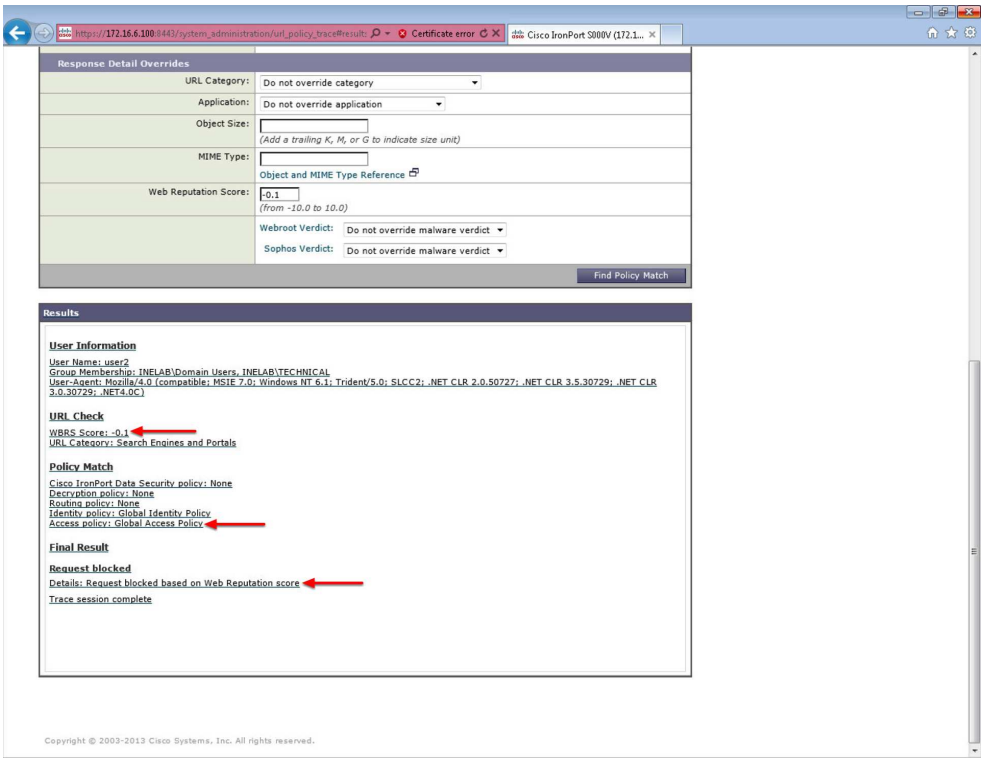


Simulate the request again, but this time override the WBRS to -3.9; access is allowed this time (the request is subject to inspections because of the WBRS score):

https://t.me/learningnets

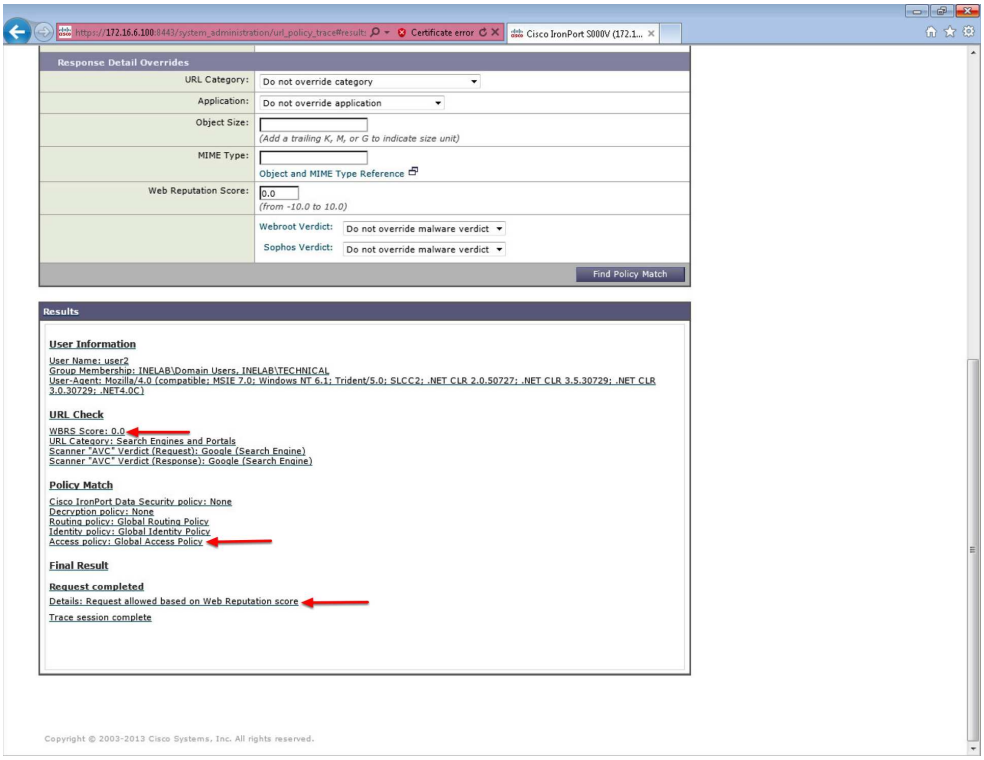


Using the Policy trace tool, simulate a request to **www.google.com**; override the WBRS to -0.1, and simulate authentication with USER2. The request should be denied because of the WBRS score; verify that correct access policy is matched:

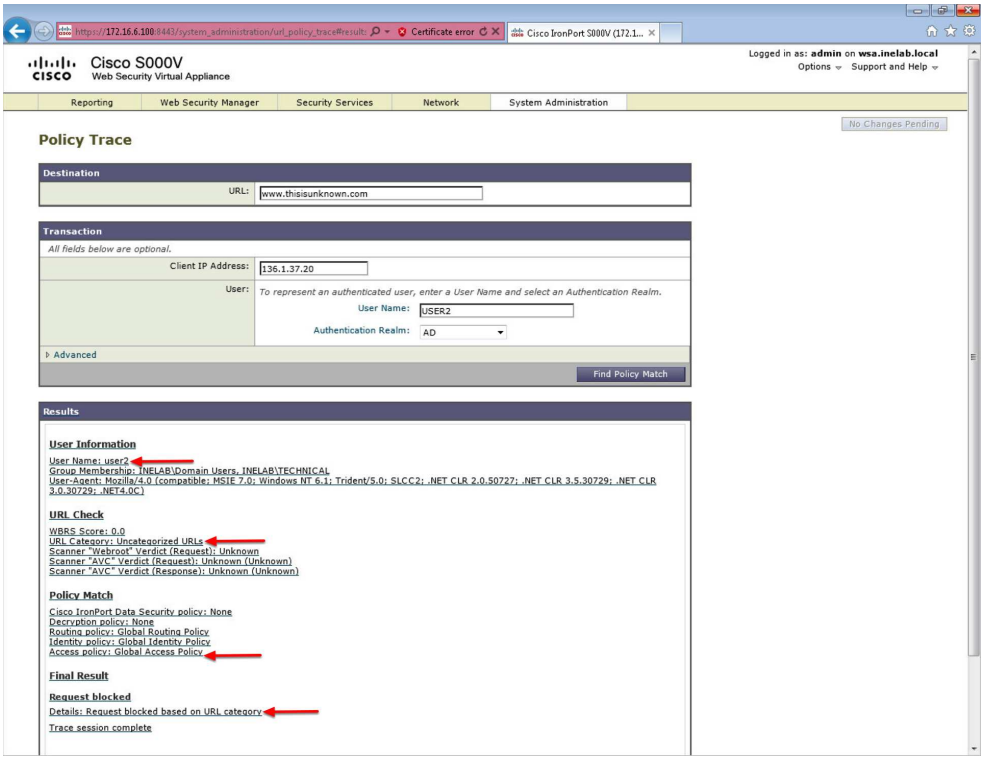


Simulate the request again, but this time override the WBRS to 0; access is allowed this time (the request is not subject to antimalware inspections because of the WBRS settings):

https://t.me/learningnets



Simulate a request to an uncategorized URL using USER2 to match the global access policy, and verify that the request is blocked:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Transparent Mode with Switch

Task

- Configure WSA and SW1 for WCCP layer 2 integration; redirect all HTTP/HTTPS traffic received inbound on VLAN 37, unless destined to R6's Loopback0 address.
- Ensure that HTTP/HTTPS traffic toward ASA3's VLAN 93 IP address is injected back in the network by the WSA without being inspected.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the WSA to the initial configuration (not to factory defaults), load it again by navigating to **System Administration > Configuration File > Load Configuration**, select the **WSA-DEFAULT.xml** file, click **Load**, and commit changes afterward.

Overview

When deployed in transparent mode, the WSA receives traffic not directly from the client, but redirected from network devices through the use of WCCP (Web Cache Communication Protocol). Currently, WCCPv2 is being used, which supports traffic for various TCP/UDP ports; the original WCCPv1 supported only HTTP (TCP port 80) traffic. WSA supports WCCP integration with layer 3 switches, routers, and ASA firewalls, but there are some differences between the implementations. In general,

https://t.me/learningnets

within the WCCP communication, there are two entities:

- **WCCP Server**, which is the network device performing traffic redirection (router, switch, firewall).
- **WCCP Client**, which is the device receiving redirected traffic and performing inspection (WSA).

The communication process is initiated by the **WCCP Client**. The two devices become paired after a **Here I am** and **I see you** packet exchange; the only condition is that both use the same WCCP service group ID, which acts like an identifier for the service (traffic being redirected). MD5 authentication is supported, and it's optional. Basically, when the **WCCP Client** registers with the **WCCP Server**, it tells the server which traffic (based on port numbers) to redirect for which service group ID. Communication can be unicast on UDP port 2048 by default, or multicast. WSA and ASA do not support multicast implementation. Based on the capabilities of both the **WCCP Server** and **WCCP Client**, there are two implementation scenarios:

- **Layer2**, where the **WCCP Server** and **WCCP Client** must be within the same subnet, traffic being redirected through destination MAC address change.
- **GRE**, where the **WCCP Server** and **WCCP Client** can reside in different subnets, traffic being injected inside a GRE tunnel between the WCCP Server and WCCP Client.

Switches used for the current CCIE Security blueprint supports only layer2 implementation. Note the following about the end-to-end traffic flow, considering the two WCCP entities (Client and Server) as well as client/server representing original source and destination for the traffic:

- Client/source (for example, TEST-PC-B) initiates an HTTP session toward a server/destination (for example, R4).
- WCCP Server (SW1) intercepts the traffic and redirects it to the WCCP Client (WSA).
- WCCP Client (WSA) initiates a HTTP session toward the destination (R4).
- WCCP Client (WSA) receives the requested data from the destination (R4).
- WCCP Client (WSA) sends the response straight back to the client (TEST-PC-B), bypassing the WCCP Server; the response is basically forwarded back based on the WSA routing table.

Based on the above, we can conclude that with WCCP, traffic flow is asymmetric; WCCP is therefore used only to redirect the initial request to the WCCP Client and not for return traffic. WCCP GRE tunnel could be used for return packets in isolated cases, such as when WSA is configured with a proxy bypass list, so it must allow the traffic without inspecting it—in which case it injects it back into the network via

https://t.me/learningnets

layer 2 or through GRE.

To configure the switch to redirect traffic toward the WSA, the following steps are required:

1. Define traffic to be redirected within an access-list (standard or extended, with extended preferred because the layer 4 port is available).
2. Define the service group ID (the same as defined on the WSA) and assign the previously configured redirect access-list; optionally, authentication can be configured as well.
3. Enable traffic interception for the configured service group ID inbound on the interface where traffic is received.

When grouping the service ID with the redirection ACL, a group-list ACL can be optionally specified that defines the WSAs allowed to register for this service ID. This is not a requirement, but rather a security enhancement for the WCCP configuration. Note that based on the IOS code running on the switches, deny statements may not be allowed/functional in the redirect ACL.

In previous scenarios, we did not enable HTTPS proxy, but HTTPS traffic was allowed by the proxy functioning in forward mode; this is because of the `CONNECT` method used by the client initiating the session when it has a proxy configured in the browser for HTTPS. Basically, this is an HTTP header command, similar to `GET`, used by the client for HTTP tunneling (such as FTP over HTTP) or HTTPS sessions; WSA is configured in the access policy to listen for some default ports for `CONNECT` actions, such as TCP 20,21,443, for example. In transparent mode, the client directly initiates the HTTPS session to the destination; there is no `CONNECT` method, so the WSA cannot serve HTTPS requests unless HTTPS proxy is activated, with or without HTTPS decryption.

The following actions can be performed by the WSA to an HTTP session:

- **Pass Through**: The session is allowed and WSA does not perform decryption.
- **Monitor**: The session is allowed and is passed to further inspection policies.
- **Decrypt**: WSA performs like a MiTM; it performs session decryption and establishes two HTTPS sessions, one with the destination/server and one with the source/client.
- **Drop**: The session is denied and the user is notified.

Note that when HTTPS proxy is activated, all HTTPS sessions are controlled by the decryption policies. But when a session is decrypted, it is being inspected by the access policies. The actual scope of the decryption policies, as implied by the name itself, is to determine how the HTTPS session will be processed, the available actions being specified earlier; there are multiple factors in this decision. The most

https://t.me/learningnets

important factors, in order of evaluation, are:

- Custom URL Category action
- Is WBRS action set to drop?
- Pre-defined URL Category action
- Application Visibility Control action (if HTTPS decryption is enabled for AVC)
- Is WBRS action set to decrypt or pass-through?
- Default action

Proxy bypass lists work only for transparent mode deployments, because the WSA needs to inject traffic back into the network without inspecting it, and this is not possible in forward mode. In forward mode, the proxy bypass list (also known as proxy exceptions) is configured at the client level in the browser, specifying for which destinations traffic will be sent directly and not to the proxy. For this task, there is no need to perform any configurations (identities, access policies, or decryption policies); the default configurations are enough because there are no specific requirements.

Configuration

```
SW1:


ip access-list extended REDIRECTED_TRAFFIC
 deny tcp  any host 150.1.6.6 eq 80
 deny tcp  any host 150.1.6.6 eq 443
 permit tcp any any eq 80
 permit tcp any any eq 443
!
ip wccp 92 redirect-list REDIRECTED_TRAFFIC
!
interface Vlan37
 ip wccp 92 redirect in
```

Because WCCP shares the same TCAM region as PBR (Policy Based Routing), with the default SDM, WCCP will not work; depending on the IOS Code, the switch may or may not log a warning message. Change the default SDM mode to **access**, **routing**, or **dual-ipv4-and-ipv6**, which gives some TCAM space to WCCP. A device reload is required for the SDM change to take effect and TCAM to be reinitialized.

```
SW1#show sdm prefer
 The current template is "desktop default" template.
 The selected template optimizes the resources in
```

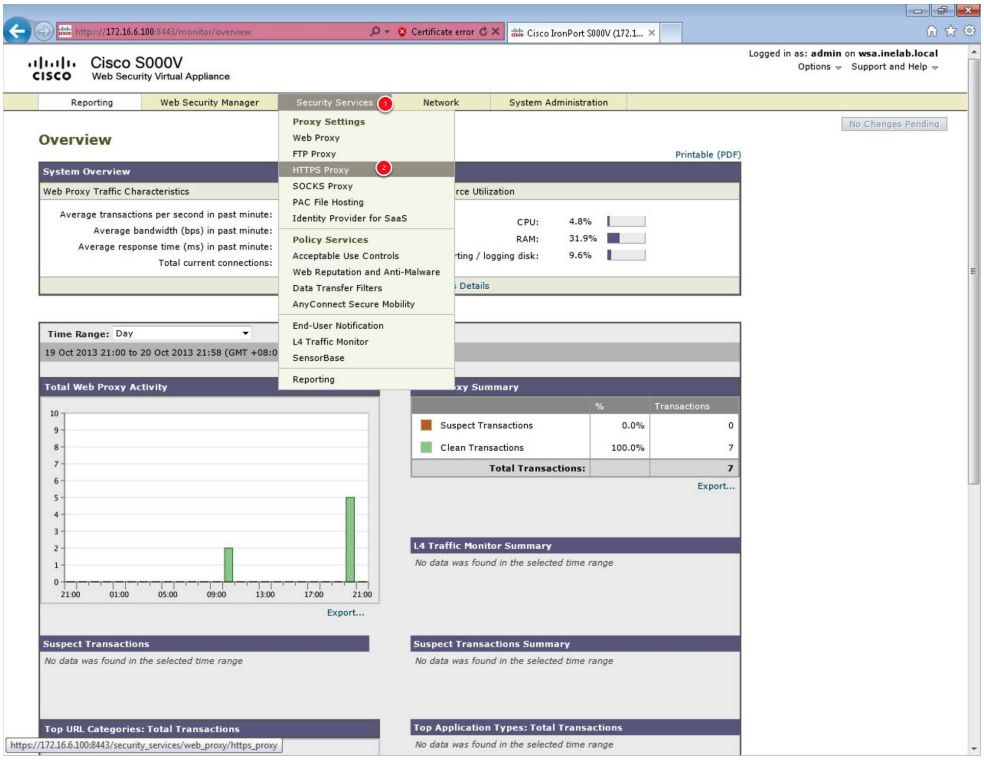
https://t.me/learningnets

```
 the switch to support this level of features for
 8 routed interfaces and 1024 VLANs.

   number of unicast mac addresses:                  6K
   number of IPv4 IGMP groups + multicast routes:    1K
   number of IPv4 unicast routes:                    8K
     number of directly-connected IPv4 hosts:        6K
     number of indirect IPv4 routes:                 2K
 number of IPv4 policy based routing aces:           0
   number of IPv4/MAC qos aces:                      0.5K
   number of IPv4/MAC security aces:                 1K
!
!SW1:

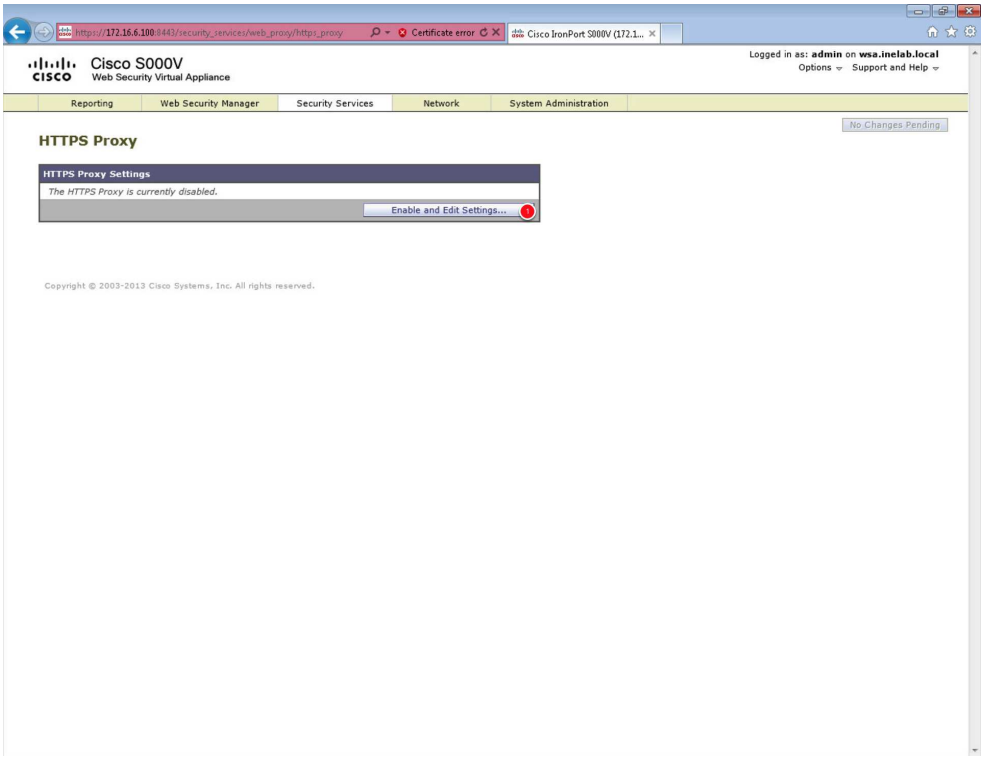
sdm prefer routing
```

Navigate through the GUI for HTTPS proxy activation:

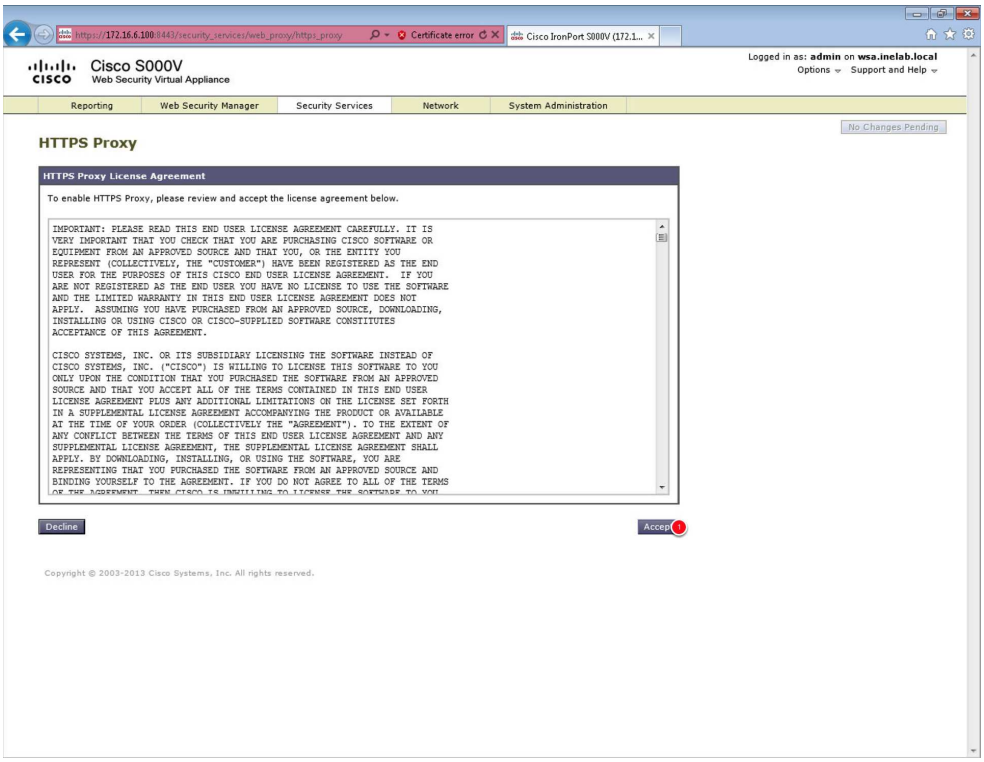


Click the **Enable and Edit Settings** button:

https://t.me/learningnets

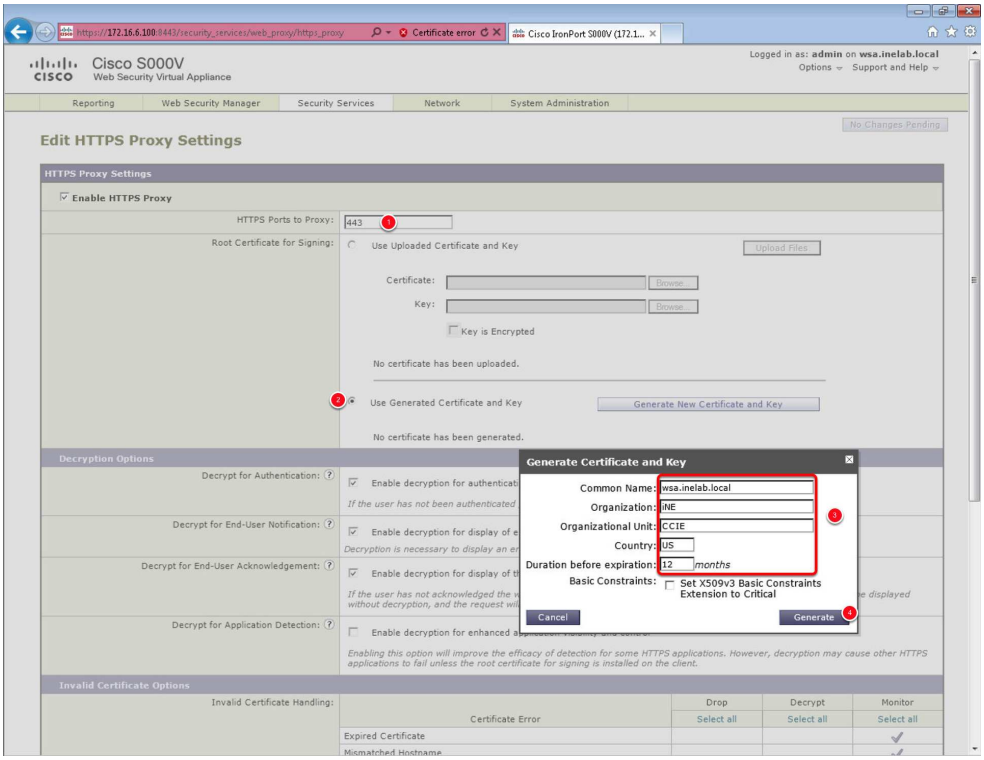


Accept the license agreement:

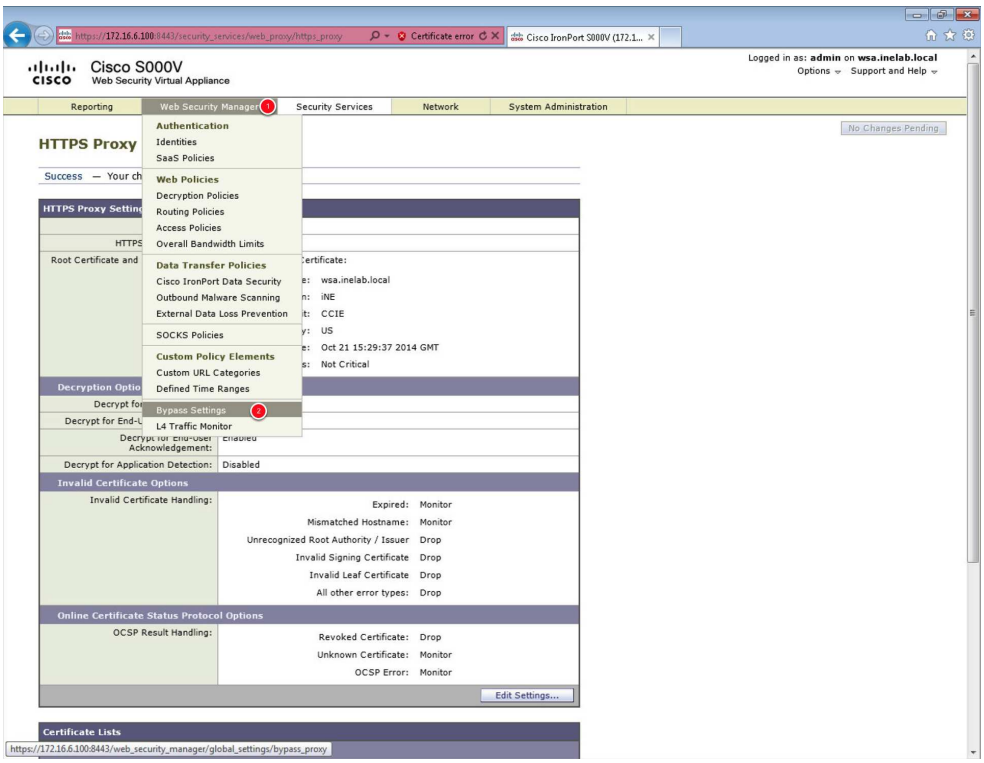


Specify the HTTPS proxy port (443) and generate a certificate to be used for the HTTPS decryption process:

https://t.me/learningnets

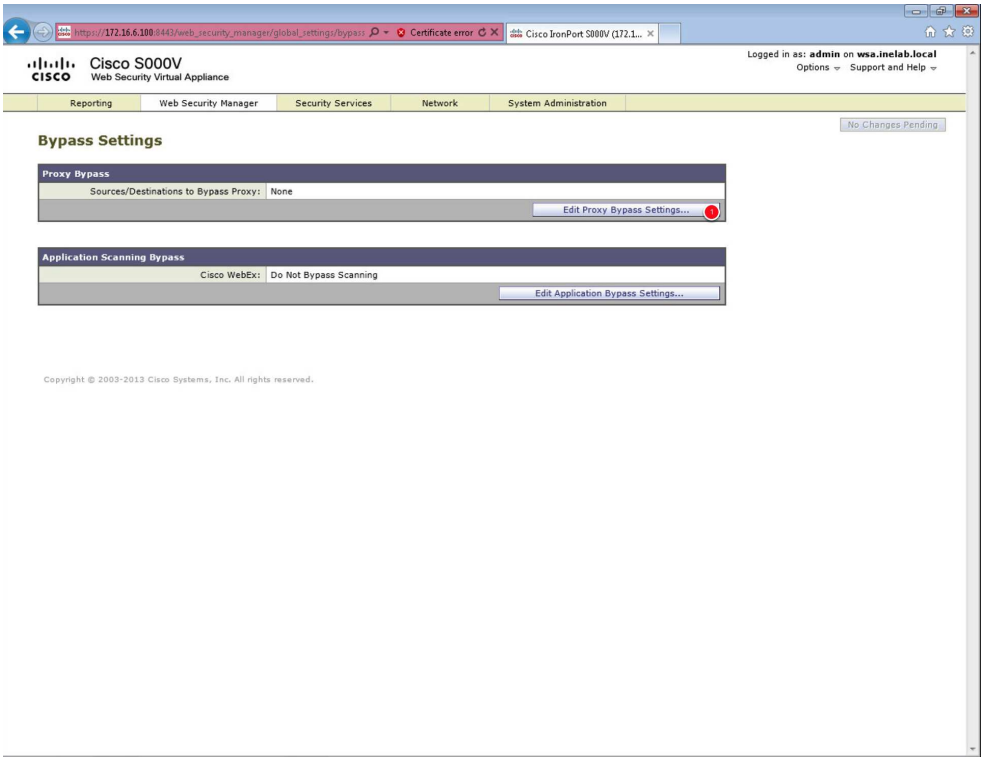


Wait for the certificate to be generated and click the **Submit** button. Navigate through the GUI to configure the proxy bypass list:

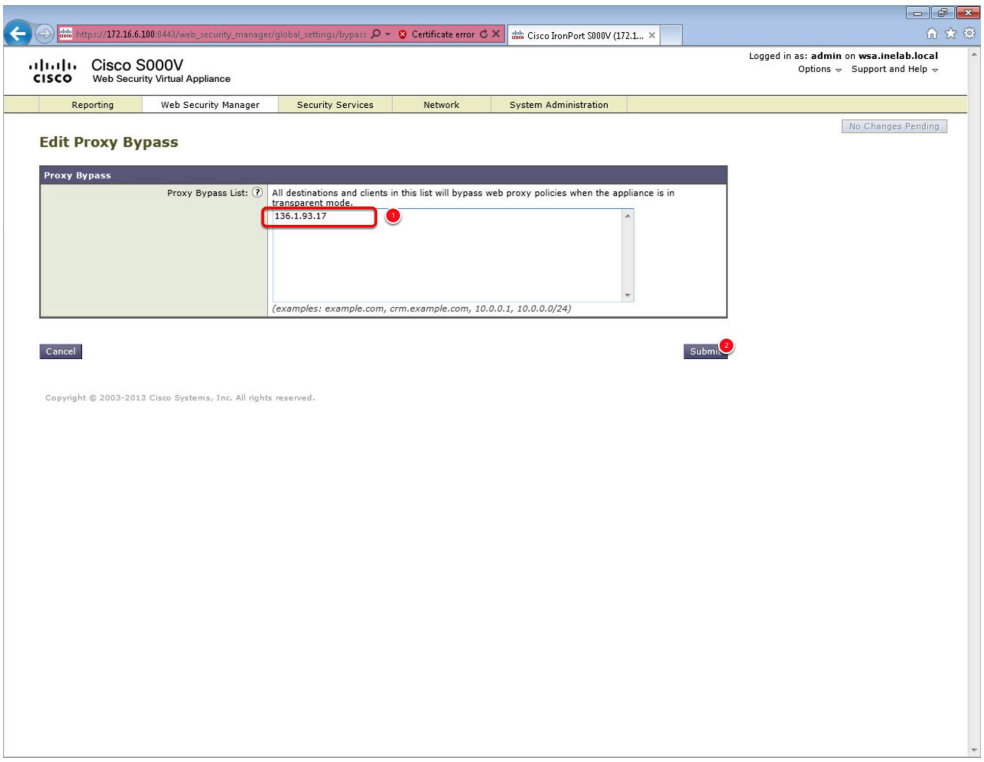


Click the **Edit Proxy Bypass Settings** button:

https://t.me/learningnets

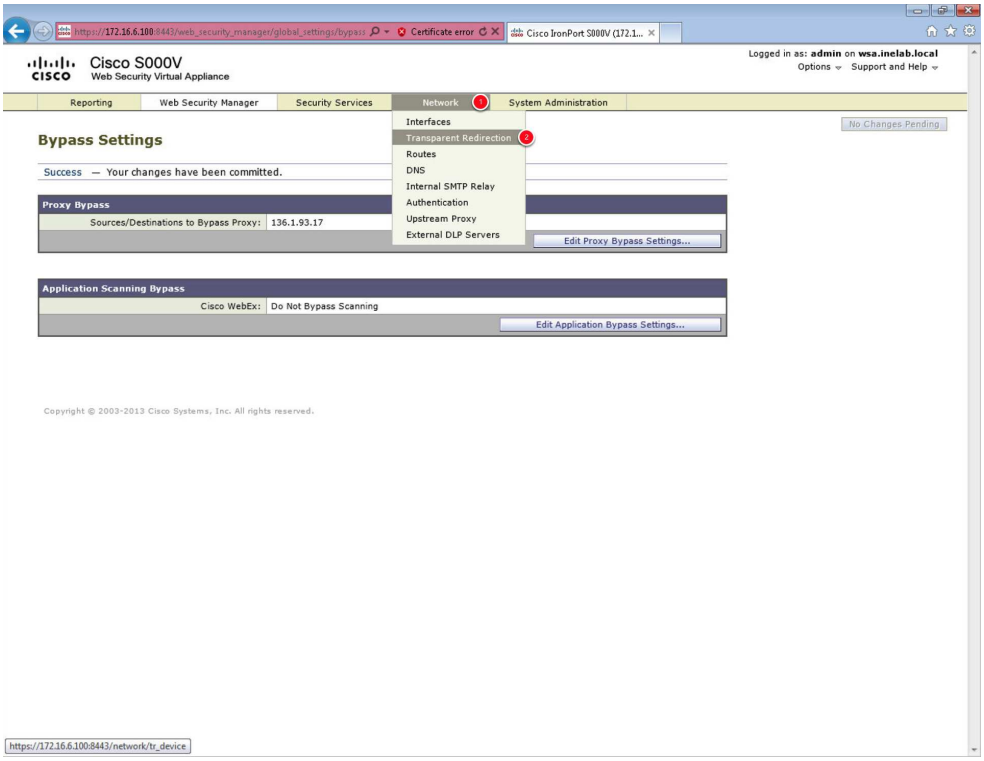


Add ASA3's VLAN 93 IP address to the list:

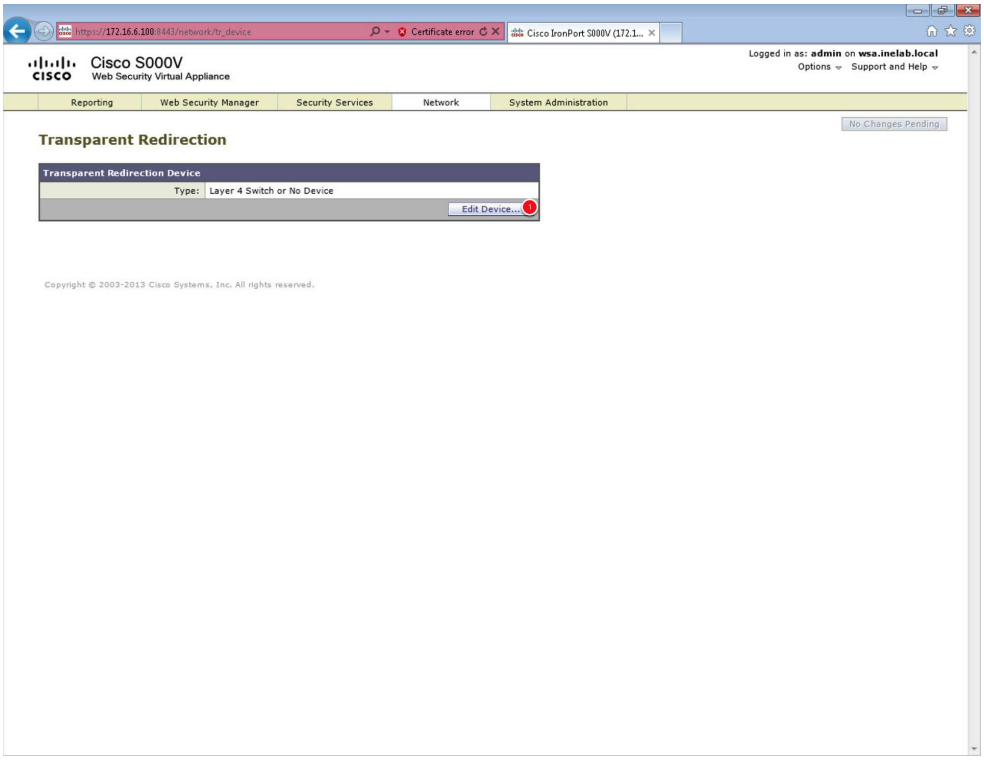


Browse through the GUI to configure the transparent mode functionality:

https://t.me/learningnets

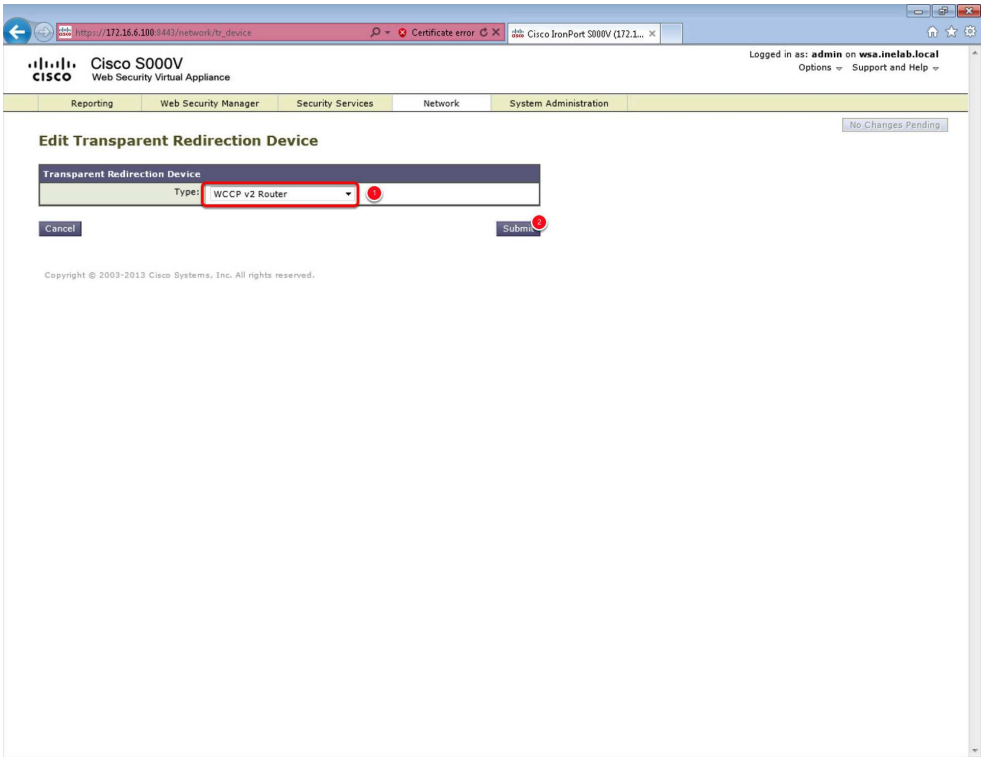


Click the **Edit Device** button:

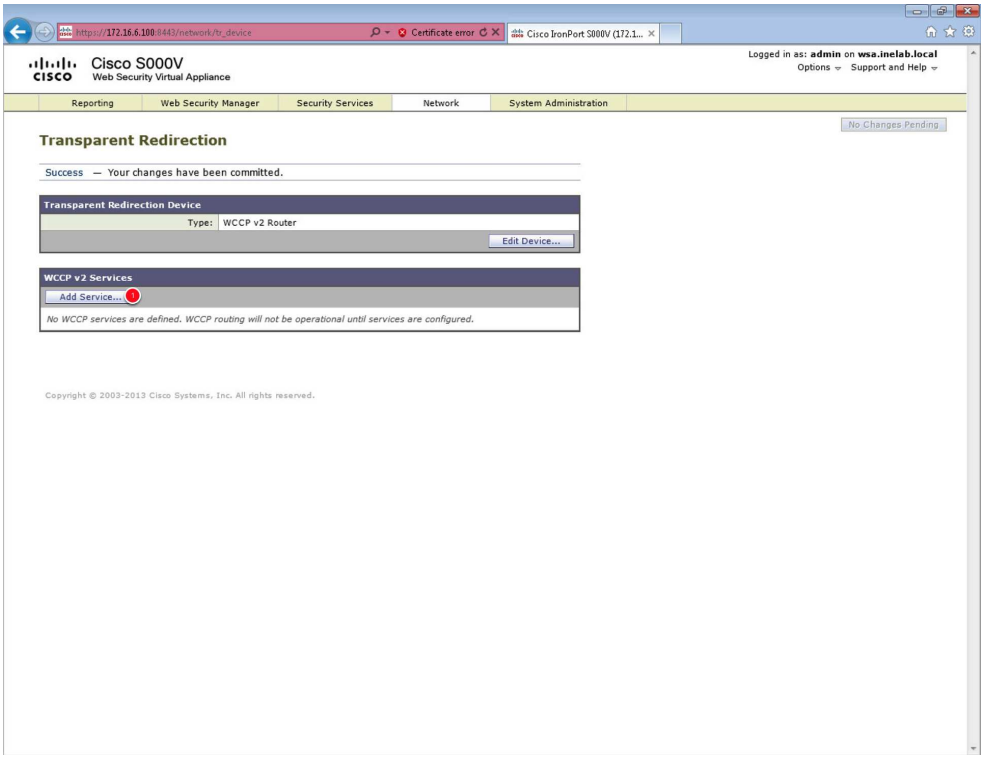


Select the **WCCP v2 Router** option:

https://t.me/learningnets

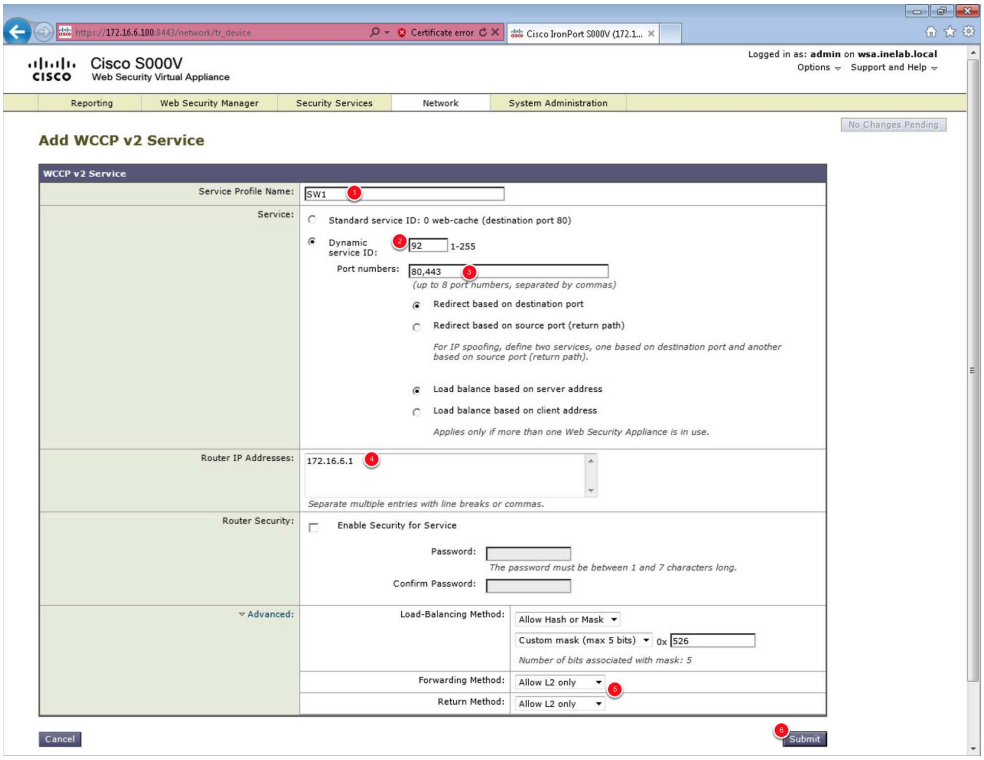


Click the **Add Service** button:



Configure the service group ID (the same as on SW1) and the ports to be used for traffic redirection (80 and 443), and allow layer 2 mode only:

https://t.me/learningnets



Verification

Verify that TCAM memory space is allocated for WCCP/PBR after reload:

```
SW1#show sdm  prefer

 The current template is "desktop routing" template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1024 VLANs.

  number of unicast mac addresses:                  3K
  number of IPv4 IGMP groups + multicast routes:    1K
  number of IPv4 unicast routes:                    11K
    number of directly-connected IPv4 hosts:        3K
    number of indirect IPv4 routes:                 8K
number of IPv4 policy based routing aces:          0.5K

  number of IPv4/MAC qos aces:                      0.5K
  number of IPv4/MAC security aces:                 1K
```

Verify the WCCP method being used, and whether service is active:

```
SW1#show ip wccp 92  detail | exclude 00
WCCP Client information:
        Protocol Version:        2.0
```

https://t.me/learningnets

```
State:                    Usable
 Redirection:              L2
Packet Return:             L2

        Packets Redirected:    0
        Assignment:            MASK

        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
```

Verify WCCP statistics before initiating any traffic; note that there are zero packets being denied from redirection:

```
SW1#show ip wccp
Global WCCP information:
    Router information:
        Router Identifier:                  150.1.9.9 Protocol Version:                    2.0
 Service Identifier: 92
Number of Service Group Clients:     1
Number of Service Group Routers:     1
        Total Packets s/w Redirected:       0
          Process:                          0
          CEF:                              0 Redirect access-list:                REDIRECTED_TRAFFIC
Total Packets Denied Redirect:       0

        Total Packets Unassigned:           0
        Group access-list:                  -none-
        Total Messages Denied to Group:     0
        Total Authentication failures:      0
        Total Bypassed Packets Received:    0
```

Verify the WSA proxy logs (number 21) for any WCCP error messages; there should be none:

```
Press Ctrl-C to stop.
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:92:[secret]
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:92: L2
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:92:ret L2

Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:92:TCP
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:92: 172.16.6.1
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:SG:92: Security enable <- 0
```

https://t.me/learningnets

```
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:SG:92: Hash enable <- 1
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:SG:92: Mask enable <- 1
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:SG:92: Service direction <- 0
Mon Oct 21 23:51:38 2013 Info: WCCP : - : CONFIG:SG:92: Hash/mask on client <- 0
```

Initiate from **TEST-PC-B** both HTTP and HTTPS sessions toward R6 Loopback0 address; these should not be redirected to WSA. Check the statistics again:

```
SW1#show ip wccp
Global WCCP information:
    Router information:
        Router Identifier:                  150.1.9.9
        Protocol Version:                   2.0

    Service Identifier: 92
        Number of Service Group Clients:    1
        Number of Service Group Routers:    1
        Total Packets s/w Redirected:       0
          Process:                          0
          CEF:                              0
        Redirect access-list:               REDIRECTED_TRAFFIC Total Packets Denied Redirect:       23

        Total Packets Unassigned:           0
        Group access-list:                  -none-
        Total Messages Denied to Group:     0
        Total Authentication failures:      0
        Total Bypassed Packets Received:    0
```

Initiate from **TEST-PC-B** a HTTPS session toward ASA3's VLAN 93 IP address; it should be functional and verify the WSA proxy bypass logs. The following messages should appear:

```
wsa.inelab.local> tail

Currently configured logs:
1. "accesslogs" Type: "Access Logs" Retrieval: FTP Poll
2. "authlogs" Type: "Authentication Framework Logs" Retrieval: FTP Poll
3. "avc_logs" Type: "AVC Engine Logs" Retrieval: FTP Poll
4. "bypasslogs" Type: "Proxy Bypass Logs" Retrieval: FTP Poll
5. "cli_logs" Type: "CLI Audit Logs" Retrieval: FTP Poll
6. "configdefragd_logs" Type: "Configuration Logs" Retrieval: FTP Poll
7. "dca_logs" Type: "DCA Engine Logs" Retrieval: FTP Poll
8. "external_auth_logs" Type: "External Authentication Logs" Retrieval: FTP
Poll
9. "feedback_logs" Type: "Feedback Logs" Retrieval: FTP Poll
```

https://t.me/learningnets

```
10. "fips_logs" Type: "FIPS Logs" Retrieval: FTP Poll
11. "ftpd_logs" Type: "FTP Server Logs" Retrieval: FTP Poll
12. "gui_logs" Type: "GUI Logs" Retrieval: FTP Poll
13. "haystackd_logs" Type: "Haystack Logs" Retrieval: FTP Poll
14. "httpslog" Type: "HTTPS Logs" Retrieval: FTP Poll
15. "idsdataloss_logs" Type: "Data Security Logs" Retrieval: FTP Poll
16. "logderrorlogs" Type: "Logging Logs" Retrieval: FTP Poll
17. "mcafee_logs" Type: "McAfee Logs" Retrieval: FTP Poll
18. "musd_logs" Type: "AnyConnect Secure Mobility Daemon Logs" Retrieval: FTP
Poll
19. "ocspd_logs" Type: "OCSP Logs" Retrieval: FTP Poll
20. "pacd_logs" Type: "PAC File Hosting Daemon Logs" Retrieval: FTP Poll
21. "proxylogs" Type: "Default Proxy Logs" Retrieval: FTP Poll
22. "reportd_logs" Type: "Reporting Logs" Retrieval: FTP Poll
23. "reportqueryd_logs" Type: "Reporting Query Logs" Retrieval: FTP Poll
24. "saas_auth_log" Type: "SaaS Auth Logs" Retrieval: FTP Poll
25. "shd_logs" Type: "SHD Logs" Retrieval: FTP Poll
26. "snmp_logs" Type: "SNMP Logs" Retrieval: FTP Poll
27. "sntpd_logs" Type: "NTP Logs" Retrieval: FTP Poll
28. "sophos_logs" Type: "Sophos Logs" Retrieval: FTP Poll
29. "status" Type: "Status Logs" Retrieval: FTP Poll
30. "system_logs" Type: "System Logs" Retrieval: FTP Poll
31. "trafmon_errlogs" Type: "Traffic Monitor Error Logs" Retrieval: FTP Poll
32. "trafmonlogs" Type: "Traffic Monitor Logs" Retrieval: FTP Poll
33. "uds_logs" Type: "UDS Logs" Retrieval: FTP Poll
34. "updater_logs" Type: "Updater Logs" Retrieval: FTP Poll
35. "wbnp_logs" Type: "WBNP Logs" Retrieval: FTP Poll
36. "webcat_logs" Type: "Web Categorization Logs" Retrieval: FTP Poll
37. "webrootlogs" Type: "Webroot Logs" Retrieval: FTP Poll
38. "welcomeack_logs" Type: "Welcome Page Acknowledgement Logs" Retrieval:
FTP Poll
Enter the number of the log you wish to tail.[]> 4
Press Ctrl-C to stop.
Mon Oct 21 23:55:16 2013 Info: Connection from 136.1.37.51:11471 to 136.1.93.17:443 bypassed
```

Initiate from **TEST-PC-B** an HTTP session toward R4's Loopback0 address; it should be functional and verify the WSA access logs. The following message should appear:

https://t.me/learningnets

```
1382371141.403 29 136.1.37.51
TCP_MISS/401 253 GET http://150.1.4.4/ - DIRECT/150.1.4.4 text/plain DEFAULT_CASE_12-DefaultGroup-DefaultGroup-NONE-
 <nc,-3.5,0,"-",0,0,0,1,"-",-,-,-,"-",0,262686,"-","-",-,-,nc,nc,"Unknown","-","Unknown","Unknown","-","-",69.79,0,-
```

Initiate from **TEST-PC-B** an HTTPS session toward R4's Loopback0 address; it should be functional and verify the WSA access logs because the session was decrypted (because of an untrusted issuer for R4's certificate, which is self-signed, the session will match on the default HTTPS action, which is set to decrypt):

```
1382371369.151 2218 136.1.37.51
TCP_MISS_SSL/200 0 TCP_CONNECT 150.1.4.4:443 - DIRECT/150.1.4.4 - DECRYPT_ADMIN-UNRECOGNIZED_ROOT_CERT_7-DefaultGrou
 <nc,ns,-,"-",-,-,-,-,"-",-,-,-,"-",-,-,"-","-",-,-,nc,-,"-","-","Unknown","Unknown","-","-",0.00,0,-,"-","-"> -
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Transparent Mode with Router

Task

- Configure WSA and R4 for WCCP GRE integration, and redirect all outbound HTTP traffic on VLAN 32.
- If the **WCCP Client** is not registered, traffic should be discarded.
- Ensure that R4 allows only **WCCP Client** from IP address 172.16.6.100 to register for the service and enforce authentication.
- WSA should enforce user authentication through the LDAP server at 172.16.20.100 (TEST-SRV-B) based on the `sAMAccountName` attribute.
  - Use **ldapuser/1qaz@WSX** as the username/password pair for LDAP queries; **ldapuser** is a member of the **Users** Organization Unit.
  - Use **USER1/1qaz@WSX** as the username/password pair to simulate user authentication.
- Configure WSA to block GOOGLE mail service file uploads/downloads and completely restrict access to Yahoo mail service.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the WSA to the initial configuration (not to factory defaults), load it again by navigating to **System Administration > Configuration File > Load Configuration**, select the **WSA-DEFAULT.xml**

https://t.me/learningnets

file, click **Load**, and commit changes afterward.

Overview

ISR G1/G2 routers support both GRE and layer2 WCCP deployments. We will use the GRE method because R4 and WSA are not sharing the same subnet. By default, WCCP functions in **open** mode, so if no **WCCP Client** registers with the **WCCP Server**, traffic will not be redirected and will be allowed to pass. WCCP can be configured to function in **closed** mode; traffic matching the redirection access-list will be denied if no **WCCP Client** is registered. Compared with the previous task, with switch integration, the ISR router supports both inbound and outbound traffic redirection. From a configuration perspective, the steps are the same as for the switch configuration.

WSA supports the AVC (Application Visibility and Control) service for both HTTP and HTTPS sessions. It has a set of signatures with which it can recognize web 2.0 applications using HTTP as the transport protocol, or applications using HTTP tunneling. For each supported application there are several actions available; with Facebook, for example, it can allow login and chat services but restrict playing Facebook games and uploading. It is similar to how IOS CBAC and ZBPF support application inspection, and it relies on the NBAR2 engine for application discovery.

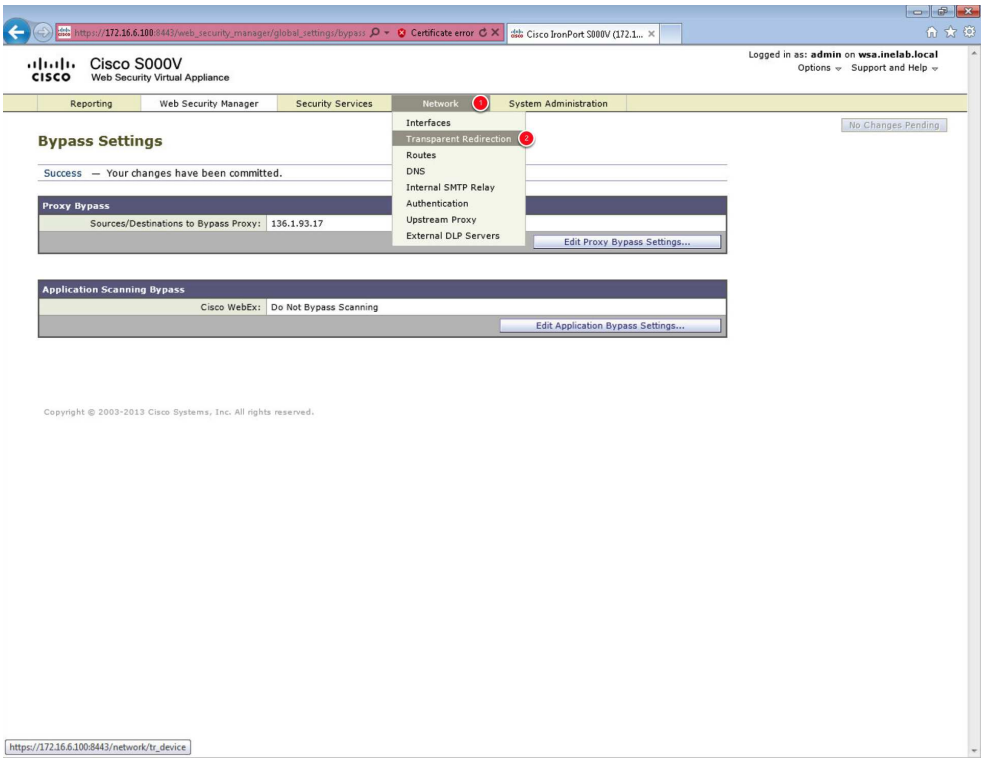
Configuration

R4:

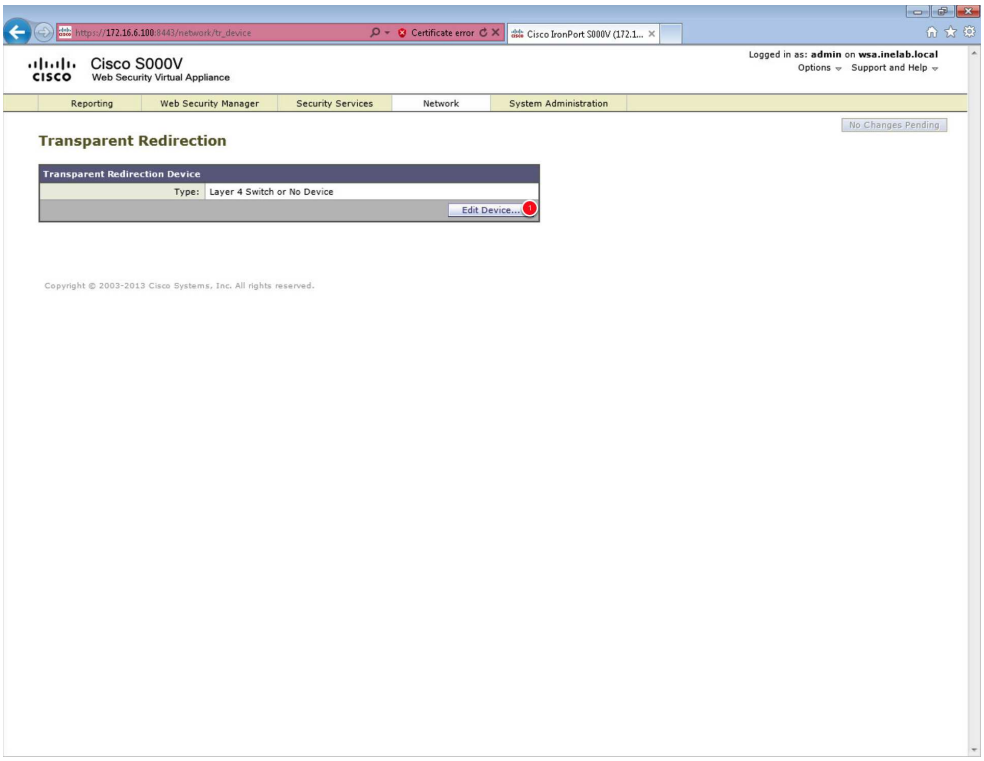
```
ip access-list extended REDIRECTED_TRAFFIC
 permit tcp any any eq 80
!
ip access-list standard WSA_GROUP
 permit host 172.16.6.100
!
ip access-list extended WCCP_CLOSED
 permit tcp any any eq 80
!
ip wccp version 2
ip wccp 93 redirect-list REDIRECTED_TRAFFIC group-list WSA_GROUP password cisco mode closed service-list WCCP_CLOSED
!
interface FastEthernet0/1.32
 ip wccp 93 redirect out
```

https://t.me/learningnets

Browse through the GUI to configure the transparent mode functionality:

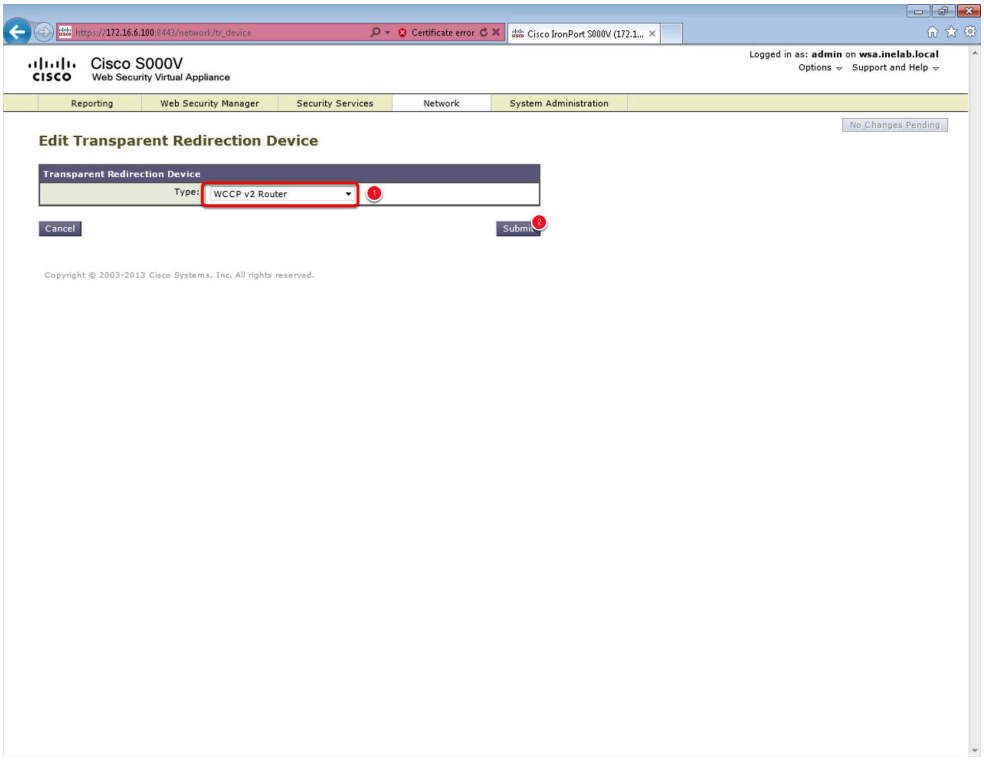


Click the **Edit Device** button:

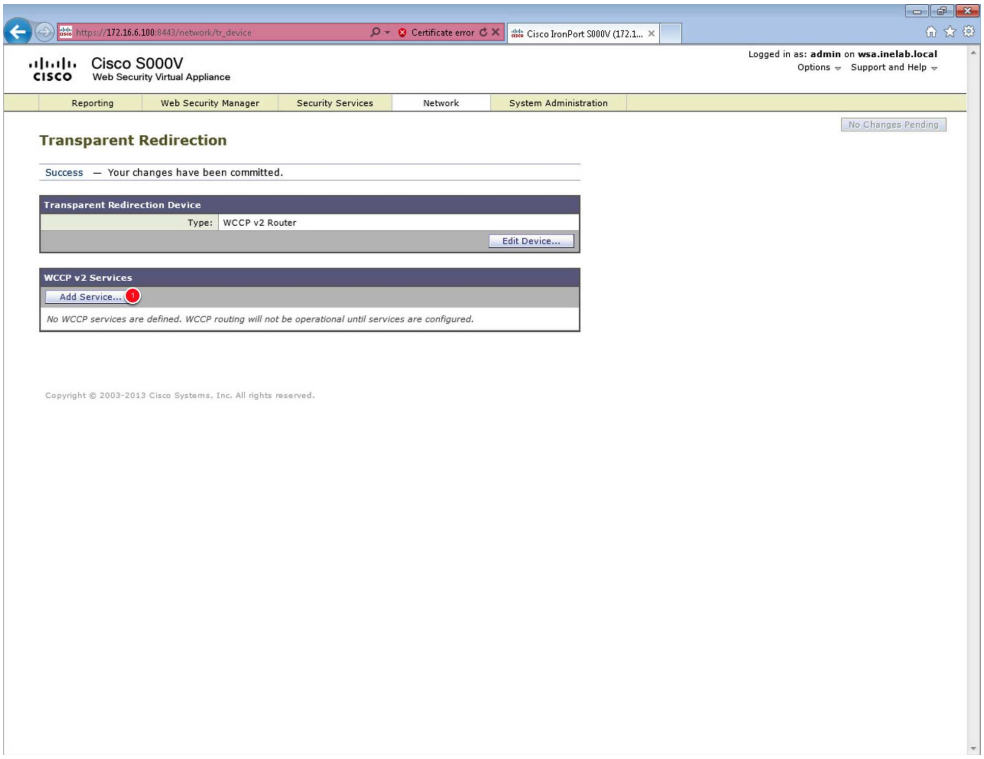


Select the **WCCP v2 Router** option:

https://t.me/learningnets

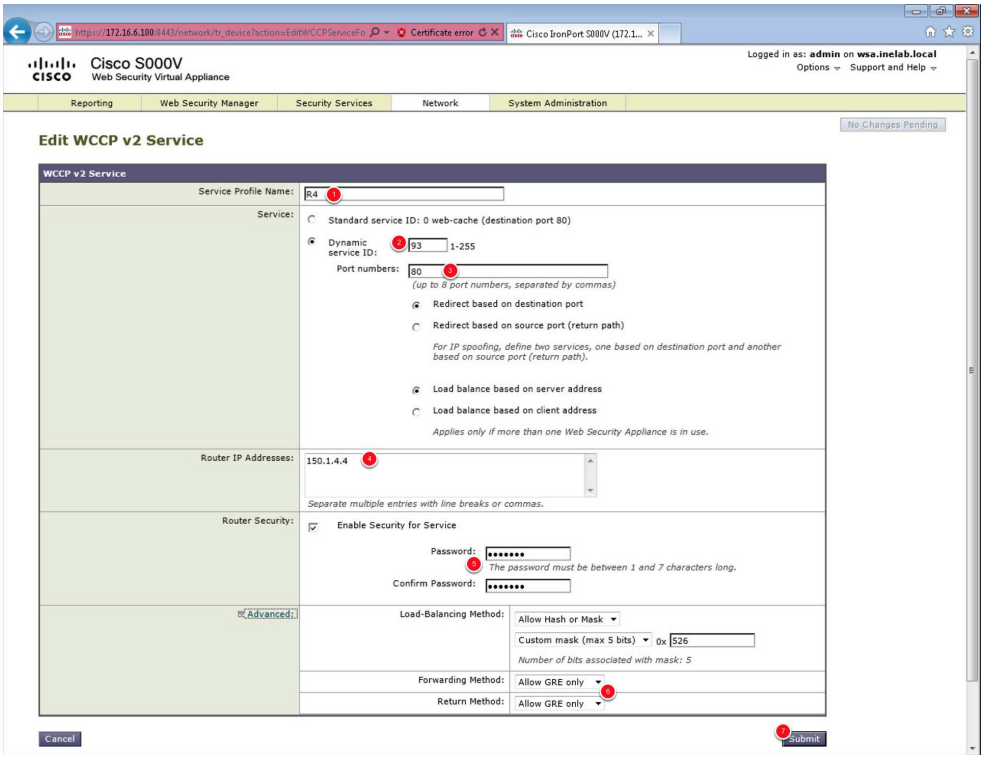


Click the **Add Service** button:

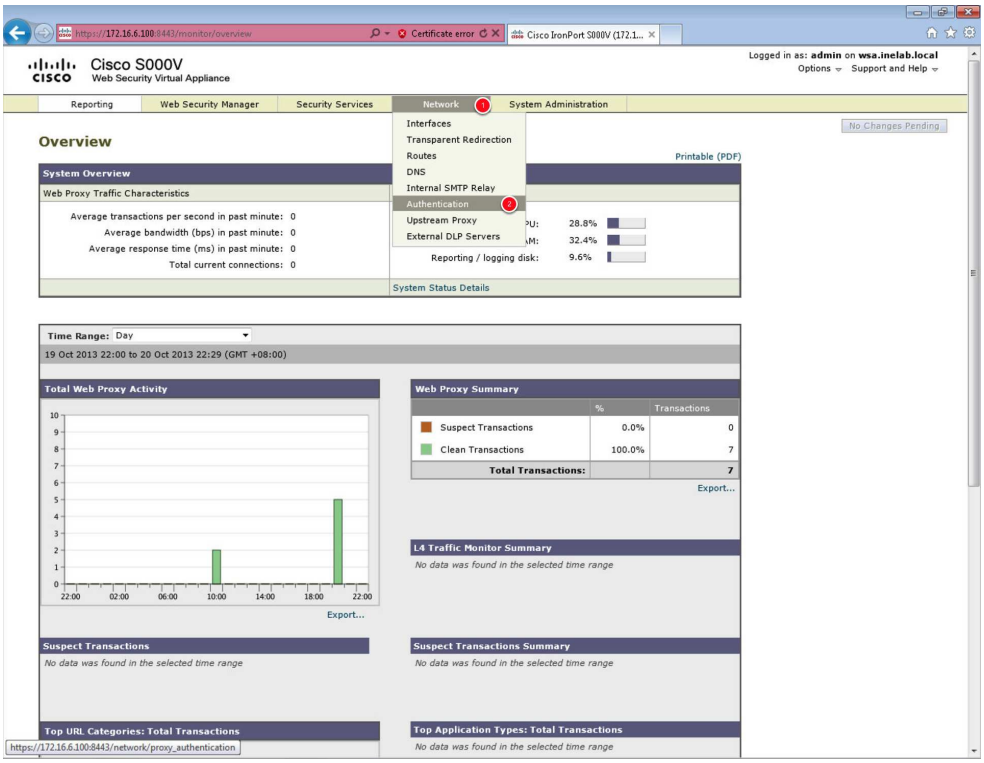


Configure the service group ID (the same as on R4), the ports to be used for traffic redirection (80), and the authentication password, and allow GRE mode only; otherwise, GRE will not be functional because WSA will always try to negotiate layer2:

https://t.me/learningnets

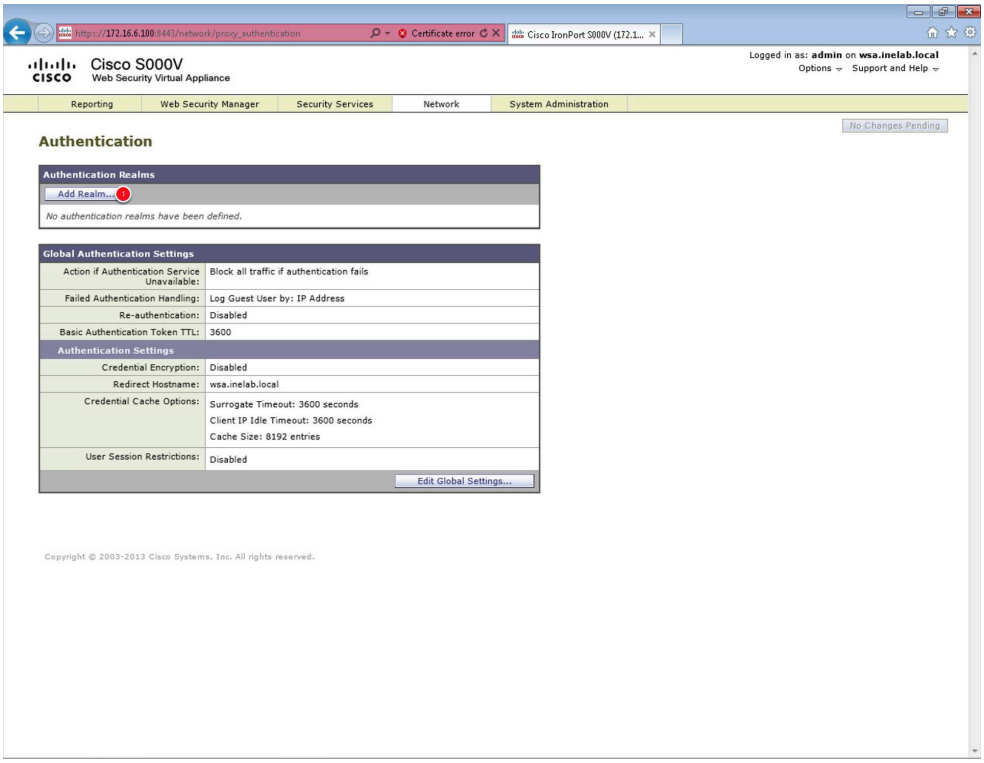


Navigate through the GUI to configure integration with an LDAP/AD server for user authentication:

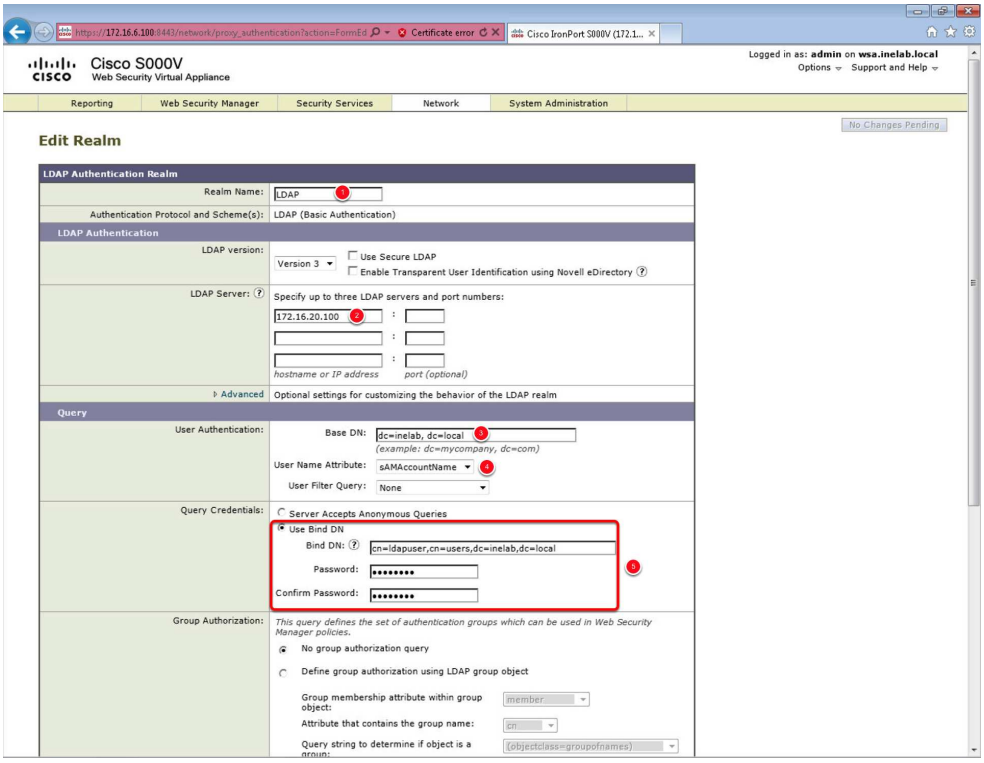


Click the **Add Realm** button:

https://t.me/learningnets

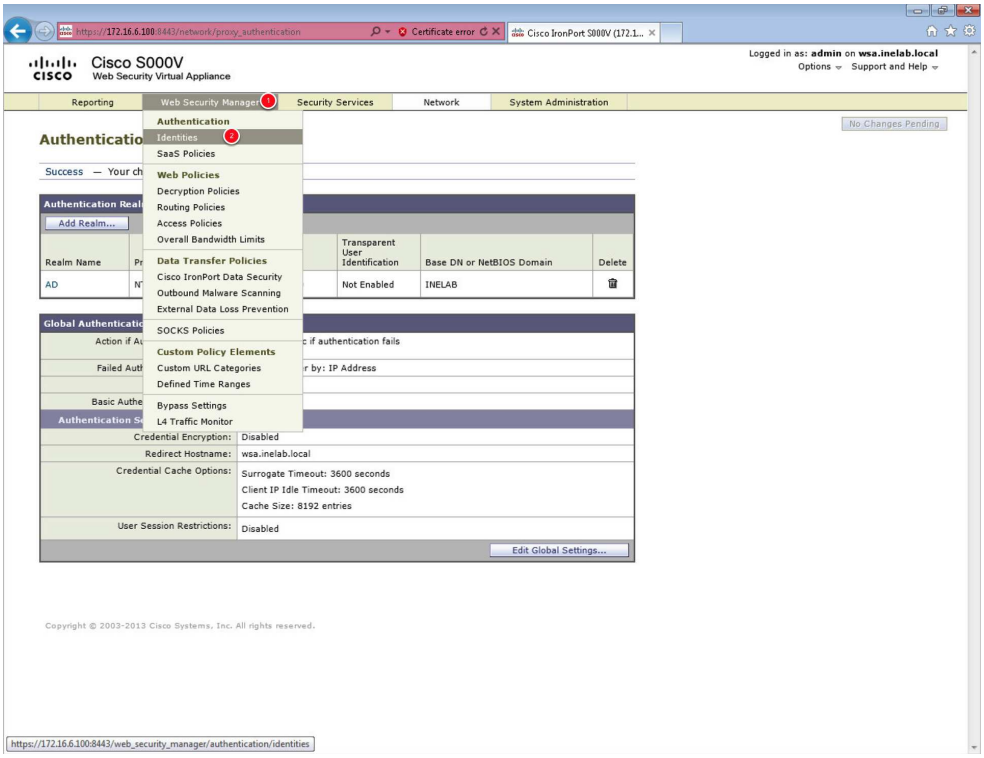


Configure the LDAP server IP address, Based DN (domain name), user attribute to perform authentication against, and Bind DN (LDAP location for the user performing LDAP queries):

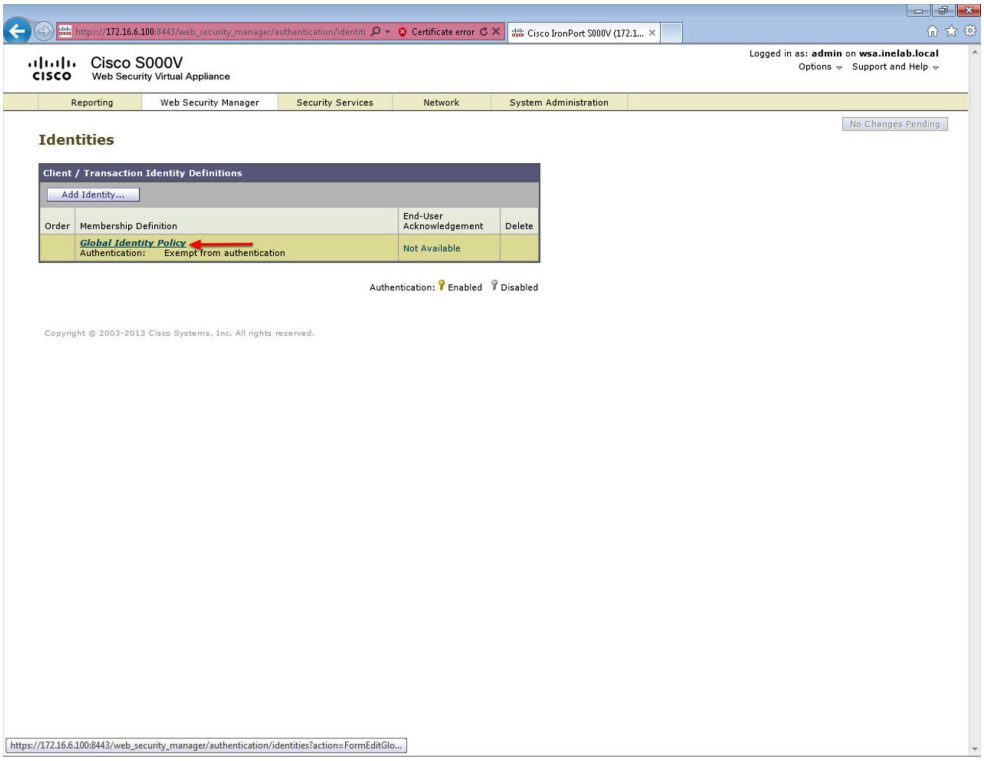


Navigate through the GUI to enable authentication for the global identity policy:

https://t.me/learningnets

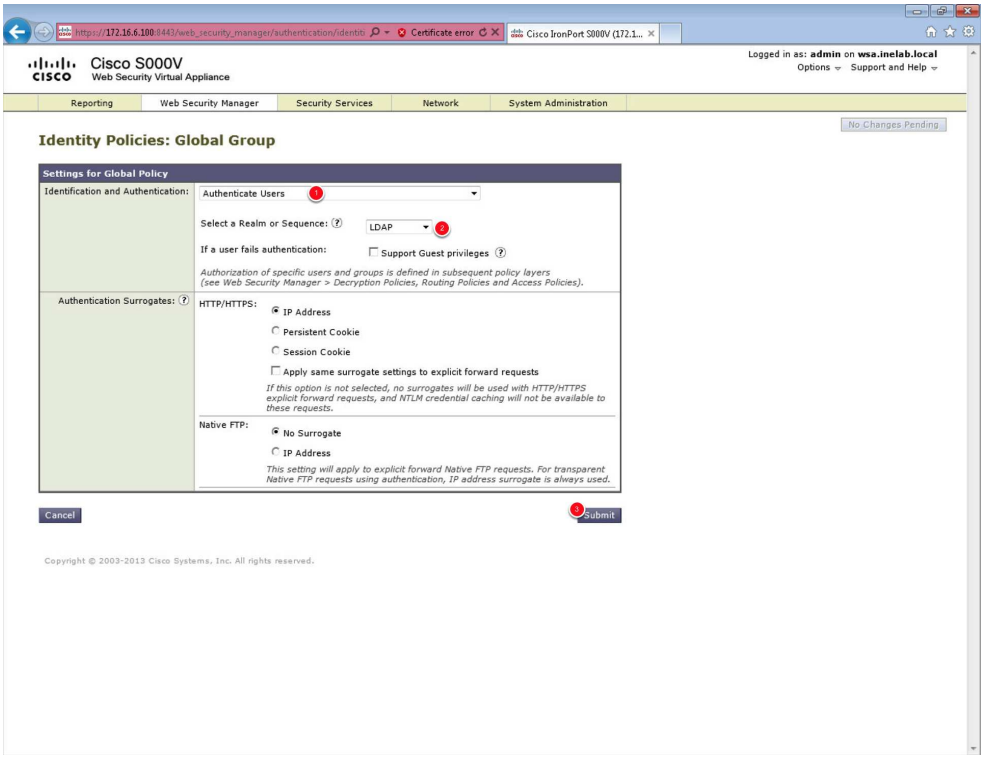


Click the **Global Identity Policy** hyperlink:

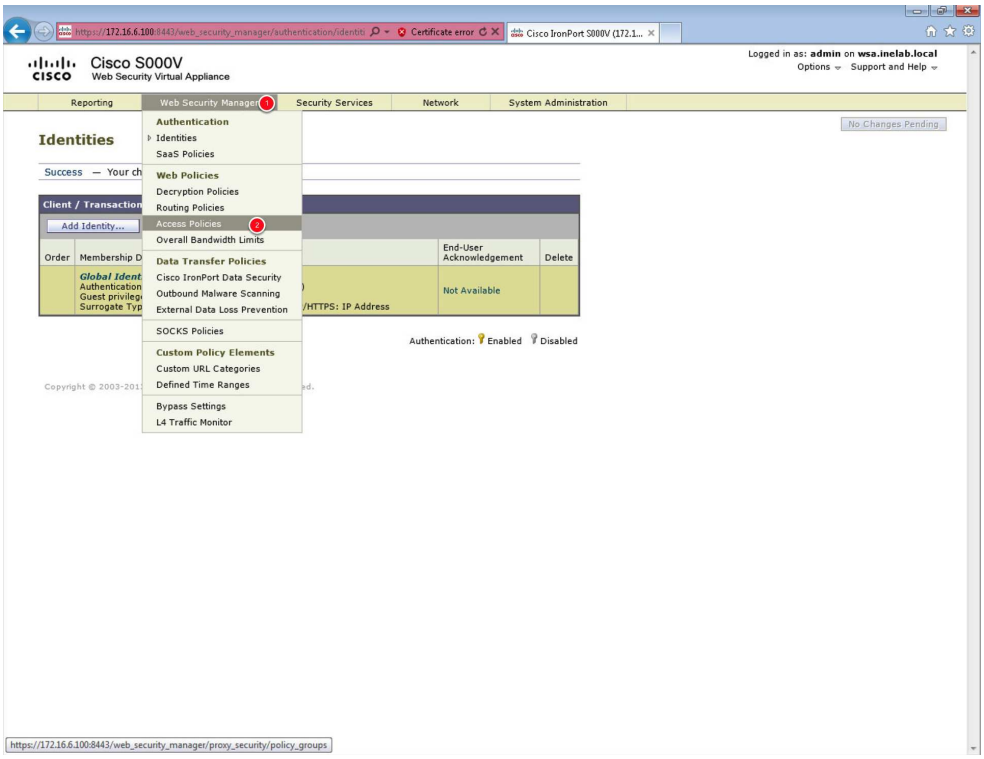


Enable Basic authentication, using the previously defined LDAP realm:

https://t.me/learningnets

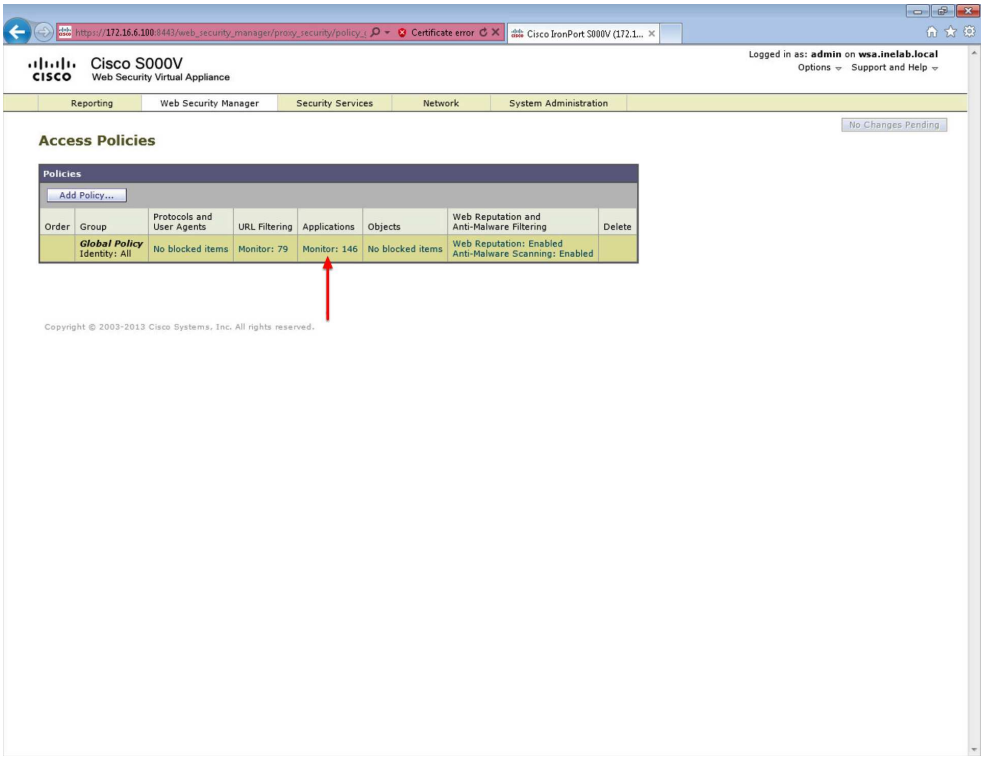


Navigate through the GUI to modify access policy settings for the authenticated users:

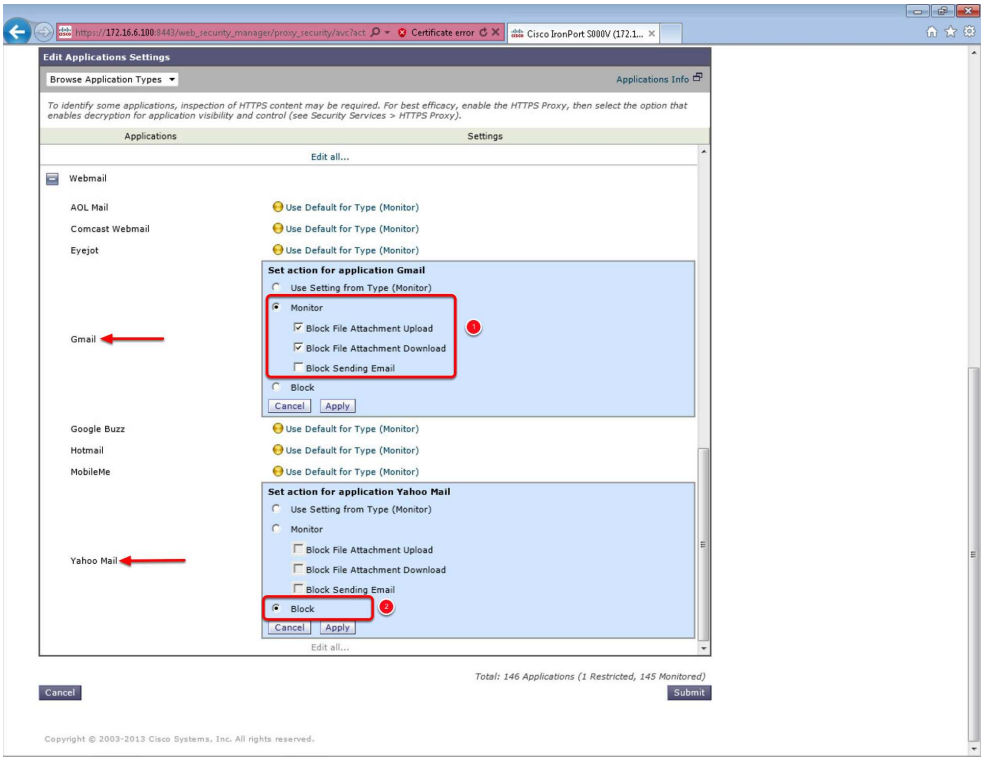


Click to change the application's visibility and control settings:

https://t.me/learningnets



Configure to block Google mail file upload/download and Yahoo mail completely:



Verification

Verify WCCP settings on R4, and confirm that WCCP Client registered successfully:

```
R4#show ip wccp summary
```

https://t.me/learningnets

```
WCCP version 2 enabled, 1 service

Service     Clients   Routers   Assign      Redirect   Bypass
-------     -------   -------   ------      --------   ------
Default routing table (Router Id: 150.1.4.4):93          1         1         MASK        GRE         GRE
!
!R4#show ip wccp
Global WCCP information:
    Router information:
        Router Identifier:                  150.1.4.4
        Protocol Version:                   2.0
 Service Identifier: 93
        Number of Service Group Clients:    1
        Number of Service Group Routers:    1
        Total Packets s/w Redirected:       99
          Process:                          0
          CEF:                              99 Service mode:                          Closed
Service Access-list:                  WCCP_CLOSED
        Total Packets Dropped Closed:       19 Redirect Access-list:                REDIRECTED_TRAFFIC
        Total Packets Denied Redirect:      0
        Total Packets Unassigned:           0 Group Access-list:                    WSA_GROUP
        Total Messages Denied to Group:     0
        Total Authentication failures:      0
        Total GRE Bypassed Packets Received: 0
!
!R4#show ip wccp 93 detail | e 00
WCCP Client information:
        Protocol Version:       2.0 State:                    Usable
Redirection:            GRE
Packet Return:          GRE

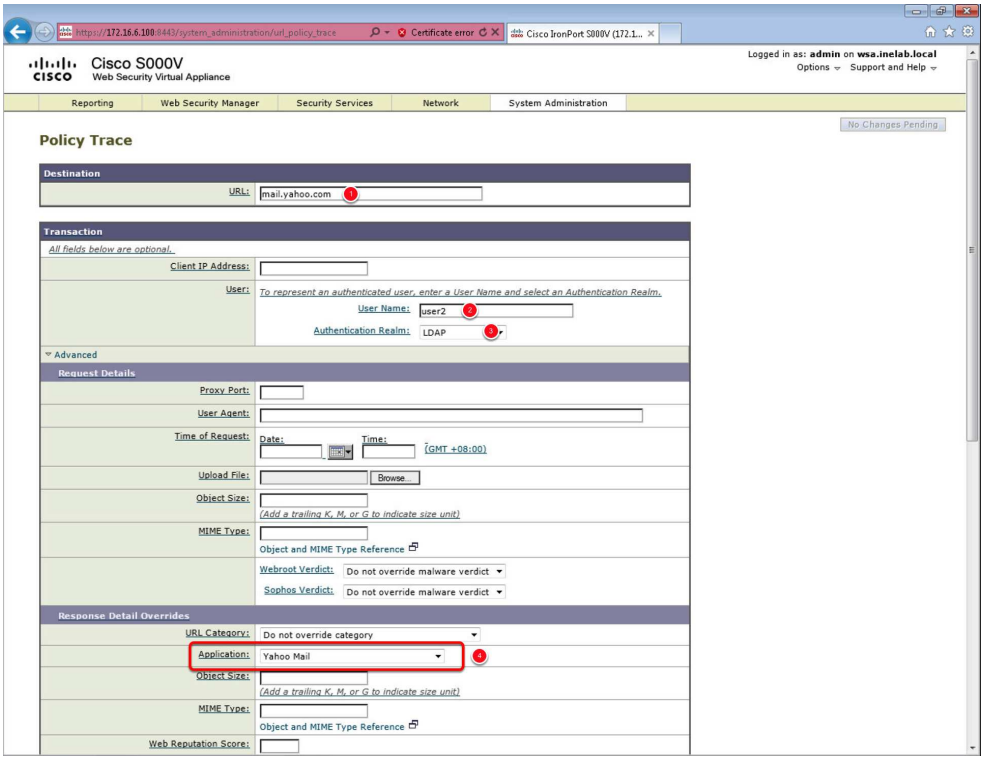
        Packets Redirected:     99
        Connect Time:           01:33:41
        Assignment:             MASK

        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------

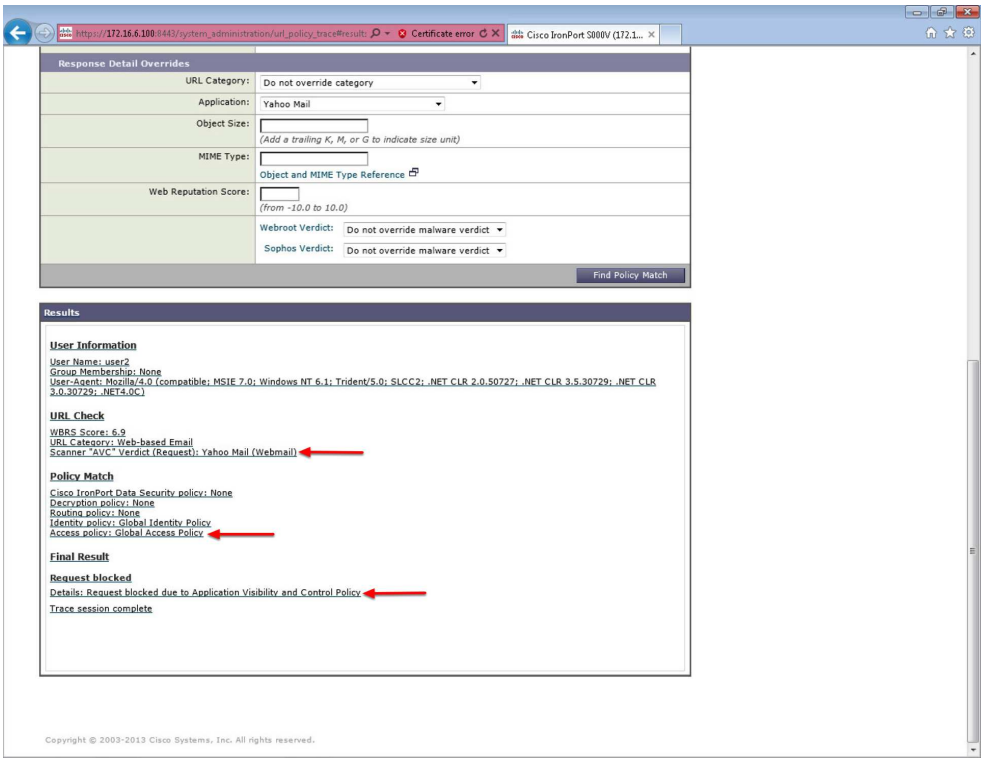
        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
```

Simulate an authenticated connection request to Yahoo mail service:

https://t.me/learningnets

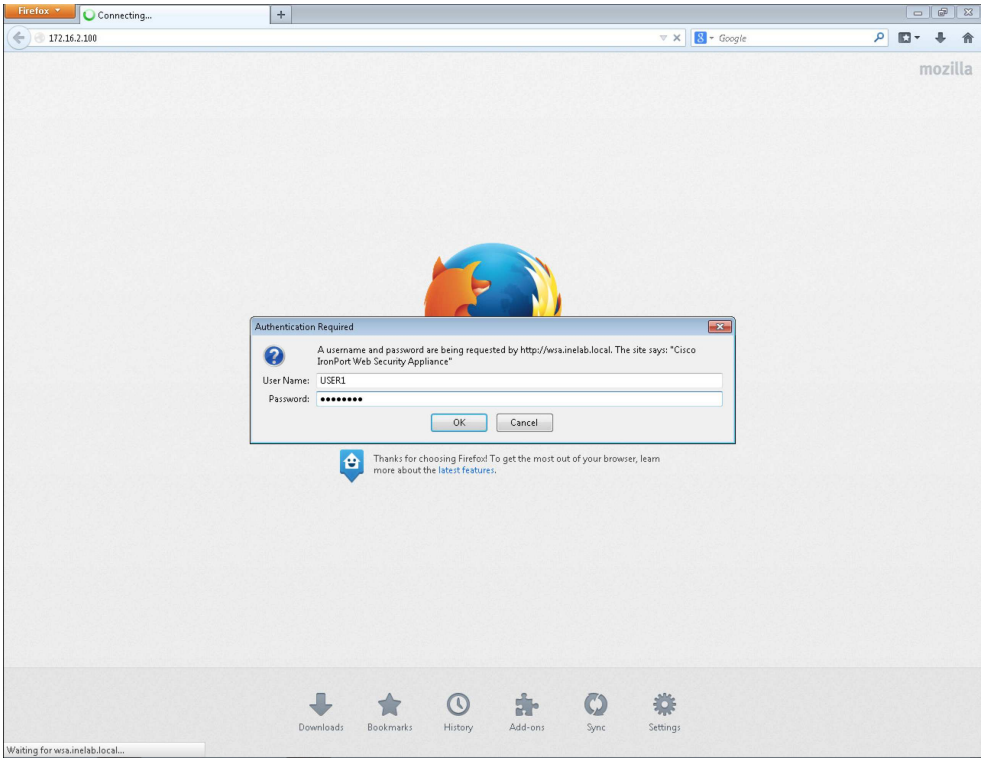


Based on the configured AVC settings, connection is not allowed:

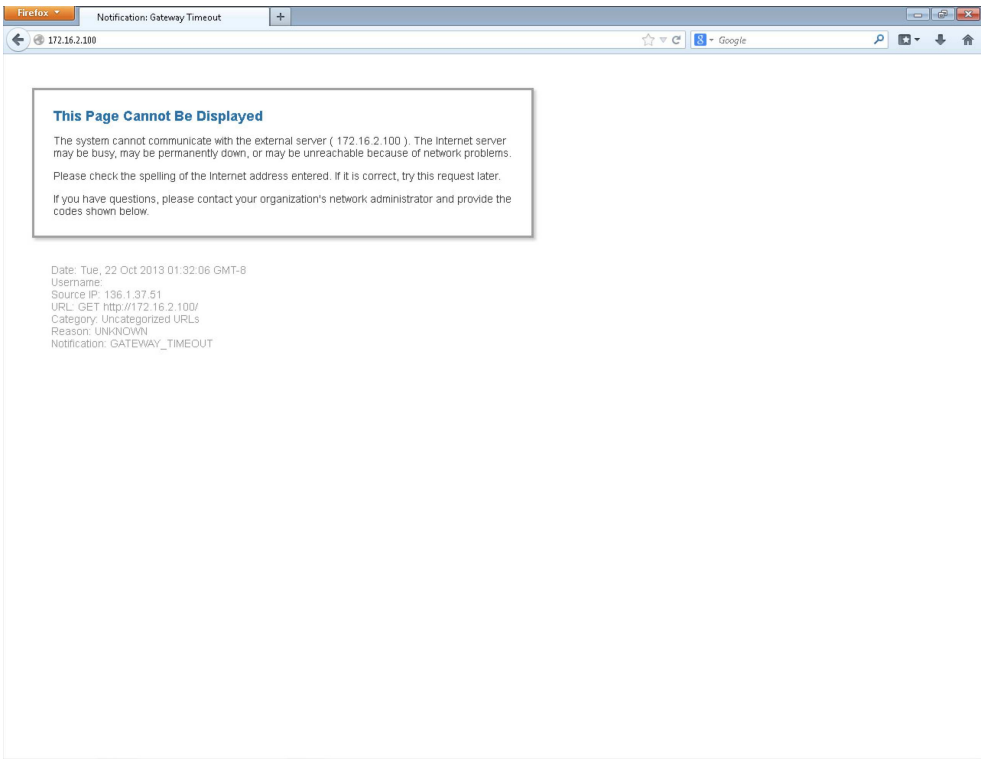


From **TEST-PC-B**, initiate an HTTP session toward the ACS server to trigger the WCCP redirection; WSA will ask for authentication:

https://t.me/learningnets



Because the HTTP service is not active on the ACS, the connection will eventually time out and an error message will be presented to the user:



Verify the WSA access logs to ensure that traffic was proxied through the WSA:

https://t.me/learningnets

```
1382378166.727 150005
136.1.37.51 NONE/504 0 GET http://172.16.2.100/ "user1@LDAP" DIRECT/172.16.2.100 - OTHER-NONE-DefaultGroup-NONE-NONE
 <nc,-3.5,0,"-",0,0,0,-,"-",-,-,-,"-",-,-,"-","-",-,-,nc,-,"-","-","Unknown","Unknown","-","-",0.00,0,-,"Unknown","-
```

Verify the authentication cache on the WSA:

```
wsa.inelab.local> authcache



Choose the operation you want to perform:
- FLUSHALL - Flush all entries from auth cache
- FLUSHUSER - Flush specific user entry from auth cache
- LIST - List all entries from auth cache
- SEARCH - Search all entries from auth cache[]> LIST

List may print a lot of entries. Are you sure? [Y]> y
 user1@LDAP




1 entries in authentication cache

Choose the operation you want to perform:
- FLUSHALL - Flush all entries from auth cache
- FLUSHUSER - Flush specific user entry from auth cache
- LIST - List all entries from auth cache
- SEARCH - Search all entries from auth cache
[]>
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

WSA Transparent Mode with ASA

Task

- Configure WSA and ASA3 for WCCP integration, and redirect all HTTP/HTTPS traffic received inbound on VLAN 93.
- Your configuration should permit enabling WCCP authentication only for HTTP traffic if desired.
- Ensure that the default action for HTTPS traffic is set to **pass through**, and do not inspect files larger than 10 MB.
- ASA3 should control which **WCCP Clients** can register.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the WSA to the initial configuration (not to factory defaults), load it again by navigating to **System Administration > Configuration File > Load Configuration**, select the **WSA-DEFAULT.xml** file, click **Load**, and commit changes afterward.

Overview

ASA supports only GRE integration for WCCP, and only for the inbound direction. To allow WCCP authentication for only HTTP traffic, two service groups have been created, one for HTTP and one for HTTPS. The default action for an HTTPS

https://t.me/learningnets

session is taken only if no prior decision has been made based on the certificate, URL categories, or WBRS results.

For the **group-list** access-list controlling which **WCCP Clients** can register a service, ASA supports only extended access-lists, where the source must be the **WCCP Client**—in our case WSA. WCCP service group IDs between 90 and 97 are user configurable and can identify any protocol, but there are some dedicated service group ID, the most important being:

- service 0 for HTTP
- service 60 for FTP
- service 70 for HTTPS

While in general you can use any group ID number between 1-255 for any protocol you want to redirect to WSA (HTTP or HTTPS), it is recommended to use values between 90-97 to avoid possible incompatibilities between WCCP clients and WCCP servers.

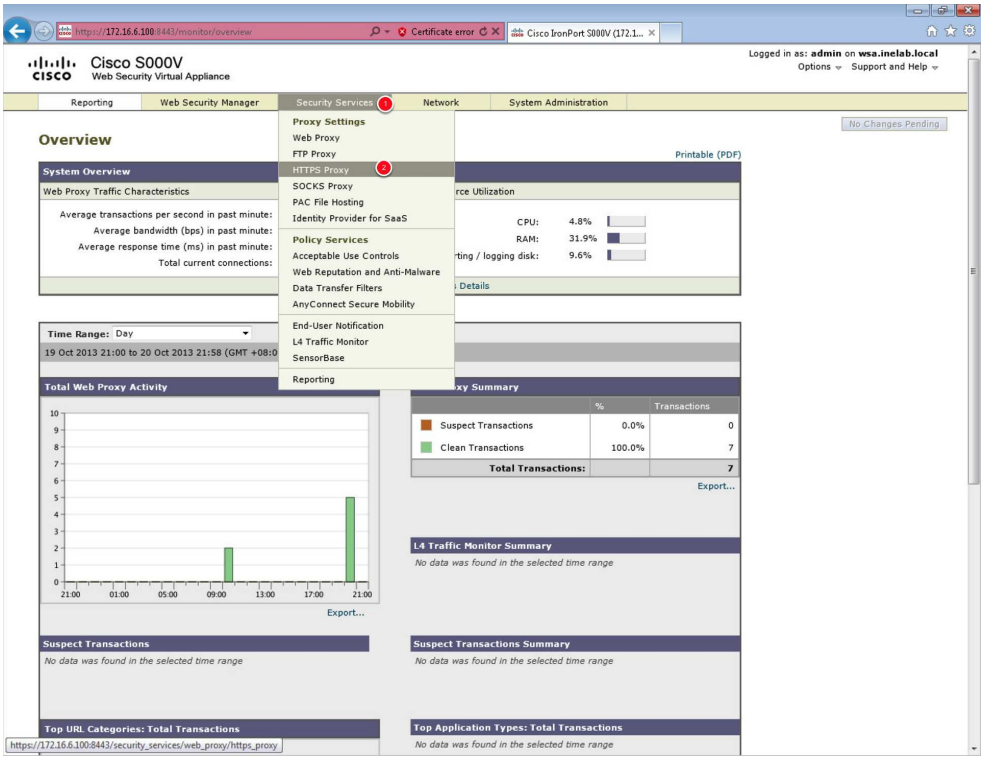
Configuration

ASA3:

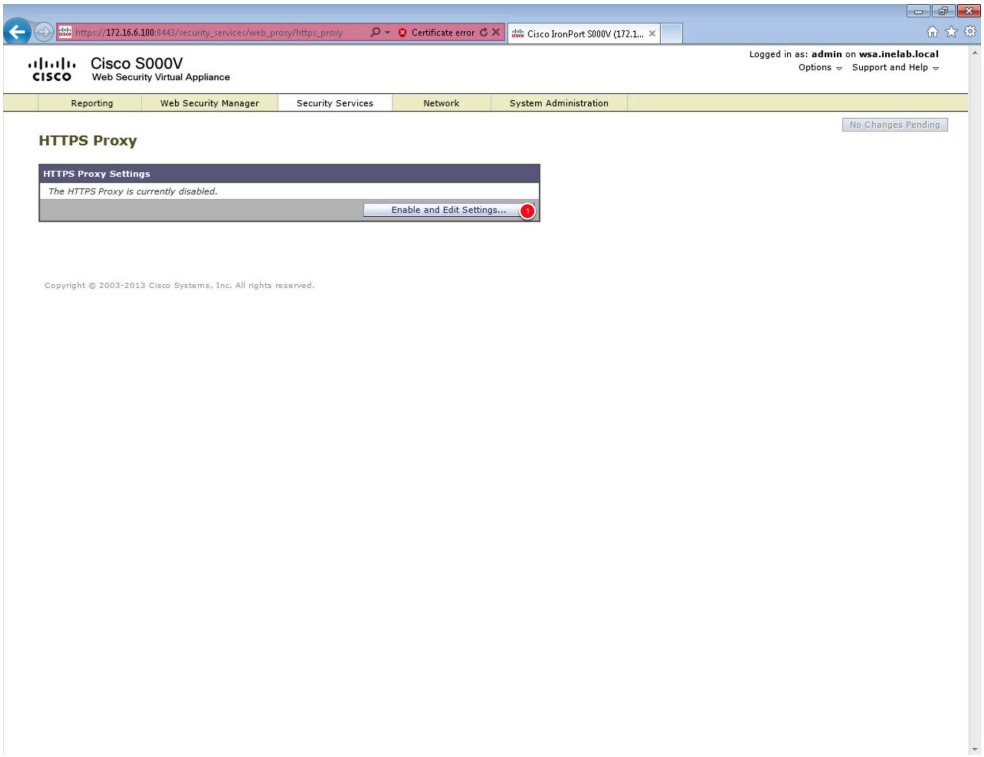
```
access-list REDIRECTED_HTTP extended permit tcp any any eq 80
access-list REDIRECTED_HTTPS extended permit tcp any any eq 443
!
access-list WSA_GROUP extended permit ip host 172.16.6.100 any
!
wccp 94 redirect-list REDIRECTED_HTTP group-list WSA_GROUP
wccp 95 redirect-list REDIRECTED_HTTPS group-list WSA_GROUP
!
wccp interface VLAN93 94 redirect in
wccp interface VLAN93 95 redirect in
```

Navigate through the GUI for HTTPS proxy activation:

https://t.me/learningnets

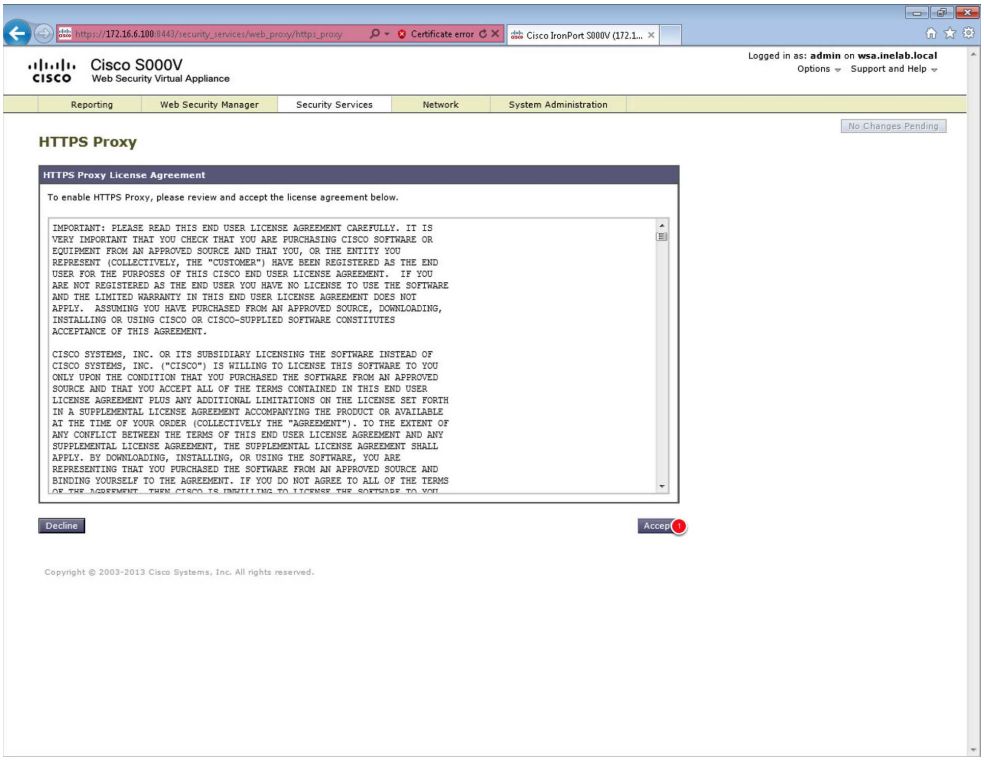


Click the **Enable and Edit Settings** button:

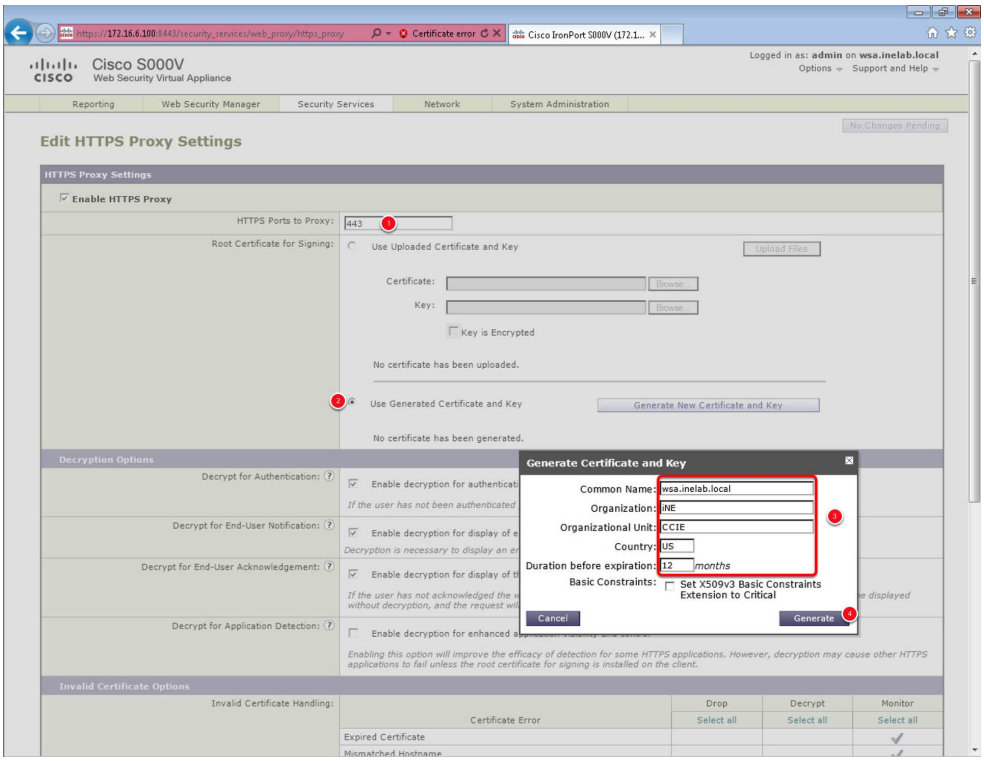


Accept the license agreement:

https://t.me/learningnets

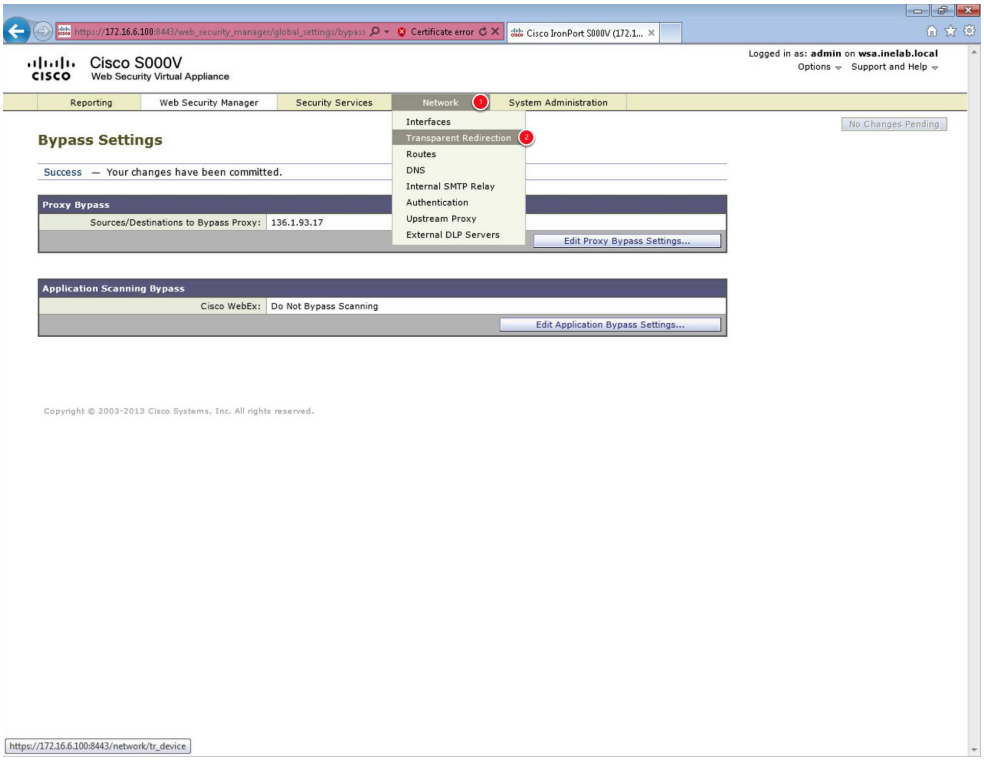


Specify the HTTPS proxy port (443) and generate a certificate to be used for HTTPS decryption process; wait for the certificate to be generated and click the **Submit** button:

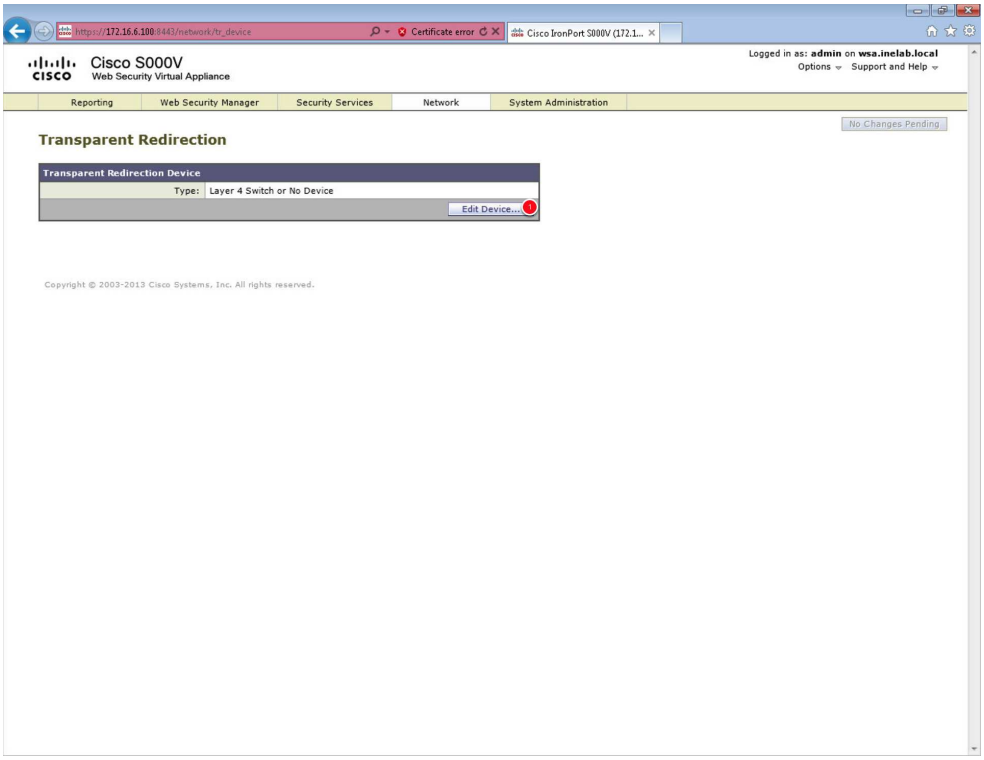


Browse through the GUI to configure the transparent mode functionality:

https://t.me/learningnets

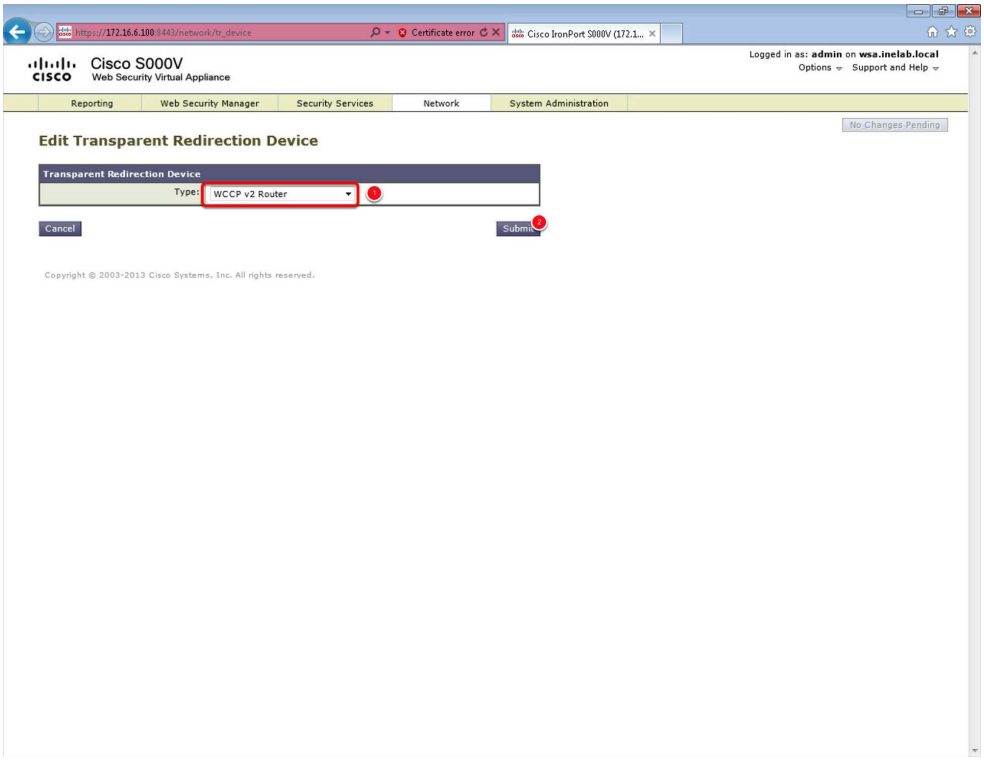


Click the **Edit Device** button:

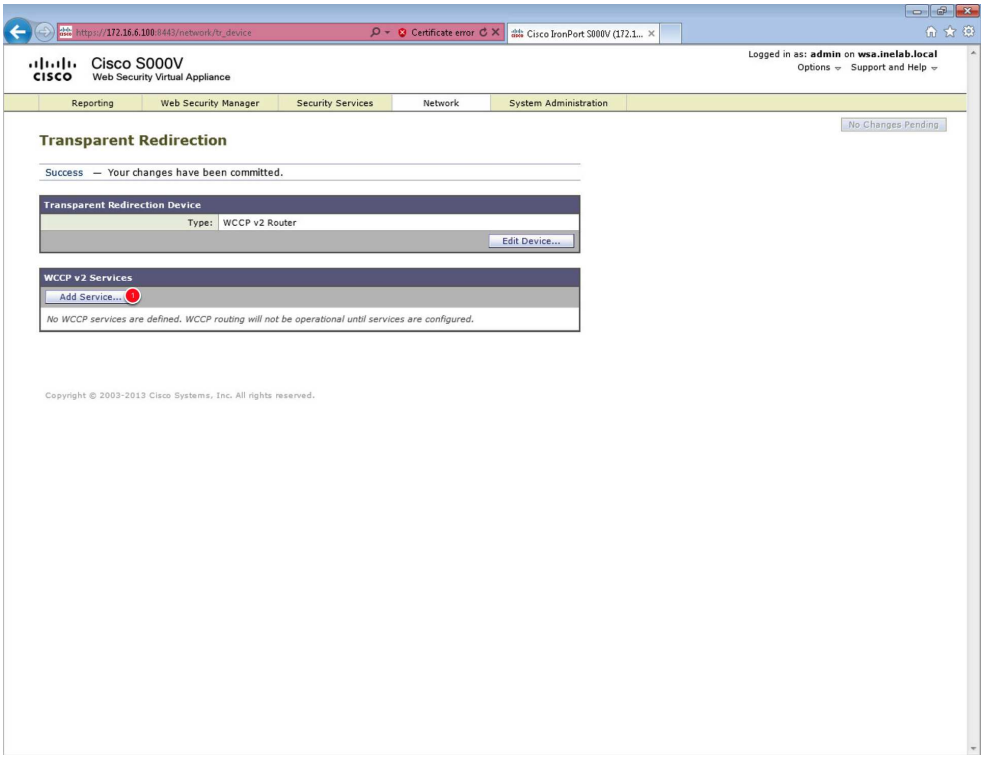


Select the **WCCP v2 Router** option:

https://t.me/learningnets

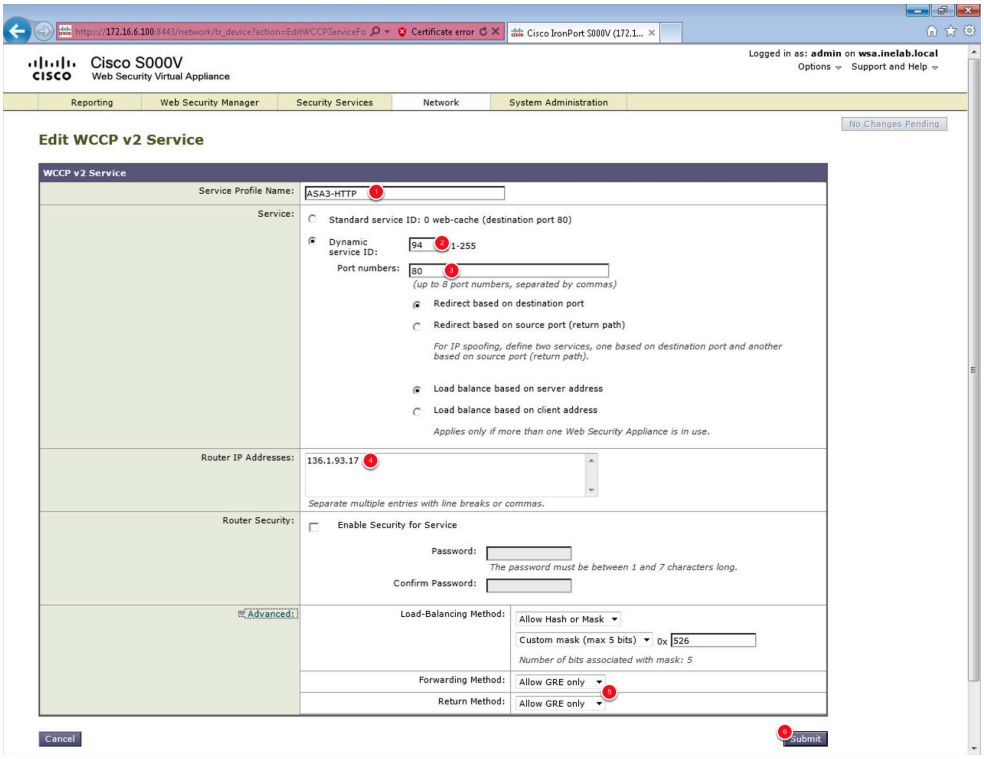


Click the **Add Service** button:

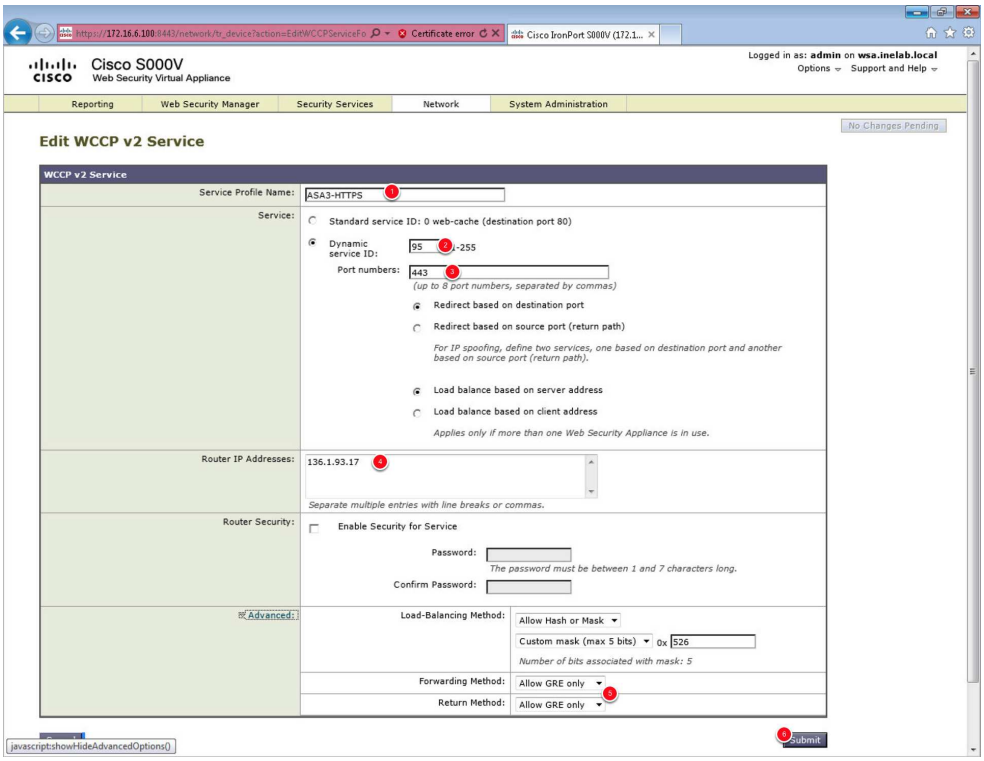


Configure the service group ID for HTTP traffic (the same as on ASA3) and the ports to be used for traffic redirection (80), and enforce GRE mode:

https://t.me/learningnets

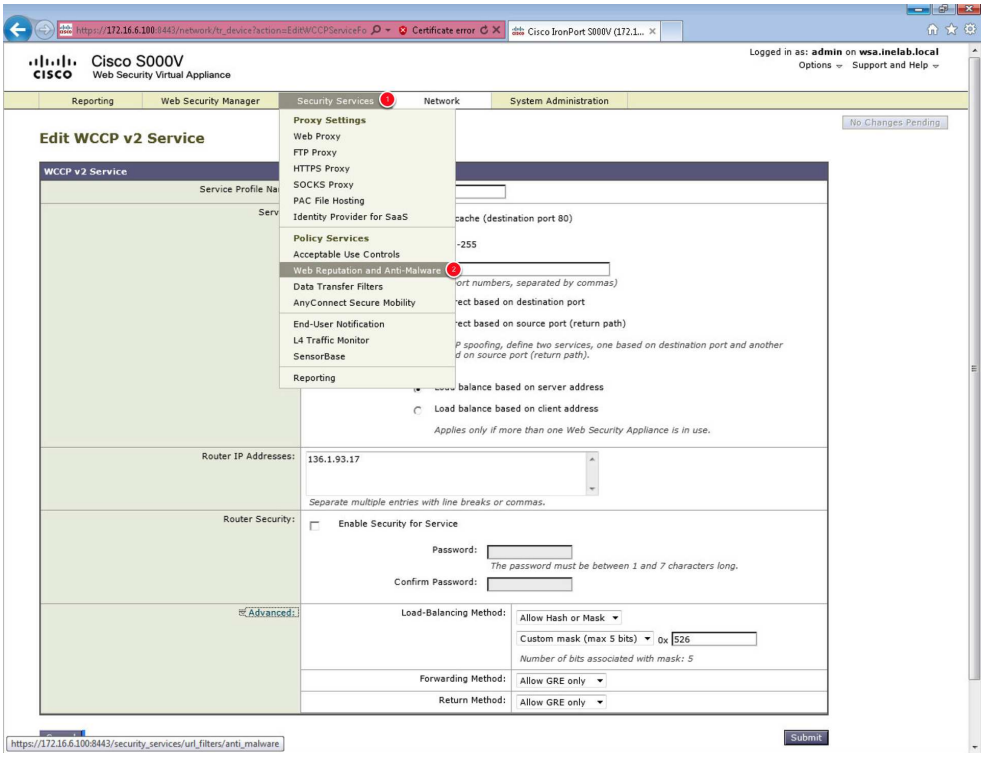


Configure the service group ID for HTTPS traffic (the same as on ASA3) and the ports to be used for traffic redirection (443), and enforce GRE mode:

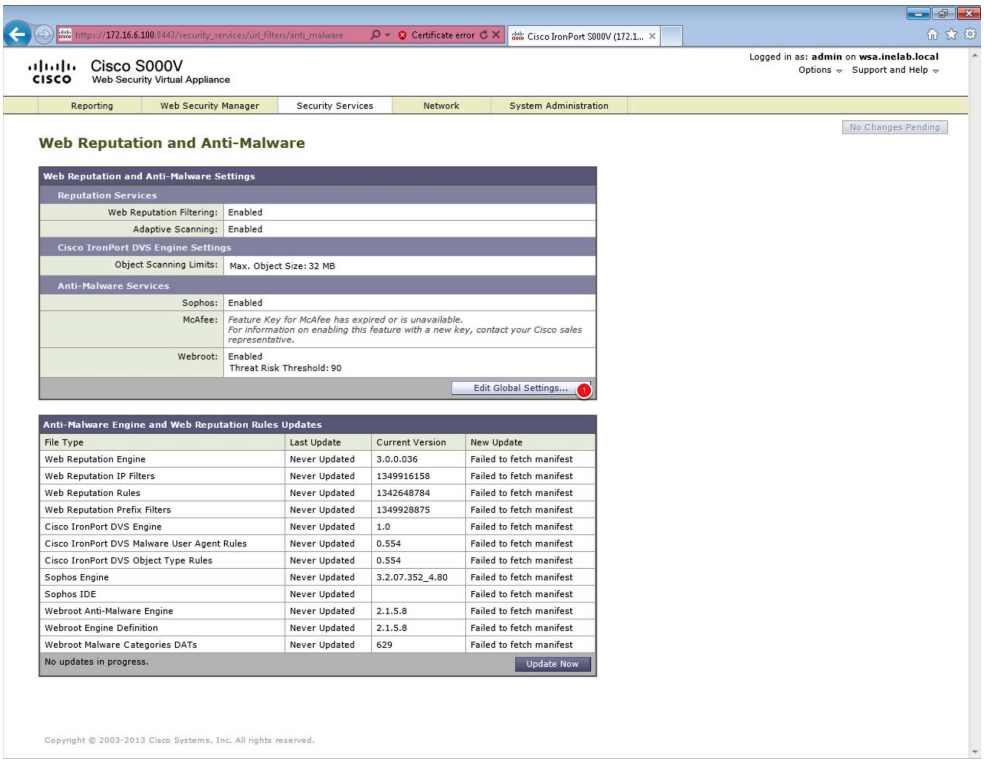


Navigate through the GUI for antimalware settings:

https://t.me/learningnets

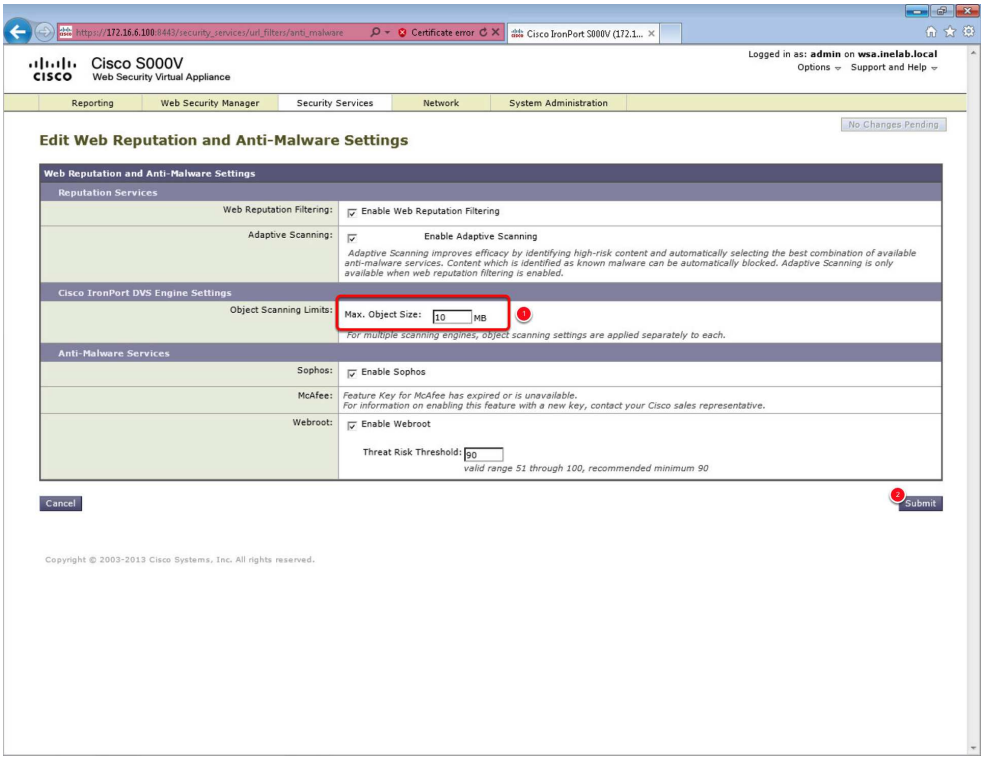


Click the **Edit Global Settings** button:

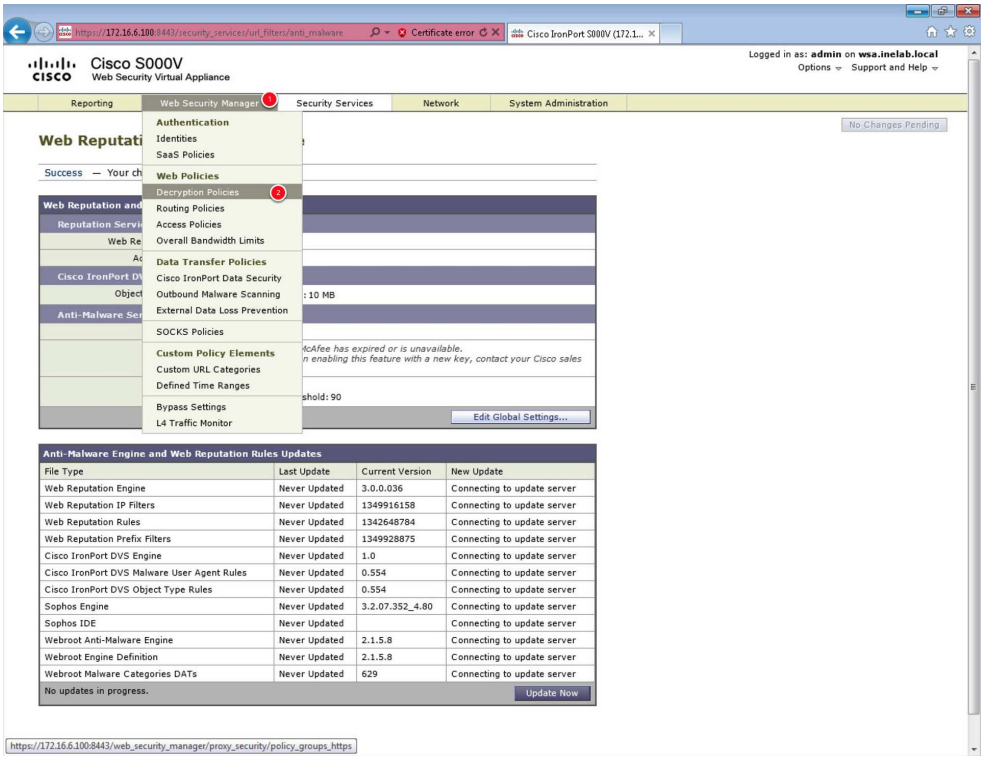


Modify the maximum size for object scanning:

https://t.me/learningnets

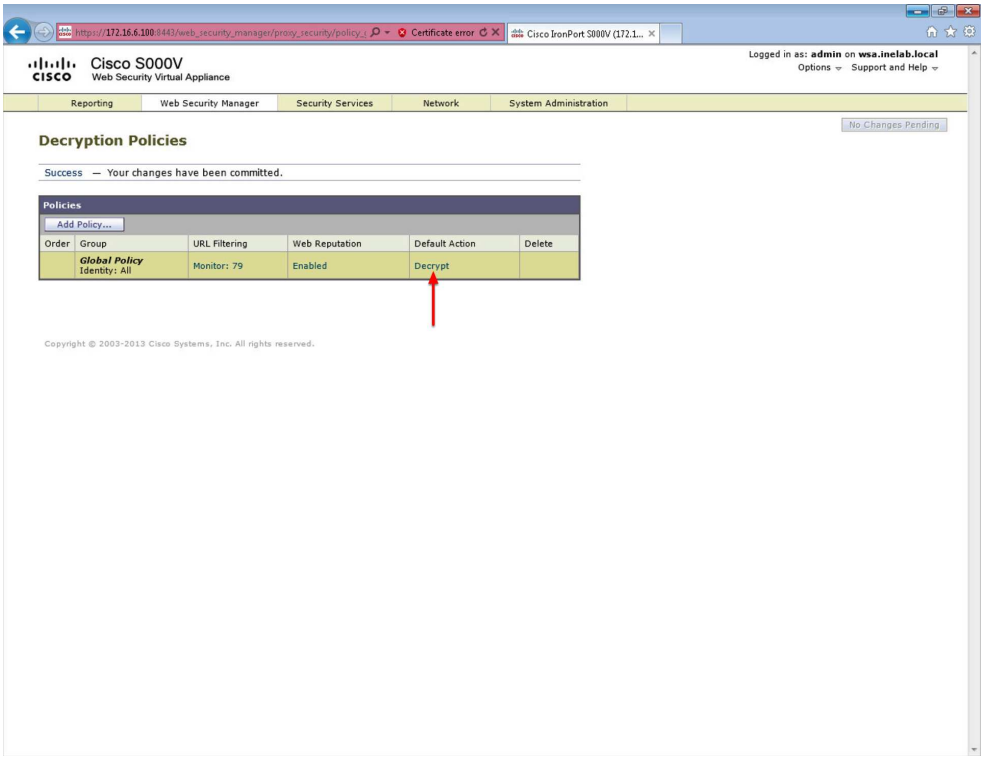


Navigate through the GUI for decryption policies settings:

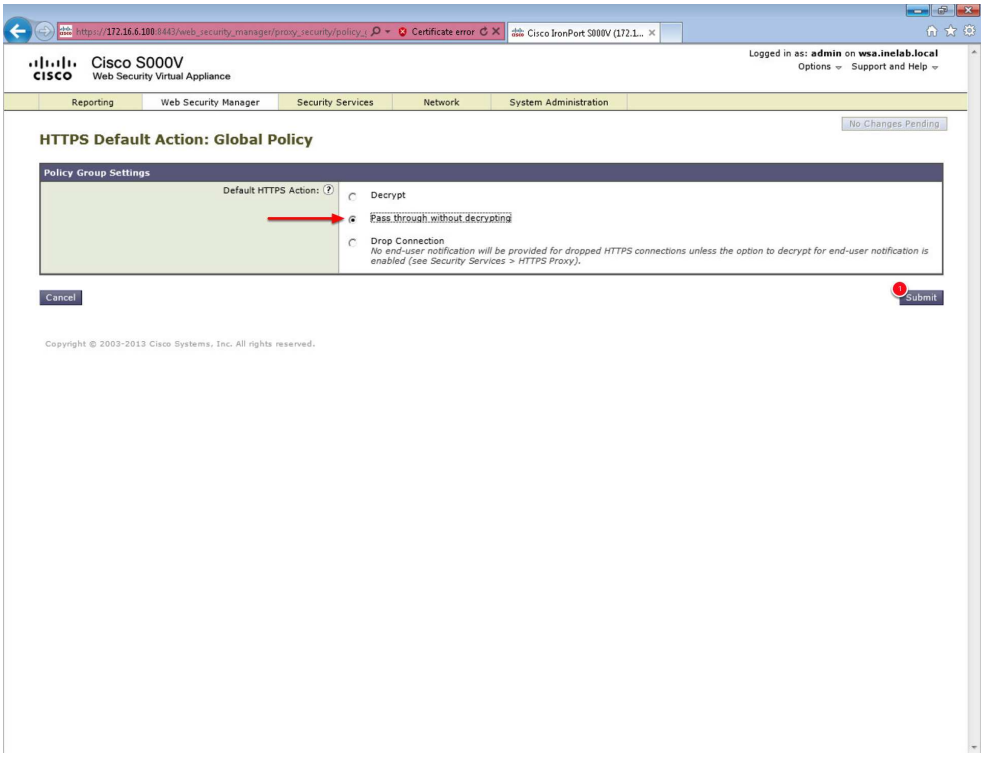


Click the decrypt default action:

https://t.me/learningnets



Change the default setting to **pass through**:



Verification

Verify WCCP settings on ASA3, and confirm that both services are active:

```
ASA3# show wccp
```

https://t.me/learningnets

```
Global WCCP information:
    Router information:
        Router Identifier:                  136.1.93.17
        Protocol Version:                   2.0

    Service Identifier: 94
        Number of Cache Engines:            1
        Number of routers:                  1
        Total Packets Redirected:           0 Redirect access-list:                 REDIRECTED_HTTP
        Total Connections Denied Redirect:  0
        Total Packets Unassigned:           0 Group access-list:                    WSA_GROUP
        Total Messages Denied to Group:     0
        Total Authentication failures:      0
        Total Bypassed Packets Received:    0

    Service Identifier: 95
        Number of Cache Engines:            1
        Number of routers:                  1
        Total Packets Redirected:           0 Redirect access-list:                 REDIRECTED_HTTPS
        Total Connections Denied Redirect:  0
        Total Packets Unassigned:           0 Group access-list:                    WSA_GROUP
        Total Messages Denied to Group:     0
        Total Authentication failures:      0
        Total Bypassed Packets Received:    0
!
!ASA3# show wccp 94 detail

WCCP Cache-Engine information: Web Cache ID:          172.16.6.100
        Protocol Version:      2.0 State:                  Usable
        Initial Hash Info:     00000000000000000000000000000000
                               00000000000000000000000000000000
        Assigned Hash Info:    FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
                               FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
        Hash Allotment:        256 (100.00%) Packets Redirected:    0
        Connect Time:          00:39:36
!
!ASA3# show wccp 95 detail

WCCP Cache-Engine information: Web Cache ID:          172.16.6.100
        Protocol Version:      2.0 State:                  Usable
        Initial Hash Info:     00000000000000000000000000000000
                               00000000000000000000000000000000
        Assigned Hash Info:    FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
                               FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
        Hash Allotment:        256 (100.00%)
```

https://t.me/learningnets

```
Packets Redirected:     0

        Connect Time:           00:40:10
```

Initiate an HTTP connection from **TEST-PC-B** to R6's Loopback0 address (it should be allowed), and then verify WCCP statistics:

```
ASA3# show wccp 94 detail

WCCP Cache-Engine information:
        Web Cache ID:           172.16.6.100
        Protocol Version:       2.0
        State:                  Usable
        Initial Hash Info:      00000000000000000000000000000000
                                00000000000000000000000000000000
        Assigned Hash Info:     FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
                                FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
        Hash Allotment:         256 (100.00%) Packets Redirected:     4

        Connect Time:           00:43:22
```

Verify the access logs from the WSA:

```
1382391594.421 40 136.1.37.51
TCP_MISS/200 1975 GET http://150.1.6.6/ - DIRECT/150.1.6.6 text/html DEFAULT_CASE_12-DefaultGroup-DefaultGroup-NONE-
 <nc,-3.5,0,"-",0,0,0,1,"-",-,-,-,"-",0,262686,"-","-",-,-,nc,nc,"Unknown","-","Unknown","Unknown","-","-",395.00,0,
```

Initiate an HTTPS connection from **TEST-PC-B** to R6's Loopback0 address (it should be allowed), and then verify WCCP statistics:

https://t.me/learningnets

```
ASA3# show wccp 95 detail



WCCP Cache-Engine information:
        Web Cache ID:          172.16.6.100
        Protocol Version:      2.0
        State:                 Usable
        Initial Hash Info:     00000000000000000000000000000000
                               00000000000000000000000000000000
        Assigned Hash Info:    FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
                               FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
        Hash Allotment:        256 (100.00%) Packets Redirected:     15

        Connect Time:          00:45:37
```

Verify the access logs from the WSA:

```
1382391703.119 2191 136.1.37.51
TCP_MISS_SSL/200 0 TCP_CONNECT 150.1.6.6:443 - DIRECT/150.1.6.6 - DECRYPT_ADMIN-UNRECOGNIZED_ROOT_CERT_7-DefaultGrou
 <nc,ns,-,"-",-,-,-,-,"-",-,-,-,"-",-,-,"-","-",-,-,nc,-,"-","-","Unknown","Unknown","-","-",0.00,0,-,"-","-"> -
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Initialization

Task

- Set up the IPS hostname to **IPS4240**.
- Configure the management IP address to be 136.1.43.100/24, set default gateway to 136.1.43.9.
- Ensure HTTP management access is allowed only from VLAN 37 on port 446.
- Both telnet and ssh services should be active, present a login banner of **Authorized Access Only**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.

Overview

IPS can be managed via telnet/ssh/http/https. Telnet and HTTP services are disabled by default, whereas ssh and HTTPS are enabled. IPS restricts remote management sessions based on an access-list, which by default has no entries; therefore, it denies all sessions. Optionally, the IPS can be initialized by using the `setup` command.

Configuration

IPS:

```
sensor# configure terminal
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-name IPS4240
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
```

https://t.me/learningnets

```
sensor(config-hos-net)# telnet-option enabled
sensor(config-hos-net)# login-banner-text Authorized Access Only
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service web-server
sensor(config-web)# enable-tls false
sensor(config-web)# port 446
sensor(config-web)# exit
Apply Changes?[yes]:
```

SW3:

```
interface GigabitEthernet1/0/10
 switchport mode access
 switchport access vlan 43
 spanning-tree portfast
 no shutdown
```

Verification

Test to verify that default gateway and **TEST-PC-B** (check the DHCP assigned IP address for it; it may differ on your rack) are reachable:

```
sensor# ping 136.1.43.9
PING 136.1.43.9 (136.1.43.9): 56 data bytes
64 bytes from 136.1.43.9: seq=0 ttl=255 time=1.000 ms
64 bytes from 136.1.43.9: seq=1 ttl=255 time=3.000 ms
64 bytes from 136.1.43.9: seq=2 ttl=255 time=2.000 ms
64 bytes from 136.1.43.9: seq=3 ttl=255 time=1.000 ms

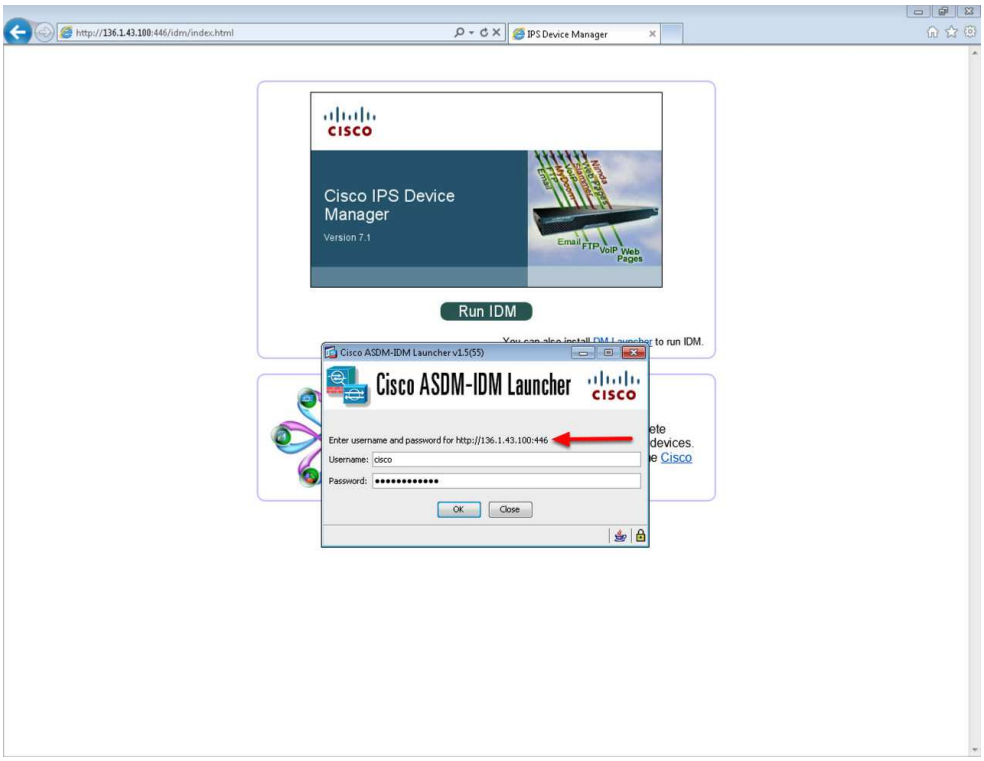
--- 136.1.43.9 ping statistics ---4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max = 1.000/1.750/3.000 ms
!
!sensor# ping 136.1.37.51
PING 136.1.37.51 (136.1.37.51): 56 data bytes
64 bytes from 136.1.37.51: seq=0 ttl=127 time=20.000 ms
64 bytes from 136.1.37.51: seq=1 ttl=127 time=5.000 ms
64 bytes from 136.1.37.51: seq=2 ttl=127 time=4.000 ms
64 bytes from 136.1.37.51: seq=3 ttl=127 time=4.000 ms
```

https://t.me/learningnets

```
--- 136.1.37.51 ping statistics ---4 packets transmitted, 4 packets received, 0% packet loss

round-trip min/avg/max = 4.000/8.250/20.000 ms
```

Verify that HTTP connection is successful on port 446:



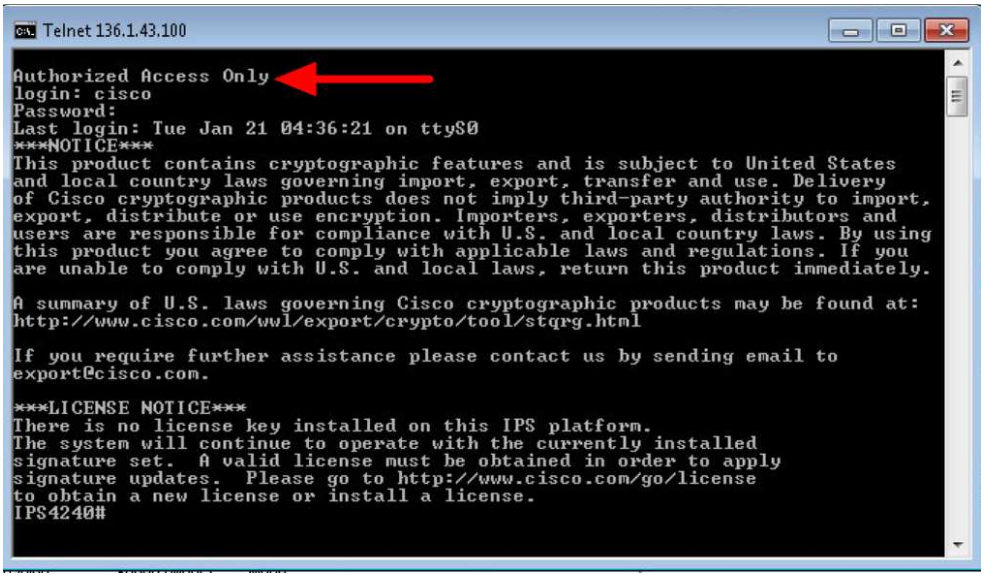
Verify the configured settings:

```
IPS4240(config-hos-net)#show settings
   network-settings
   ------------------------------------------------ host-ip: 136.1.43.100/24,136.1.43.9
 default: 192.168.1.2/24,192.168.1.1 host-name: IPS4240
 default: sensor telnet-option: enabled
 default: disabled
      access-list (min: 0, max: 512, current: 1)
      ------------------------------------------------ network-address: 136.1.37.0/24
         ------------------------------------------------
      ------------------------------------------------
      ftp-timeout: 300 seconds <defaulted> login-banner-text: Authorized Access Only
 default:
      dns-primary-server
      ------------------------------------------------
         disabled
         ------------------------------------------------
         ------------------------------------------------
      ------------------------------------------------
```

https://t.me/learningnets

```
      dns-secondary-server
      ------------------------------------------------
         disabled
         ------------------------------------------------
         ------------------------------------------------
      ------------------------------------------------
      dns-tertiary-server
      ------------------------------------------------
         disabled
         ------------------------------------------------
         ------------------------------------------------
      ------------------------------------------------
      http-proxy
      ------------------------------------------------
         no-proxy
         ------------------------------------------------
         ------------------------------------------------
      ------------------------------------------------
   ------------------------------------------------
```

Verify that telnet service is functional; check the login banner:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Management Access

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9.
- Ensure that HTTP management access is allowed only from VLAN 37.
- Disable a local user account after 5 failed authentication attempts and disconnect a session after 10 minutes of inactivity.
- Limit the maximum number of characters for local user passwords to 20.
- Configure IPS for RADIUS authentication:
  - RADIUS server is ACS2 at IP address 172.16.2.100.
  - User **ciscoips** should be given the highest privilege level upon successful RADIUS authentication.
  - If no role is received from RADIUS server upon successful authentication, automatically assign the **viewer** role.
  - If the RADIUS server is not available, the local username database should be used.
  - Console access should use only the local username database for authentication.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.

https://t.me/learningnets

If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.

To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

Cisco IPS supports remote user authentication through RADIUS, with fallback on the local user database if the RADIUS server is unavailable. To defend against possible access lockout, console access can be configured to use only the local user database, although by default uses both RADIUS and local user database. IPS supports RBAC (Role Based Access Control) for both local and RADIUS user authentication, with the following roles available:

- **administrator** has full access to IPS CLI and IDM.
- **operator** has unrestricted view access and can change configuration for: signature tuning, virtual sensor configuration assignment, manage routers, and change own password.
- **viewer** has unrestricted view access and can change own password.
- **service** has access directly to bash shell (no CLI access) and is used for advacned troubleshooting purposes.

When RADIUS authentication is being used, if no role is received from RADIUS as authorization, by default IPS will not assign any role and the session will therefore not be allowed because of unsuccesfull authentication; this is because IPS requires that a role be assigned for any authenticating user. The roles are assigned from RADIUS through Cisco AV Pair with one of following values:

- ips-role=administrator
- ips-role=operator
- ips-role=viewer
- ips-role=service
- ips-role=unspecified

Configuration

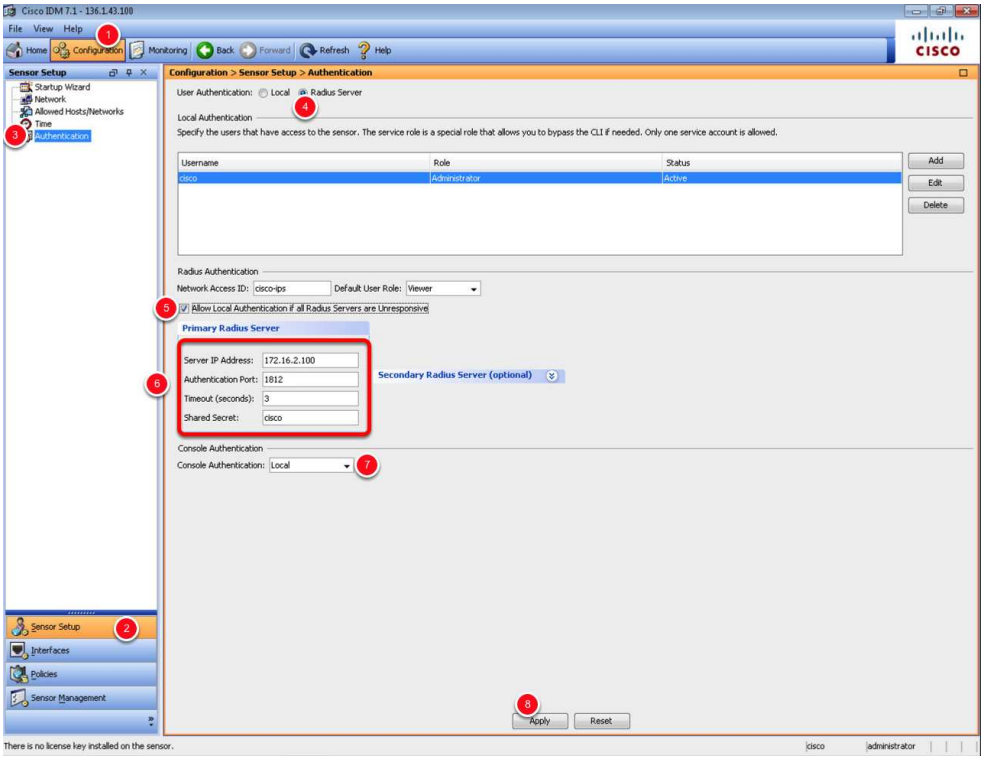
IPS:

https://t.me/learningnets

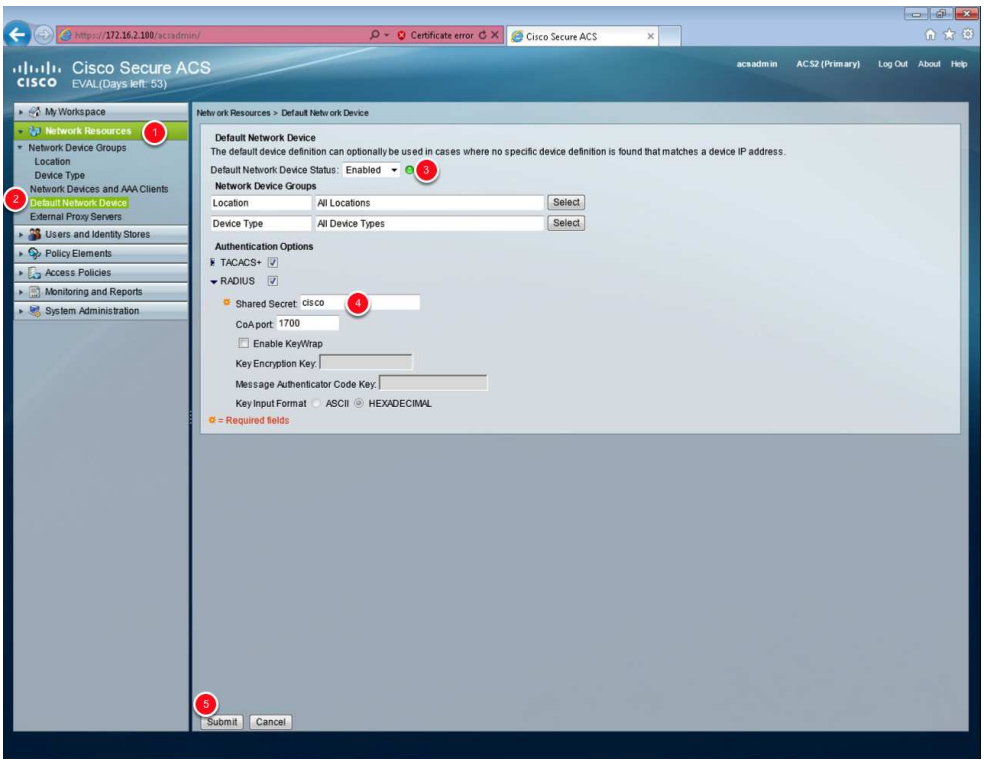
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service authentication
sensor(config-aut)# attemptLimit 5
sensor(config-aut)# cli-inactivity-timeout 10
sensor(config-aut)# password-strength
sensor(config-aut-pas)# size 6-20
sensor(config-aut-pas)# exit
sensor(config-aut)# exit
Apply Changes?[yes]:
Warning: The SSH server configuration has changed. The SSH server will not allow SSH clients to connect unless they
!
sensor(config)# service aaa
sensor(config-aaa)# aaa radius
sensor(config-aaa-rad)# primary-server
sensor(config-aaa-rad-pri)# server-address 172.16.2.100
sensor(config-aaa-rad-pri)# shared-secret cisco
sensor(config-aaa-rad-pri)# exit
sensor(config-aaa-rad)# local-fallback enabled
sensor(config-aaa-rad)# console-authentication local
sensor(config-aaa-rad)# default-user-role viewer
sensor(config-aaa-rad)# exit
sensor(config-aaa)# exit
Apply Changes?[yes]:
```

Configuring AAA from the GUI instead of CLI, using IDM, requires the following steps:

https://t.me/learningnets

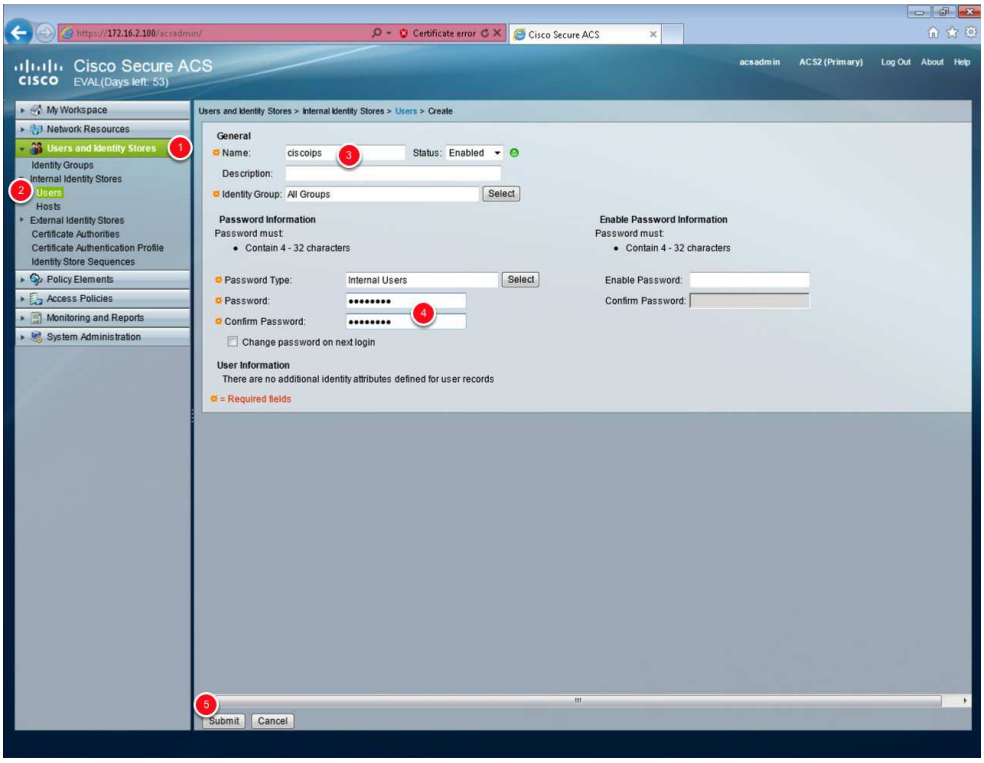


Add IPS as a NAD into ACS, or configure ACS for default network device (RADIUS only):

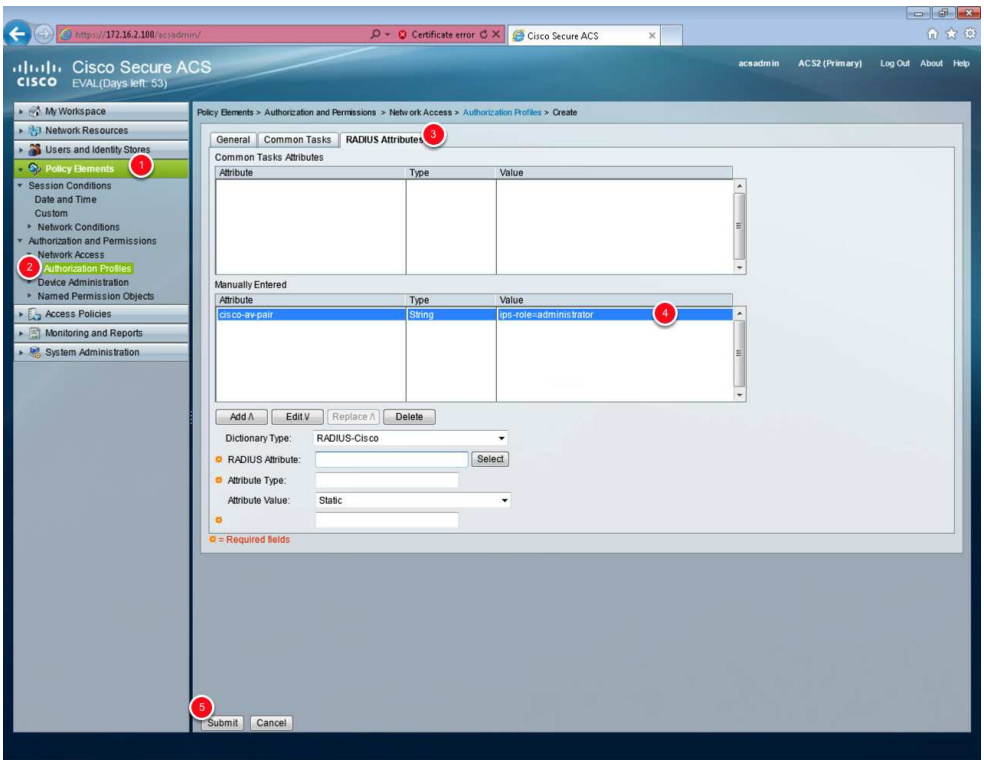


Configure the user into ACS internal store:

https://t.me/learningnets

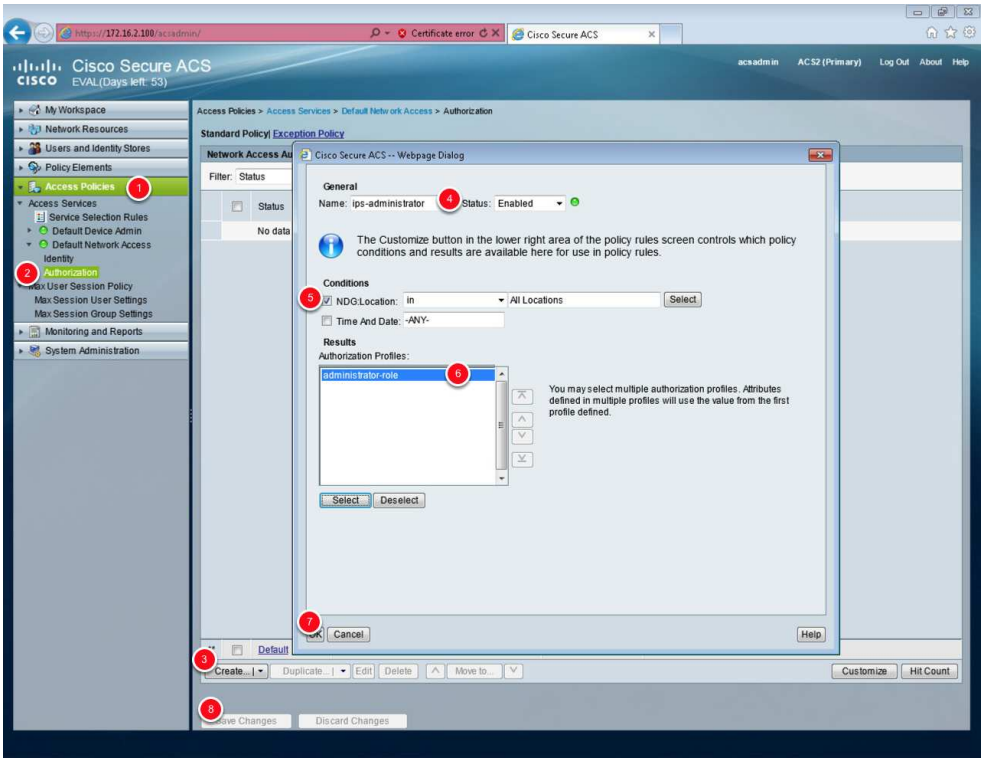


Configure the authorization profile, with the RADIUS Cisco AV pair value for the administrator role:



Configure the access policy for the RADIUS request:

https://t.me/learningnets



Verification

Verify the configured AAA settings:

```
sensor(config-aaa)# show settings
    aaa
    ----------------------------------------------
        radius
        ---------------------------------------------
            primary-server
            ---------------------------------------------- server-address: 172.16.2.100
                server-port: 1812 default: 1812
                shared-secret: cisco
                timeout: 3 default: 3
            ---------------------------------------------
            secondary-server
            ---------------------------------------------
                disabled
                ---------------------------------------------
                ---------------------------------------------
            ---------------------------------------------
            nas-id: cisco-ips <defaulted> local-fallback: enabled
 <defaulted> console-authentication: local
 default: radius-and-local default-user-role: viewer
 default: unspecified
        ---------------------------------------------
```

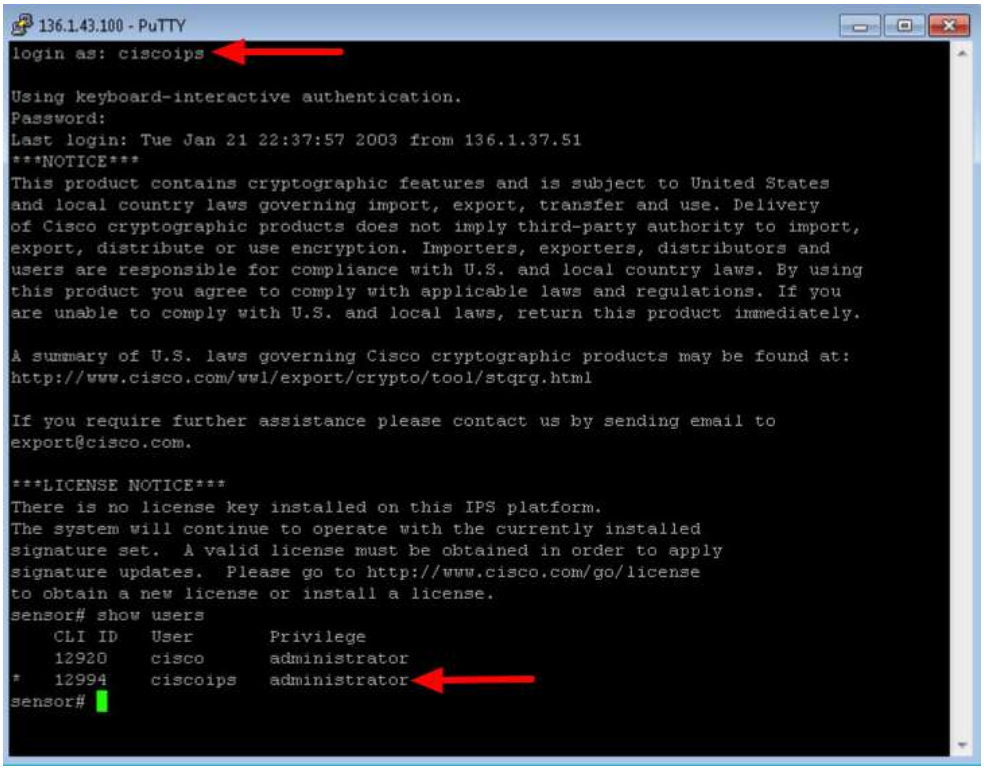
https://t.me/learningnets

```
    -----------------------------------------------
```

Verify the configured password restrictions:

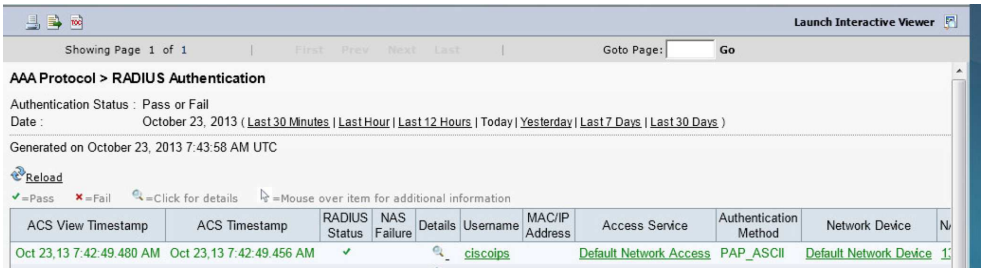
```
sensor(config-aut)# show settings
 attemptLimit: 5
 default: 0
   password-strength
   ------------------------------------------------ size: 6-20
 default: 8-64
      digits-min: 0 <defaulted>
      uppercase-min: 0 <defaulted>
      lowercase-min: 0 <defaulted>
      other-min: 0 <defaulted>
      number-old-passwords: 0 <defaulted>
   ------------------------------------------------
   permit-packet-logging: true <defaulted> cli-inactivity-timeout:
 10 default: 0
```

From **TEST-PC-B**, using putty, initiate an SSH connection to IPS and authenticate using the RADIUS user **ciscoips**:

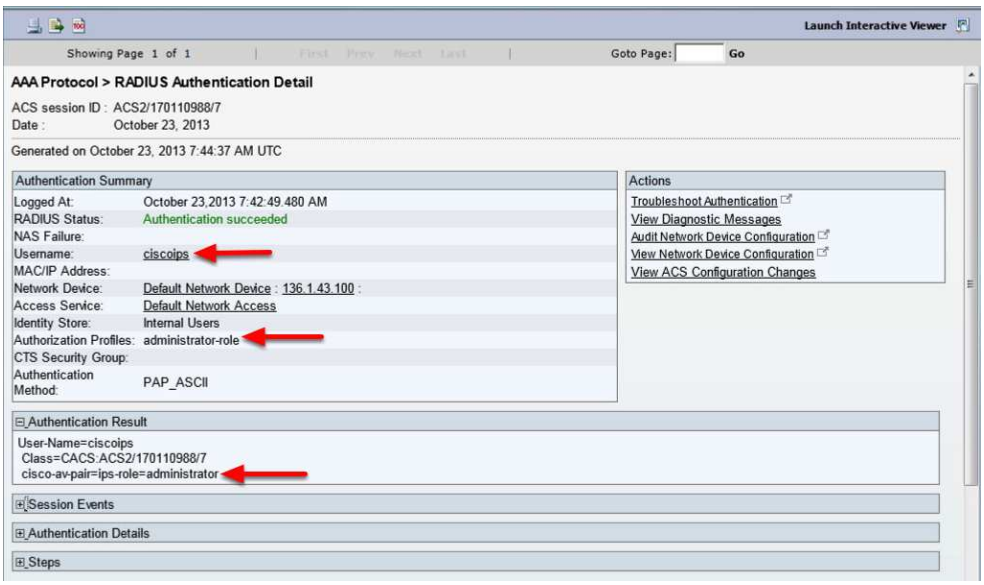


Verify the RADIUS authentication logs from ACS:

https://t.me/learningnets



Click **Details** for the authentication log and verify the matched authorization profile and assigned Cisco AV Pair value:



Verify that the console uses local authentication:

```
sensor login: cisco
Password:
Last login: Tue Jan 21 22:35:27 on ttyS0
***NOTICE***
This product contains cryptographic features and is subject to United States
and local country laws governing import, export, transfer and use. Delivery
of Cisco cryptographic products does not imply third-party authority to import,
export, distribute or use encryption. Importers, exporters, distributors and
users are responsible for compliance with U.S. and local country laws. By using
this product you agree to comply with applicable laws and regulations. If you
are unable to comply with U.S. and local laws, return this product immediately.

A summary of U.S. laws governing Cisco cryptographic products may be found at:
http://www.cisco.com/wwl/export/crypto/tool/stqrg.html

If you require further assistance please contact us by sending email to
export@cisco.com.

***LICENSE NOTICE***
There is no license key installed on this IPS platform.
```

https://t.me/learningnets

```
The system will continue to operate with the currently installed
signature set.  A valid license must be obtained in order to apply
signature updates.  Please go to http://www.cisco.com/go/license
to obtain a new license or install a license.sensor# show users
    CLI ID   User    Privilege       *   14426    cisco   administrator
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Inline Interface Pair

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9.
- Ensure that HTTP management access is allowed only from VLAN 37.
- Configure IPS to allow communication between R4 and R5 using its Gi0/0 and Gi0/1 interfaces according to the diagram.
- If an additional VLAN is required, use VLAN 54.
- Traffic between R4 and R5 should be inspected.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

IPS can be deployed using two methods: inline (in-band) or promiscuous mode (out-of-band). In inline mode traffic is actively passing through the IPS, whereas in promiscuous mode IPS receives a copy of the traffic through SPAN/RSPAN/ERSPAN. With inline mode, there are two possible deployments:

https://t.me/learningnets

- Inline interface-pair, where you must dedicate two physical interfaces and bind these into a logical interface.
- Inline VLAN-pair, where you must use a single physical interface, create a logical sub-interface, and assign two VLANs for the sub-interface; several sub-interfaces can be configured for a physical interface.

With inline interface-pair, depending on whether both interfaces are attached to same switch, the switchports must be configured:

- If both ports are attached to same switch, these must be in access mode but in different VLANs.
- If ports are attached to different switches, these can be either access in the same VLAN or trunk with all or restricted VLANs allowed.

In inline interface-pair mode, IPS does not perform VLAN and BPDU translation, so it cannot change the VLAN tag. For this reason, both BPDUs and traffic must be received either untagged or tagged with the same VLAN on both interfaces.

In all scenarios, IPS also bridges control-plane packets such as BPDU or UDLD. It inspects both IPv4 and IPv6 traffic, but non-IP traffic (such as ARP) is forwarded through uninspected.

Configuration

SW3:

```
interface GigabitEthernet1/0/8
 switchport mode access
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

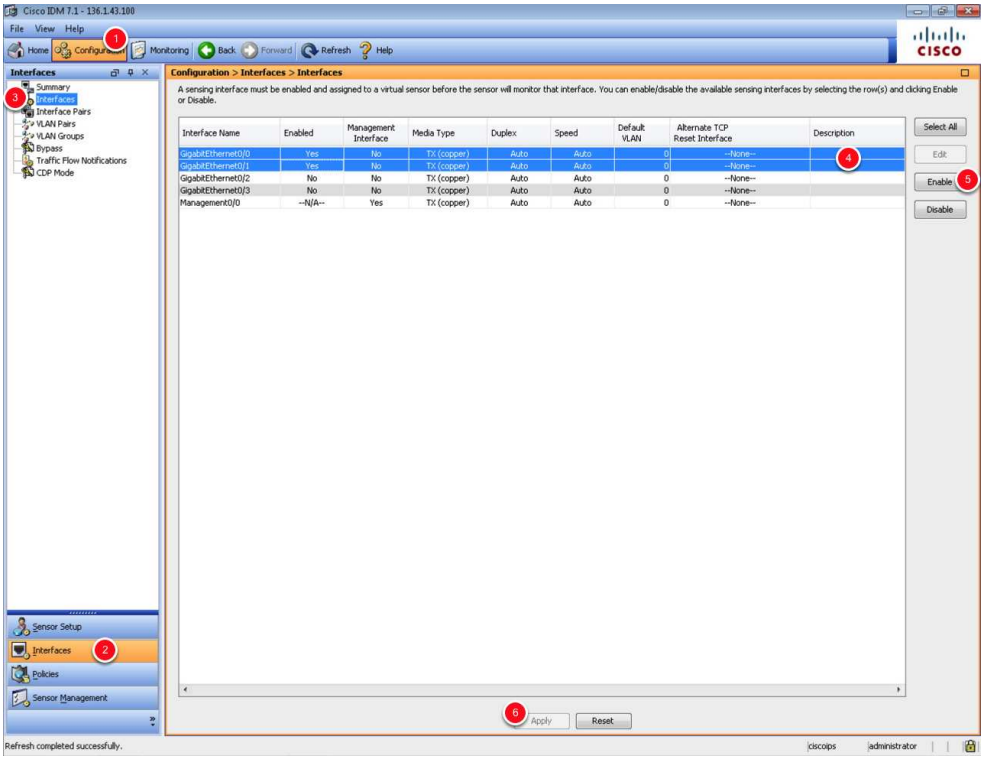
```
sensor(config)# service host
```

https://t.me/learningnets

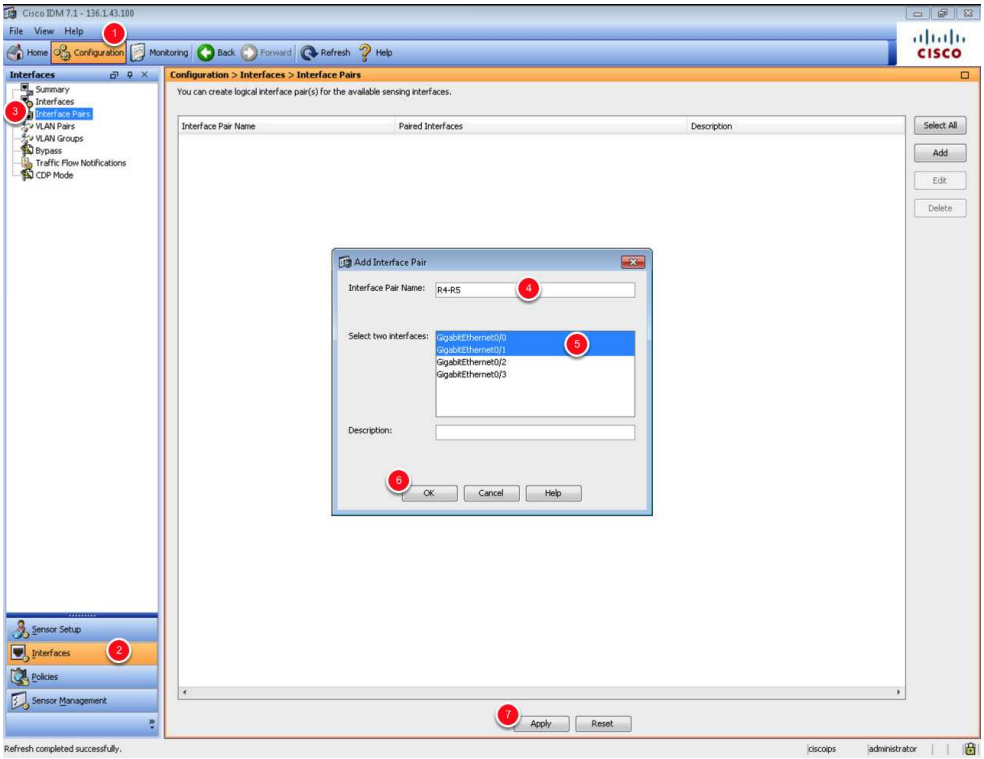
```
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor# configure terminal
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

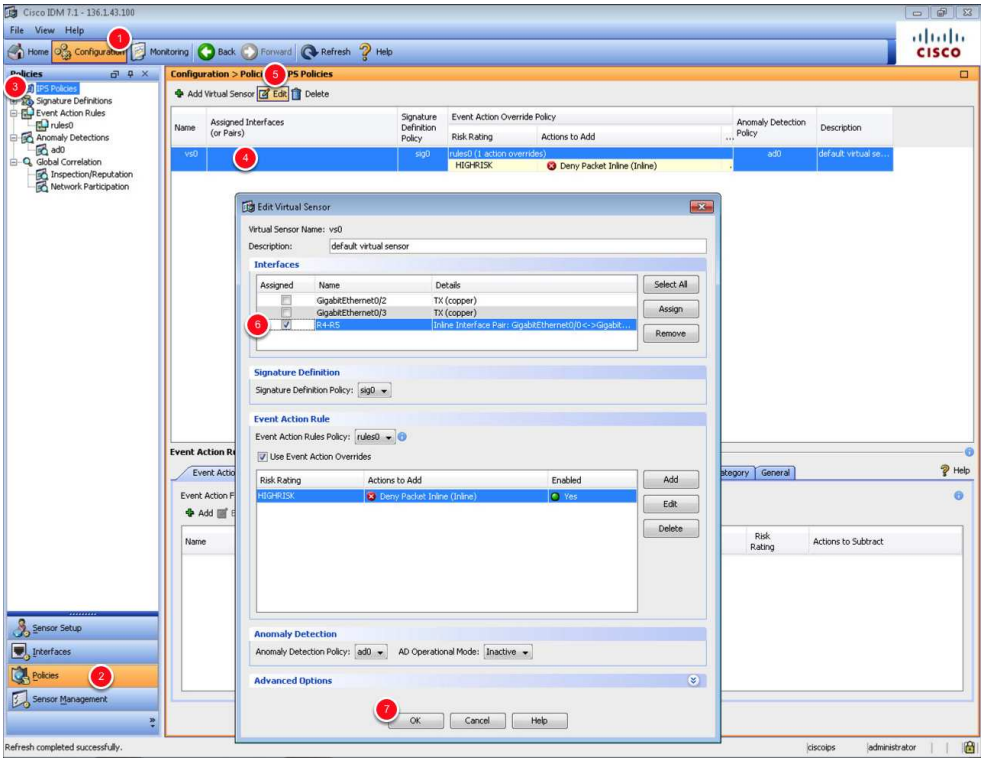


Create the inline interface-pair:



Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets



Verification

Verify interface state:

```
sensor# show interfaces brief
CC    Interface              Sensing State    Link    Inline Mode                                  Pair Status

GigabitEthernet0/0    Enabled          Up      Paired with interface GigabitEthernet0/1    Up

*     Management0/0          Disabled         Up
GigabitEthernet0/1    Enabled          Up      Paired with interface GigabitEthernet0/0    Up

      GigabitEthernet0/2     Disabled         Down    Unpaired                                     N/A
      GigabitEthernet0/3     Disabled         Down    Unpaired                                     N/A
```

Verify that IPS interfaces send and receive traffic:

```
sensor# show interfaces gigabitEthernet0/0 | i Packets
Total Packets Received = 36166
   Total Multicast Packets Received = 30912
   Total Broadcast Packets Received = 0
   Total Jumbo Packets Received = 0
   Total Undersize Packets Received = 0 Total Packets Transmitted = 14289
   Total Multicast Packets Transmitted = 9063
   Total Broadcast Packets Transmitted = 0
   Total Jumbo Packets Transmitted = 0
   Total Undersize Packets Transmitted = 0
```

https://t.me/learningnets

```
!
!sensor# show interfaces gigabitEthernet0/1 | i Packets
Total Packets Received = 15458
   Total Multicast Packets Received = 10173
   Total Broadcast Packets Received = 0
   Total Jumbo Packets Received = 0
   Total Undersize Packets Received = 0 Total Packets Transmitted = 35043

   Total Multicast Packets Transmitted = 29842
   Total Broadcast Packets Transmitted = 0
   Total Jumbo Packets Transmitted = 0
   Total Undersize Packets Transmitted = 0
```

Confirm EIGRP adjacency between R4 and R5 and test IP connectivity using ping:

```
R4#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(111)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   136.1.45.5              Fa0/0             13 11:29:16    3   200  0  10
0   136.1.43.9              Fa0/1.43          10 3d03h       5   200  0  14
!
!R4#ping 136.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.45.5, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Verify that R4 and R5 interfaces pointing to SW3 are not in the same VLAN, so traffic passes through the IPS:

https://t.me/learningnets

```
SW3#show spanning-tree interface gigabitEthernet1/0/8


Vlan                Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
VLAN0045            Desg FWD 4         128.8    P2p Edge

!
!SW3#show spanning-tree interface gigabitEthernet1/0/9


Vlan                Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
VLAN0054            Root FWD 4         128.9    P2p
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Inline VLAN Pair

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9.
- Ensure that HTTP management access is allowed only from VLAN 37.
- Configure IPS to allow communication between R4 and R5 using only its Gi0/0 interface.
- If an additional VLAN is required, use VLAN 54.
- Traffic between R4 and R5 should be inspected.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

This is very similar to the previous scenario, but it requires the use of a single physical interface, so inline VLAN-pair is being used. When functioning in this mode, because it bridges traffic between two VLANs, it needs to perform VLAN translation for both data-plane traffic and control-plane traffic (such as BPDUs), so it basically

https://t.me/learningnets

changes the VLAN tag when forwarding traffic.

Configuration

SW3:

```
interface GigabitEthernet1/0/8
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree portfast trunk
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

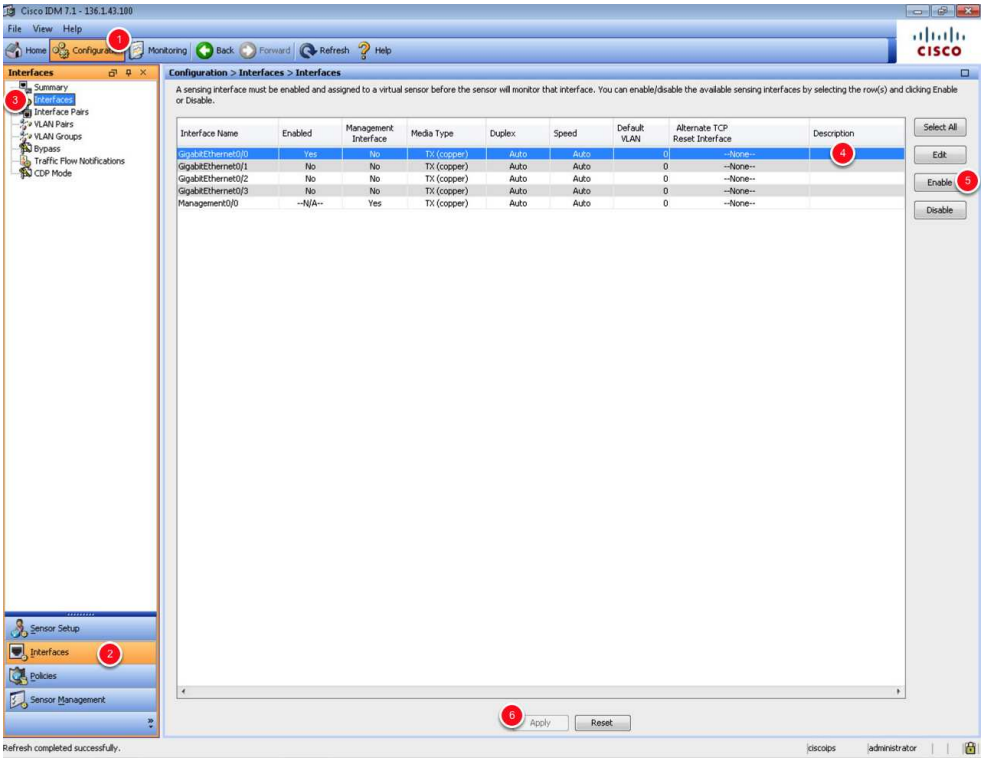
IPS:

```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor# configure terminal
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# subinterface-type inline-vlan-pair
sensor(config-int-phy-inl)# subinterface 1
sensor(config-int-phy-inl-sub)# vlan1 45
sensor(config-int-phy-inl-sub)# vlan2 54
sensor(config-int-phy-inl-sub)# exit
sensor(config-int-phy-inl)# exit
sensor(config-int-phy)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# physical-interface gigabitEthernet0/0 subinterface-number 1
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
```

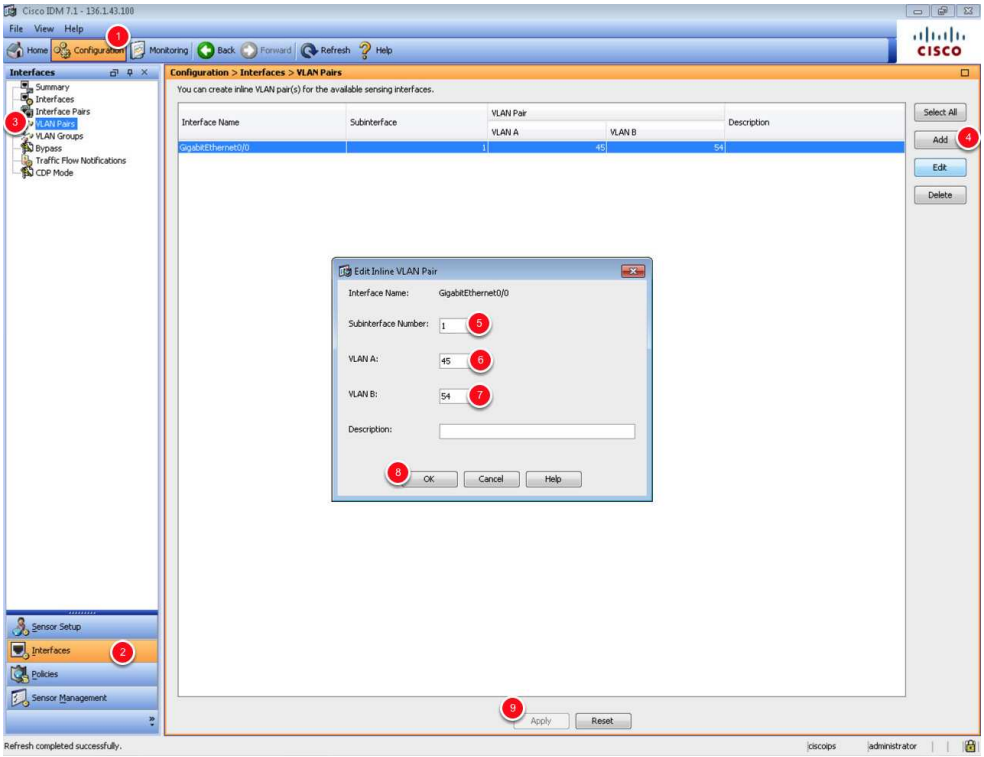
https://t.me/learningnets

```
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable Gi0/0 interface:

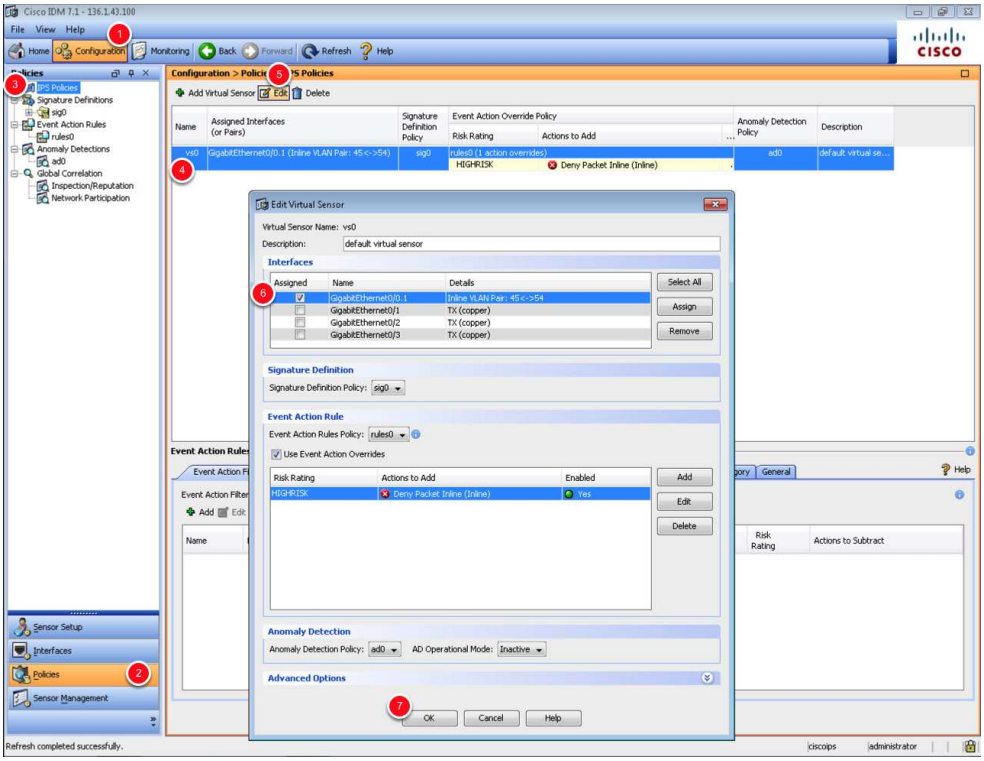


Configure the inline VLAN pair:



Assign the vlan-pair to the default virtual sensor:

https://t.me/learningnets



Verification

Verify interface state:

```
sensor# show interfaces brief
CC    Interface             Sensing State    Link    Inline Mode        Pair Status
  GigabitEthernet0/0    Enabled          Up      Inline-vlan-pair   N/A

*     Management0/0         Disabled         Up
      GigabitEthernet0/1    Disabled         Down    Unpaired           N/A
      GigabitEthernet0/2    Disabled         Down    Unpaired           N/A
      GigabitEthernet0/3    Disabled         Down    Unpaired           N/A
```

Verify that IPS interfaces send and receive traffic, and note the dropped packets as the trunk port from the switch did not restrict which VLANs are allowed:

```
sensor# show interfaces gigabitEthernet0/0 | i Packets
 Total Packets Received = 41748
    Total Multicast Packets Received = 36210
    Total Broadcast Packets Received = 10
    Total Jumbo Packets Received = 0
    Total Undersize Packets Received = 0 Total Packets Transmitted = 15289
    Total Multicast Packets Transmitted = 9890
    Total Broadcast Packets Transmitted = 0
    Total Jumbo Packets Transmitted = 0
    Total Undersize Packets Transmitted = 0
```

https://t.me/learningnets

```
          Total Packets Received On This Vlan = 479
          Total Packets Transmitted On This Vlan = 166
          Total Packets Received On This Vlan = 166
          Total Packets Transmitted On This Vlan = 479
Dropped Packets From Vlans Not Mapped To Subinterfaces = 4043
```

Confirm EIGRP adjacency between R4 and R5 and test IP connectivity using ping:

```
R4#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(111)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   136.1.45.5              Fa0/0             13 00:12:04    3   200  0  13
0   136.1.43.9              Fa0/1.43          13 3d03h       4   200  0  19
!
!R4#ping 136.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.45.5, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Promiscuous Mode

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Configure SW3 to send VLAN 37 traffic to the IPS. Traffic should not be received by SW1. Avoid traffic duplication and associate TCP resets with VLAN 37.
- Configure SW4 to send Gi1/0/7 traffic to the IPS and allow tagged inbound TCP resets.
- Enable signature 2000 for VLAN 37 traffic and signature 2004 for Gi1/0/7 traffic.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command **`erase current-config`** .

Overview

IPS can work in promiscuous mode and therefore receive a copy of the traffic through SPAN/RSPAN/ERSPAN sessions configured on the switch:

https://t.me/learningnets

- SPAN is used when both the sources for which traffic needs to be inspected and the IPS are attached to the same switch.
- RSPAN is used when the sources for which traffic needs to be inspected and the IPS are attached to different switches, but these switches are connected through trunk ports; a dedicated VLAN is used to carry traffic from the switch that has traffic sources attached to the switch that has the IPS attached.
- ERSPAN is used when the sources for which traffic needs to be inspected and the IPS are attached to different switches, but these switches are separated by layer 3 links and traffic must be routed; a GRE tunnel is used to carry traffic from the switch that has traffic sources attached to the switch that has the IPS attached.

Switch models available in the CCIE lab exam support only SPAN and RSPAN. For a SPAN session, a SPAN source and a SPAN destination must be configured:

- The SPAN source can be a local interface, a local VLAN, or a RSPAN VLAN; configured with the command `monitor session <number> source [interface <intf>|vlan <vlan_number>|remote vlan <vlan_number>] .`
- The SPAN destination can be an RSPAN VLAN or a local interface; configured with the command `monitor session <number> destination [remote vlan <vlan_number>|interface <intf>] .`

A SPAN session cannot use both interface and VLAN as source, whereas a SPAN destination cannot be used as a SPAN source. By default, both tx and rx traffic is copied for a SPAN session. When the SPAN source is a VLAN, if both tx and rx traffic is copied, it will result in traffic being duplicated, so it is common to specify only rx or tx. The VLAN tag is not preserved by default for monitored traffic, but based on the SPAN/RSPAN configuration, different options are available:

- If the SPAN source is an access interface, by default the VLAN tag is stripped out. To preserve the VLAN tag, use the `encapsulation dot1q` option when configuring the SPAN destination.
- If the SPAN source is a trunk interface, by default the VLAN tag is stripped out. To preserve the VLAN tag, use the `encapsulation dot1q` or `encapsulation replicate` option when configuring the SPAN destination; the difference is that the second option also copies layer 2 traffic such as STP, CDP, VTP, DTP, and PAgP.
- If the SPAN source is a VLAN, by default the VLAN tag is stripped out. To preserve the VLAN tag, use the `encapsulation dot1q` or `encapsulation replicate` option when configuring the SPAN destination.
- For RSPAN, if the `source remote vlan` option is used when defining the source, the VLAN tag is stripped out and there is no option to preserve it.

https://t.me/learningnets

- For RSPAN, if the `source vlan` option is used when defining the source, the behavior is the same as using a SPAN with a source VLAN, so by default the VLAN tag is stripped out but it can be preserved with one of the two options.

By default, no ingress traffic is allowed on a SPAN/RSPAN destination port, but this may be required to allow TCP resets. If IPS needs to send TCP resets, it will send it back on the interface that monitored traffic was received on (optionally, a different interface for sending TCP resets can be configured) and the VLAN tag is preserved. Therefore, if packets were received untagged, TCP resets will be sent untagged, and if packets were received tagged, TCP resets will be send tagged. To allow ingress traffic, several options are available when configuring the SPAN destination port:

- `ingress isl` will allow ISL encapsulated packets, but this is not in the scope of the CCIE lab exam.
- `ingress dot1q vlan <vlan_number>` will allow 802.1q tagged packets for all VLANs, and untagged packets will be associated to the specified VLAN.
- `ingress untagged <vlan_number>` or `ingress vlan <vlan_number>` will allow only untagged packets and associate it to the specified VLAN.

MAC address learning is disabled for VLANs configured as RSPAN, and ports configured as access in an RSPAN VLAN will be suspended. SPAN destination ports do not run STP and appear in the UP/DOWN state to be easily identified; the only requirement for SPAN destination ports is to be enabled.

VLAN groups were introduced in IPS 6.x code. It allows for partitioning the monitoring interface into logical groups, where each group is a collection of one or more VLANs. This implies that the traffic received by the IPS appliance should be tagged with VLAN headers. The benefit of using VLAN groups is that you may associate the group with a subinterface and later assign the subinterface to a virtual sensor. With the ability to use multiple virtual-sensors, VLAN groups allow for effective virtualization, by partitioning traffic streams in logical groups. Cisco IPS supports 802.1q encapsulation for VLAN groups, and all untagged frames or frames with ISL encapsulation are mapped to a special VLAN group called **unassigned**; this special group could be assigned to a subinterface as well. VLAN groups can be created for traffic received from both a SPAN/RSPAN session and an inline interface pair, but not for traffic received from an inline VLAN pair.

To create a VLAN group:

- Enter the physical interface configuration mode using the commands `service interface` ** and ** `physical-interfaces <name> .`
- Make sure to enter the command `admin-state enabled` to bring the physical interface

https://t.me/learningnets

up.

- Enter the subinterface configuration mode by using the command `subinterface-type vlan-group` ; this will put you in the VLAN group configuration mode. Create logical subinterfaces and map VLANs using the command `subinterface <number>` followed by `vlans range <range> .`
- The command `vlans unassigned` can be used to map all VLANs not explicitly selected to a subinterface; this includes untagged frames, such as frames received on the native 802.1q VLAN and frames tagged with ISL encapsulation. By default, all such frames are reported as belonging to VLAN 0 in the IPS output; the default VLAN number can be changed by using the following commands:

```
    service interface
      physical-interfaces GigabitEthernet 0/0
        default-vlan <vlan_number>
```

Since version 6.0, you may partition the analysis engine into multiple virtual sensors, with the default sensor being **vs0**. The benefit of using virtual sensors is that you may implement multiple IPS policies in the same appliance, such as policies for pre-firewall and post-firewall traffic, which usually require separate IPSs. Every virtual sensor has its own policy: signature definition, event action rules, and anomaly detection. Every sensor may be associated with a group of physical interfaces and/or logical subinterfaces. The commands to map the VLAN group subinterface to a virtual sensor are:

```
service analysis-engine
  virtual-sensor vs0
    physical-interface GigabitEthernet0/0 subinterface-number1
  virtual-sensor vs1
    physical-interface GigabitEthernet0/0 subinterface-number 2
```

Configuration

```
SW3:
vlan 666
 remote-span
!
monitor session 1 source vlan 37 rx
monitor session 1 destination remote vlan 666
!
interface GigabitEthernet1/0/20
 switchport trunk allowed vlan remove 666
```

https://t.me/learningnets

SW4:

```
vlan 666
 remote-span
!
monitor session 1 source remote vlan 666
monitor session 1 destination interface Gi1/0/9 ingress untagged vlan 37
monitor session 2 source interface Gi1/0/7
monitor session 2 destination interface Gi1/0/10 encapsulation dot1q ingress dot1q vlan 32
!
interface range GigabitEthernet1/0/9 - 10
 no shutdown
```

IPS:

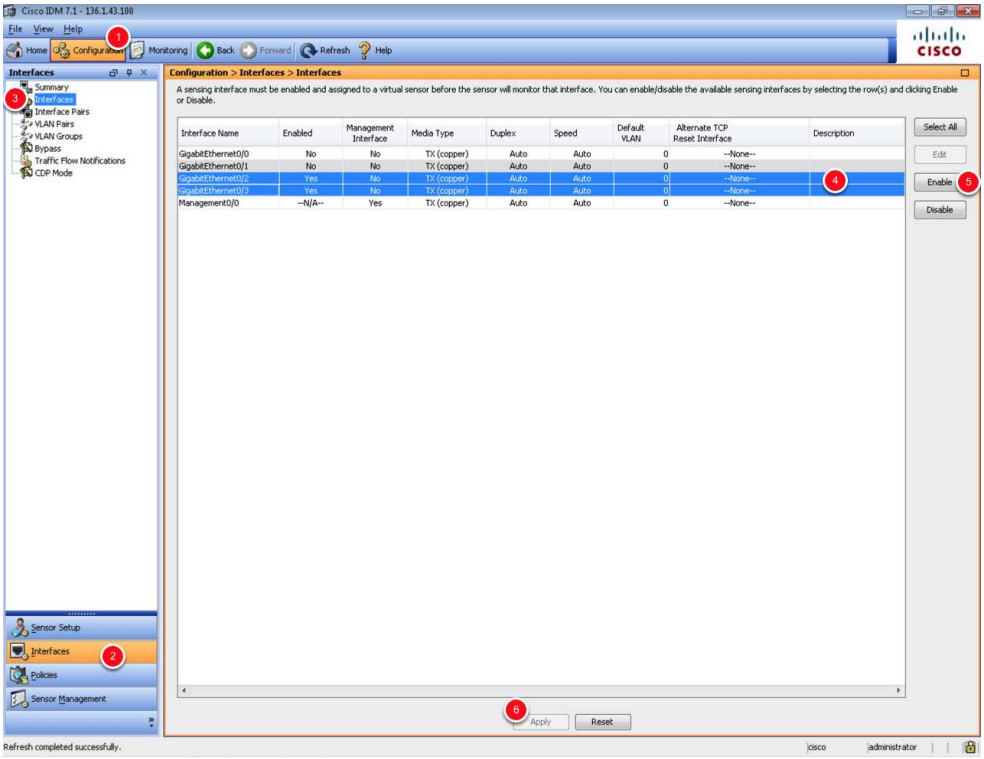
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor# configure terminal
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/2
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# subinterface-type vlan-group
sensor(config-int-phy-vla)# subinterface 1
sensor(config-int-phy-vla-sub)# vlans unassigned
sensor(config-int-phy-vla-sub)# exit
sensor(config-int-phy-vla)# exit
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/3
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# subinterface-type vlan-group
sensor(config-int-phy-vla)# subinterface 2
sensor(config-int-phy-vla-sub)# vlans range 32
sensor(config-int-phy-vla-sub)# exit
sensor(config-int-phy-vla)# exit
sensor(config-int-phy)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
```

https://t.me/learningnets

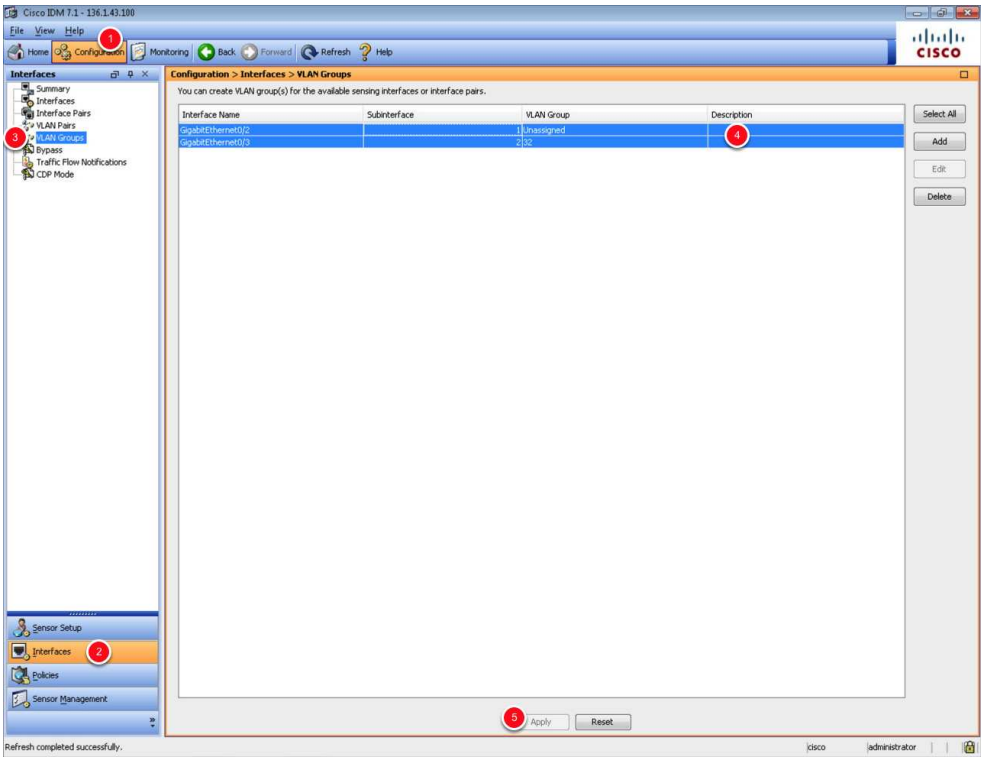
```
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 2000 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes?[yes]:
!
sensor(config)# service signature-definition sig1
sensor(config-sig)# signatures 2004 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# physical-interface gigabitEthernet0/2 subinterface-number 1
sensor(config-ana-vir)# exit
sensor(config-ana)# virtual-sensor vs1
sensor(config-ana-vir)# physical-interface gigabitEthernet0/3 subinterface-number 2
sensor(config-ana-vir)# signature-definition sig1
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable Gi0/2 and Gi0/3 interfaces:

https://t.me/learningnets

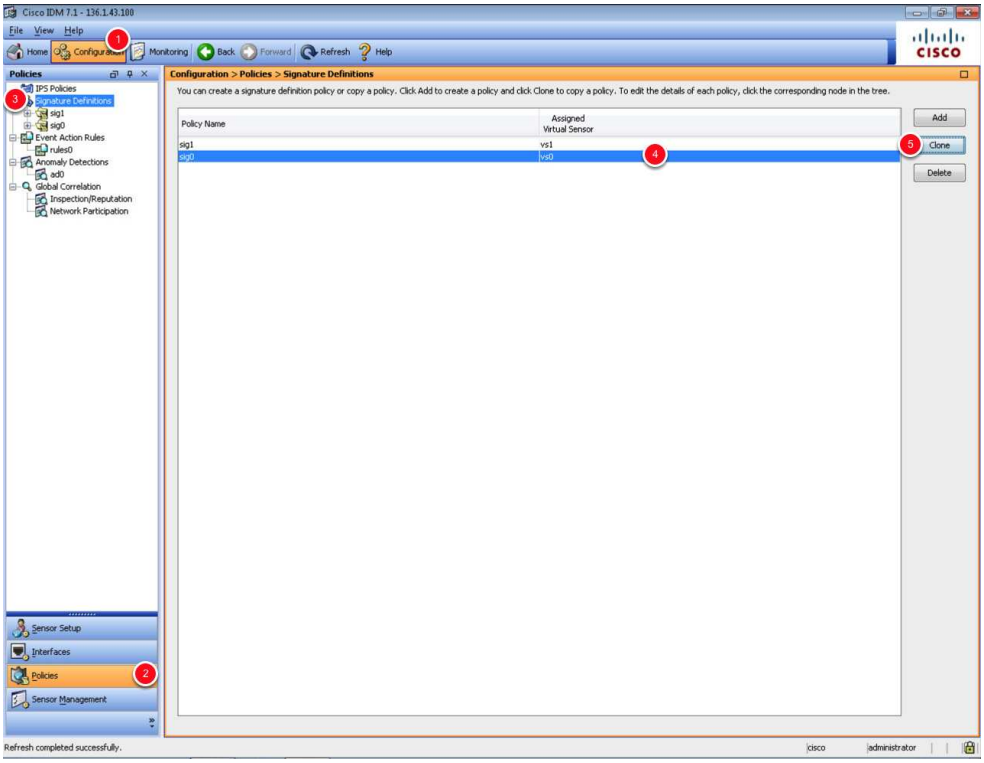


Configure the two VLAN groups, and assign correct VLANs and interfaces:

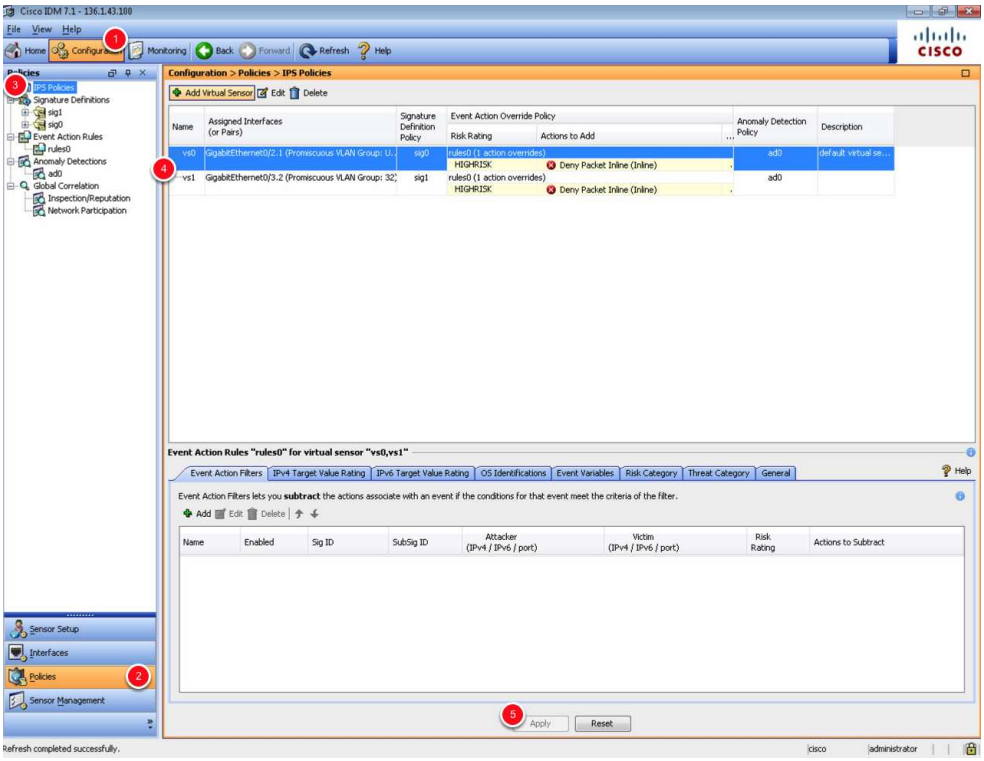


Clone the default signature definition of sig0:

https://t.me/learningnets

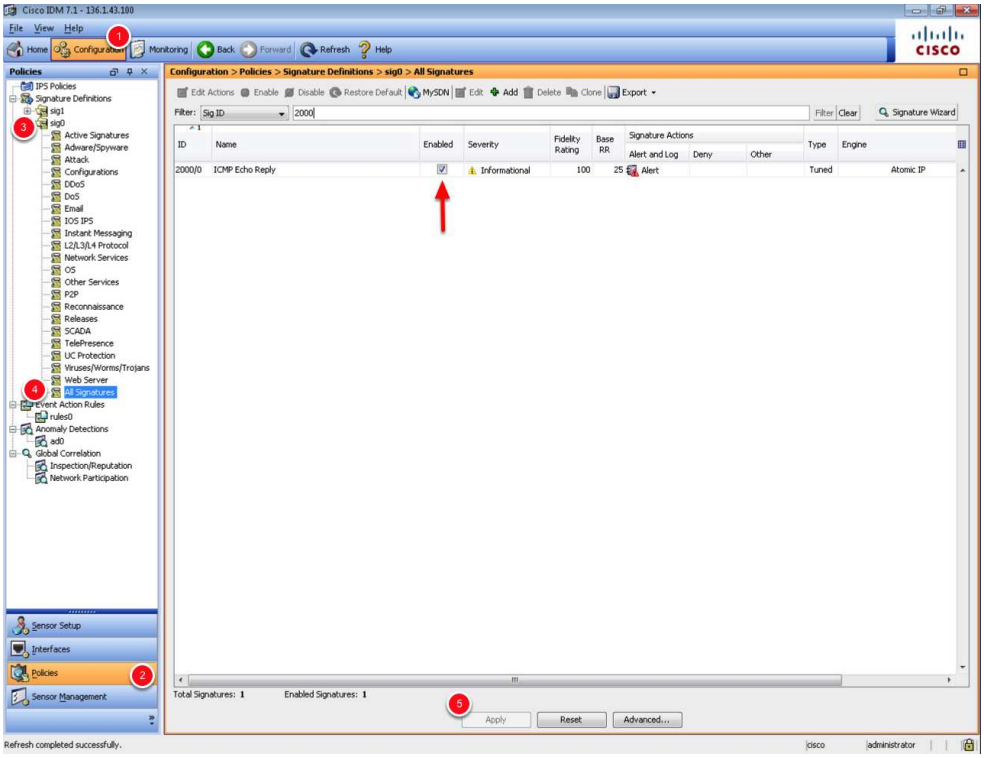


Add a new virtual sensor and assign the cloned signature definition to it:

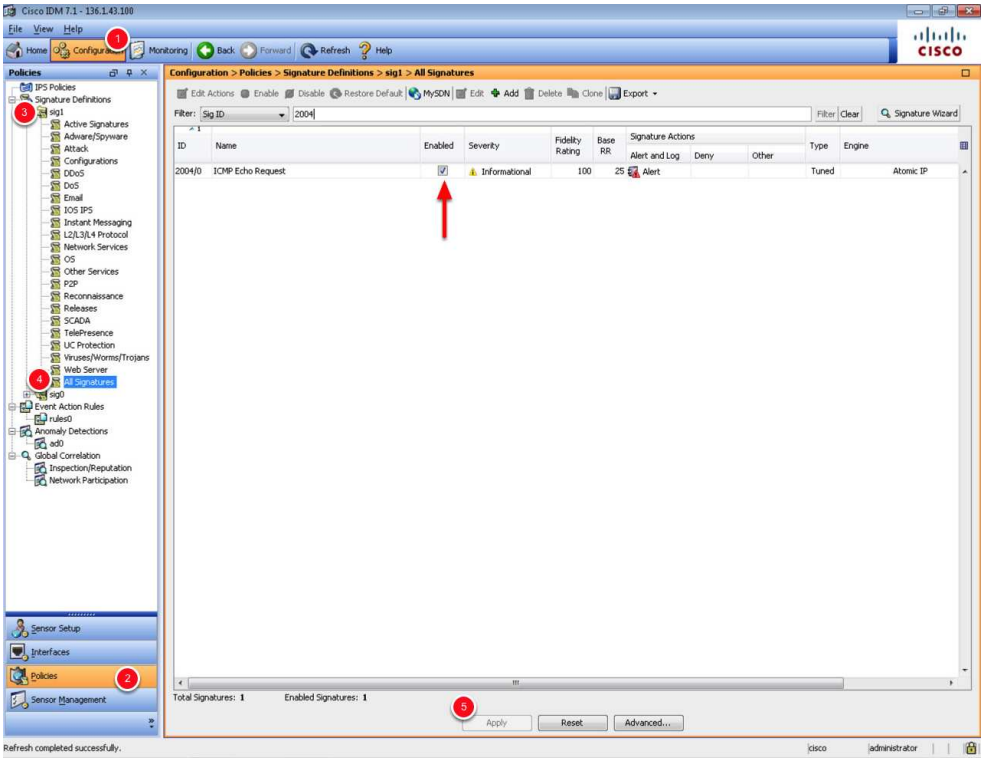


Enable signature 2000 for sig0:

https://t.me/learningnets



Enable signature 2004 for sig1:



Verification

https://t.me/learningnets

Verify that for VLAN 32 traffic only Echo Request signature is triggered; check that traffic is received tagged over the expected interface and processed by the expected virtual sensor:

```
R4#ping 172.16.2.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.100, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/12 ms
!
!sensor# show events

evIdsAlert: eventId=1376321302518003825 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1202
  time: 2003/01/30 05:43:10 2003/01/30 05:43:10 UTC
signature: description=ICMP Echo Request id=2004 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request interfaceGroup: vs1
vlan: 32
  participants:
    attacker:
      addr: locality=OUT 172.16.2.1
    target:
      addr: locality=OUT 172.16.2.100
      os: idSource=unknown relevance=relevant type=unknown
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 35 interface: ge0_3

  protocol: icmp
```

Verify that for VLAN 37 traffic only Echo Reply signature is triggered; check that traffic is received untagged over the expected interface and processed by the expected virtual sensor:

```
SW1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.37.51         0148.f8b3.2e25.32       Mar 12 1993 03:14 AM    Automatic
!
!SW1#ping 136.1.37.51
```

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.51, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!sensor# show events

evIdsAlert: eventId=1376321302518003830 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1202
  time: 2003/01/30 05:45:50 2003/01/30 05:45:50 UTC
signature: description=ICMP Echo Reply id=2000 created=20001127 type=other version=S666

    subsigId: 0
    sigDetails: ICMP Echo Reply interfaceGroup: vs0
vlan: 0
  participants:
    attacker:
      addr: locality=OUT 136.1.37.51
    target:
      addr: locality=OUT 136.1.37.9
      os: idSource=unknown relevance=relevant type=unknown
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 35 interface: ge0_2

  protocol: icmp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Custom Signature

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Create a new signature that triggers on the string **% Bad passwords** found in telnet session responses.
  - The signature should fire an alarm on every occurrence of the string.
- Configure IPS to allow communication between R4 and R5 using its Gi0/0 and Gi0/1 interfaces according to the diagram.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

Except for the default signature set, IPS allows for creating custom signatures. It's important to understand the traffic pattern you are trying to match because there are multiple options that must be selected to have a match: signature engine, direction

https://t.me/learningnets

of the traffic, and engine-specific parameters. The engine choices include String TCP Engine, Atomic IP Engine, Service HTTP Engine, and other Engines; each of these engines would have different parameters because they relate to different aspects of how traffic is matched. For example, the HTTP engine allows for the specification of certain URI information, which can be specified with a regular expression, but this regex is of no use to an Atomic IP signature. When working with regular expressions, you can use the following meta characters:

Symbol	Meaning
?	Repeat 0 or 1 times
*	Repeat 0 or more times
+	Repeat 1 or more times
.	Any single character except new line (\n) or tab (\t)
[123]	Any character inside the brackets
[^abc]	Any character not listed in the brackets
[a-z]	Any character listed in the range
|	Or
\	Removes the special meaning of the following character.

As an example, to match on the word "**CCIE**" regardless of case, we would use the regular expression `[cC][cC][iI][eE]` . To match on something like a URL, we would need to escape the meta characters listed in the table. For example, `ccielab\.com` would match the url **ccielab.com**.

In this section, the IPS is watching for a string in a telnet session, specifically for the string **% Bad passwords**. In the telnet session, this will be sent by default from the

https://t.me/learningnets

device in response to three failed password attempts. Because this is traffic coming from the device to which we are trying to connect, the direction is **from-service**, and the traffic flow from that device will have a source port of 23. If we wanted to match on something that was being typed in on the device to which we are trying to connect, the direction would be **to-service**. One additional item to consider for this type of attack notification is the **swap-attacker-victim** option under the signature. Looking at the output of the show events on the IPS, you can see who is listed as the **attacker**; this is the device that is triggering the signature, but it is not the malicious device that we are trying to identify. By adding the **swap-attacker-victim** option, the IPS will swap the two and list the real attacker.

Configuration

SW3:

```
interface GigabitEthernet1/0/8
 switchport mode access
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

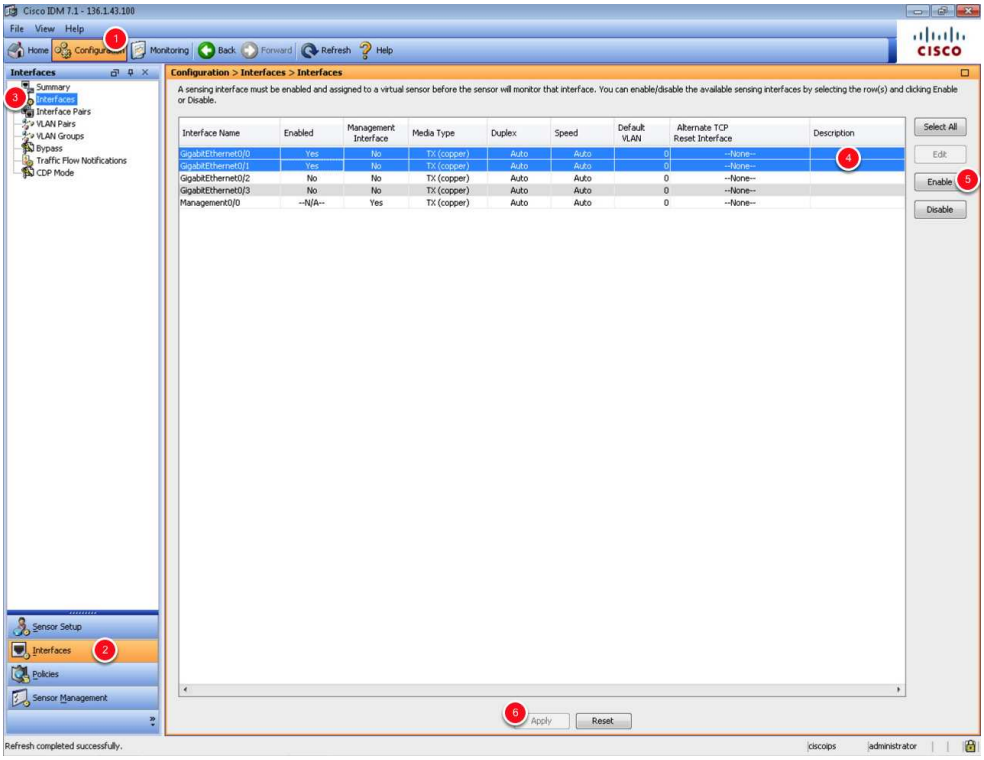
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 60100 0
sensor(config-sig-sig)# engine string-tcp
```

https://t.me/learningnets

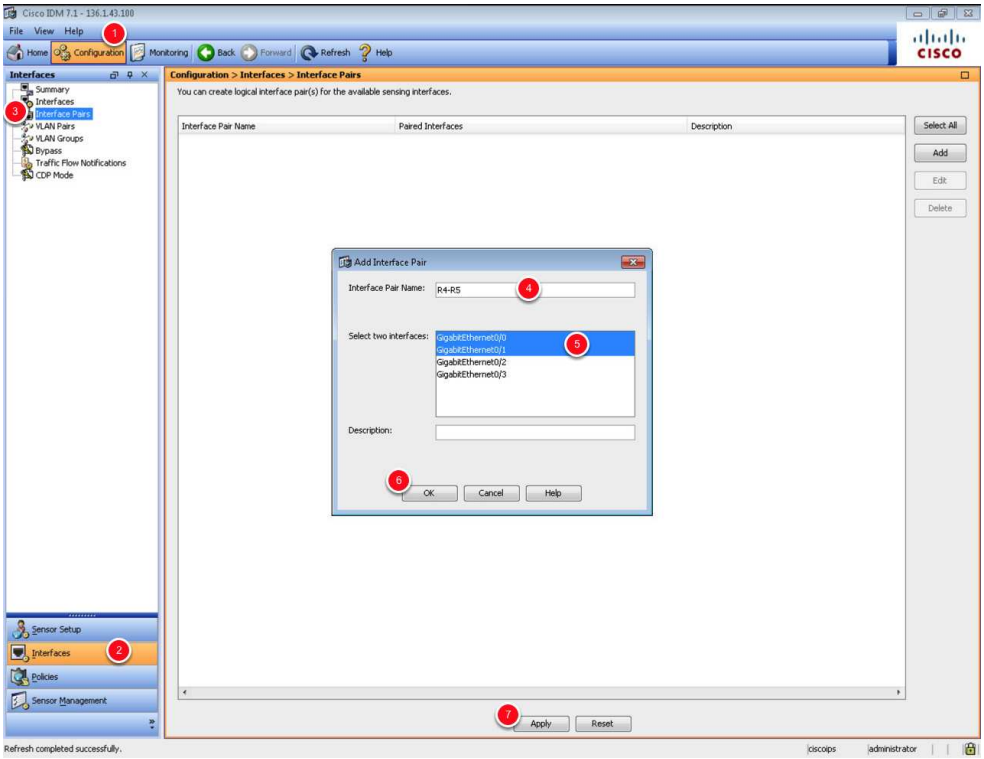
```
sensor(config-sig-sig-str)# service-ports 23
sensor(config-sig-sig-str)# direction from-service
sensor(config-sig-sig-str)# regex-string % Bad passwords
sensor(config-sig-sig-str)# exit
sensor(config-sig-sig)# alert-frequency
sensor(config-sig-sig-ale)# summary-mode fire-all
sensor(config-sig-sig-ale-fir)# exit
sensor(config-sig-sig-ale)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes:?[yes]: yes
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

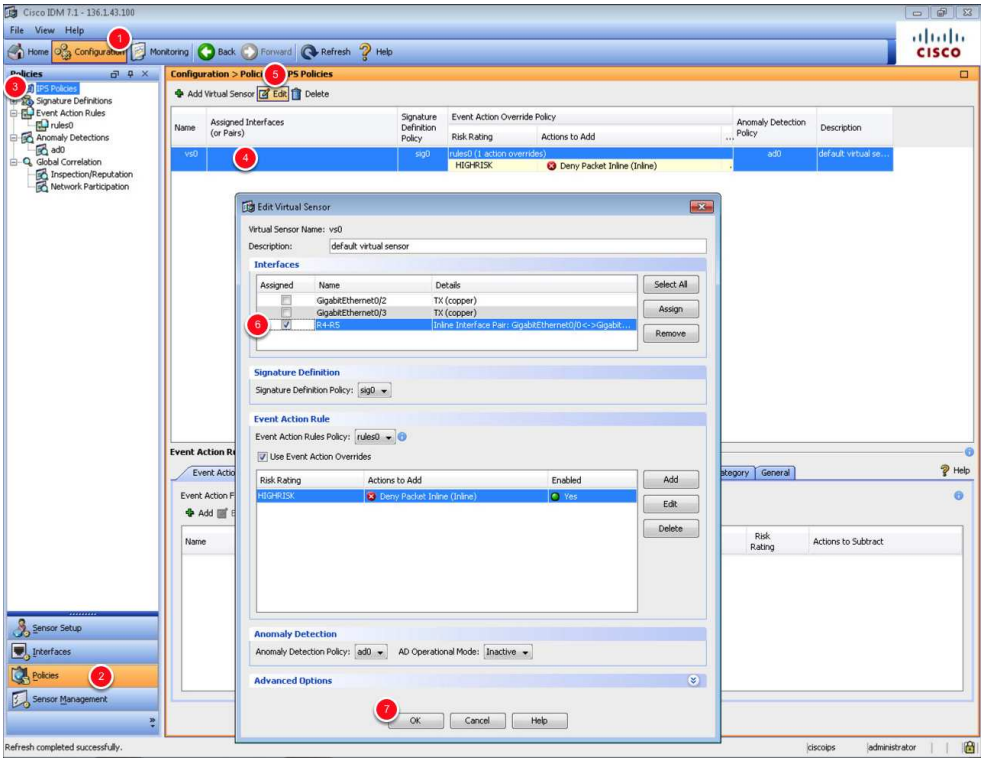


Create the inline interface-pair:

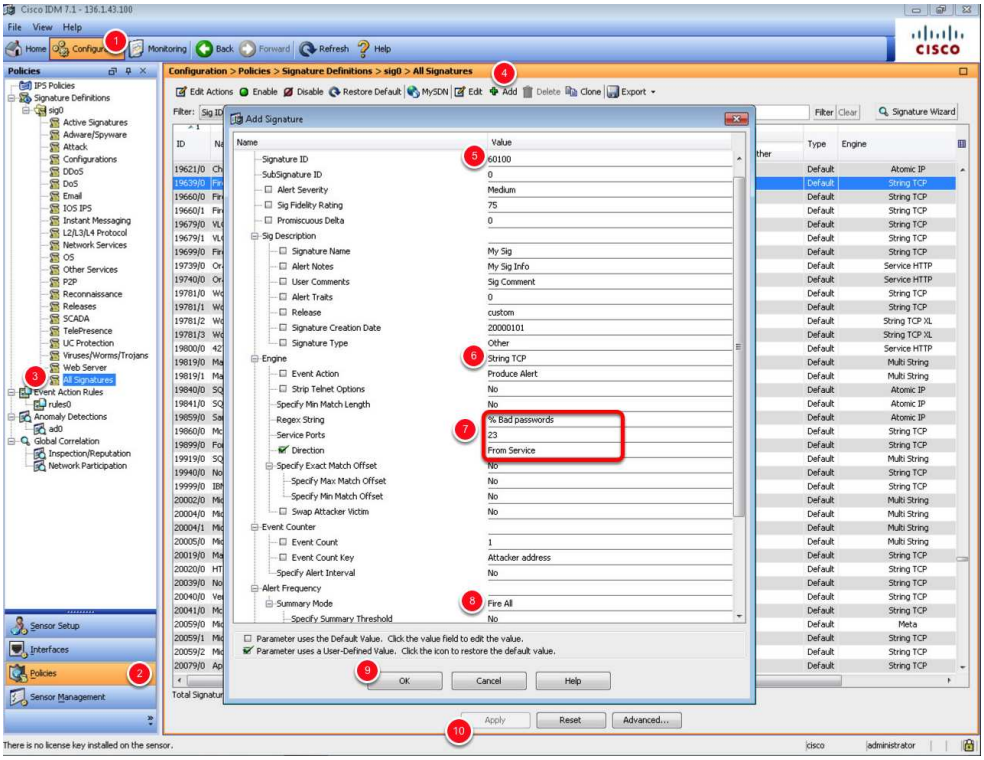


Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets



Add the custom signature:



Verification

Telnet from R4 to R5 and simulate an authentication failure to generate the message that will trigger the signature to fire:

https://t.me/learningnets

```
R4#telnet 136.1.45.5

Trying 136.1.45.5 ... Open



User Access Verification

Password:
Password:
Password:  % Bad passwords

[Connection to 136.1.45.5 closed by foreign host]
!
!sensor# show events

evIdsAlert: eventId=1376321302518003963 severity=medium vendor=Cisco alarmTraits=32768
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1202
  time: 2003/01/30 07:31:38 2003/01/30 07:31:38 UTC
signature: description=My Sig id=60100 created=20000101 type=other version=custom
    subsigId: 0
    sigDetails: My Sig Info
  interfaceGroup: vs0
  vlan: 0
  participants:  attacker:
addr: locality=OUT 136.1.45.5
      port: 23 target:
addr: locality=OUT 136.1.45.4
      port: 55316
      os: idSource=unknown relevance=relevant type=unknown
  context:
    fromAttacker:
000000  FF FB 01 FF FB 03 FF FD  18 FF FD 1F 0D 0A 0D 0A  ................
000010  55 73 65 72 20 41 63 63  65 73 73 20 56 65 72 69  User Access Veri
000020  66 69 63 61 74 69 6F 6E  0D 0A 0D 0A 50 61 73 73  fication....Pass
000030  77 6F 72 64 3A 20 FF FE  20 FF FD 21 FF FA 21 00  word: .. ..!..!.
000040  FF F0 FF FE 18 0D 0A 50  61 73 73 77 6F 72 64 3A  .......Password:
000050  20 0D 0A 50 61 73 73 77  6F 72 64 3A 20 0D 0A 25   ..Password: ..%
000060  20 42 61 64 20 70 61 73  73 77 6F 72 64 73 0D 0A   Bad passwords..
    fromTarget:
000000  FF FD 03 FF FB 20 FF FB  1F FF FB 21 FF FD 01 FF  ..... .....!....
000010  FC 18 FF FA 1F 00 50 00  18 FF F0 FF FC 20 74 65  ......P...... te
000020  6C 6E 65 74 0D 0A 74 65  6C 6E 65 74 0D 0A 74 65  lnet..telnet..te
000030  6C 6E 65 74 0D 0A                                lnet..
```

https://t.me/learningnets

```
riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 66
threatRatingValue: 66 interface: ge0_1


protocol: tcp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Inline Blocking

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Create a new signature that triggers on the string **cisco** found in telnet sessions.
  - The signature should fire only if 3 consecutive events occur during a 3-minute interval span.
  - The attacker should be denied access for 10 minutes.
- Configure IPS to allow communication between R4 and R5 using its Gi0/0 and Gi0/1 interfaces according to the diagram.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

A default IPS signature can be in the state of **retired**, which means it is not loaded in the memory, so it does not consume any resources and it is not matched by any

https://t.me/learningnets

traffic being inspected by the IPS. A non-retired signature is loaded into memory but can be in the state of **disabled**, which means it is not matched by any traffic being inspected by the IPS. When a signature is fired, there are one or more possible actions the IPS can be configured to perform, depending on the functional mode (inline or promiscuous). When deployed in inline mode, the most important actions are:

- **deny attacker inline**: blocks all traffic from the host identified as the attacker for a default period of 3600 seconds, which can be modified.
- **deny attacker service pair inline**: blocks only traffic for a specific service (TCP or UDP port) from the host identified as the attacker for a default period of 3600 seconds, which can be modified.
- **deny attacker victim pair inline**: blocks only traffic from the host identified as the attacker to the host identified as the victim for a default period of 3600 seconds, which can be modified.
- **deny connection inline**: blocks all traffic for the specific connection (source-destination socket pair) and sends a one-way TCP reset only to the victim to free up resources.
- **deny packet inline**: drops the offending packet and behaves the same as **deny connection inline**.
- **modify packet inline**: changes the packet inline and sends it to the destination, but it can be used only for specific use cases, such as clearing a TCP flag or for TCP segmentation and IP fragmentation re-ordering.

Although the above actions have a direct impact on the traffic, additional non-disruptive actions can be performed for both inline and promiscuous modes:

- **produce alert**: writes the event to the IPS internal Event store, which is a 32MB non-configurable circular buffer in-memory.
- **produce verbose alert**: writes the event to the IPS internal Event store but also includes a packet dump.
- **log attacker packets**: writes the event to the IPS internal Event store and starts IP logging for attacker packets.
- **log victim packets**: writes the event to the IPS internal Event store and starts IP logging for victim packets.
- **log attacker/victim pair packets**: writes the event to the IPS internal Event store and starts IP logging for victim and attacker packets.
- **request SNMP trap**: writes the event to the IPS internal Event store and triggers an SNMP trap, given that SNMP is correctly configured.

https://t.me/learningnets

Configuration

SW3:

```
interface GigabitEthernet1/0/8
 switchport mode access
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

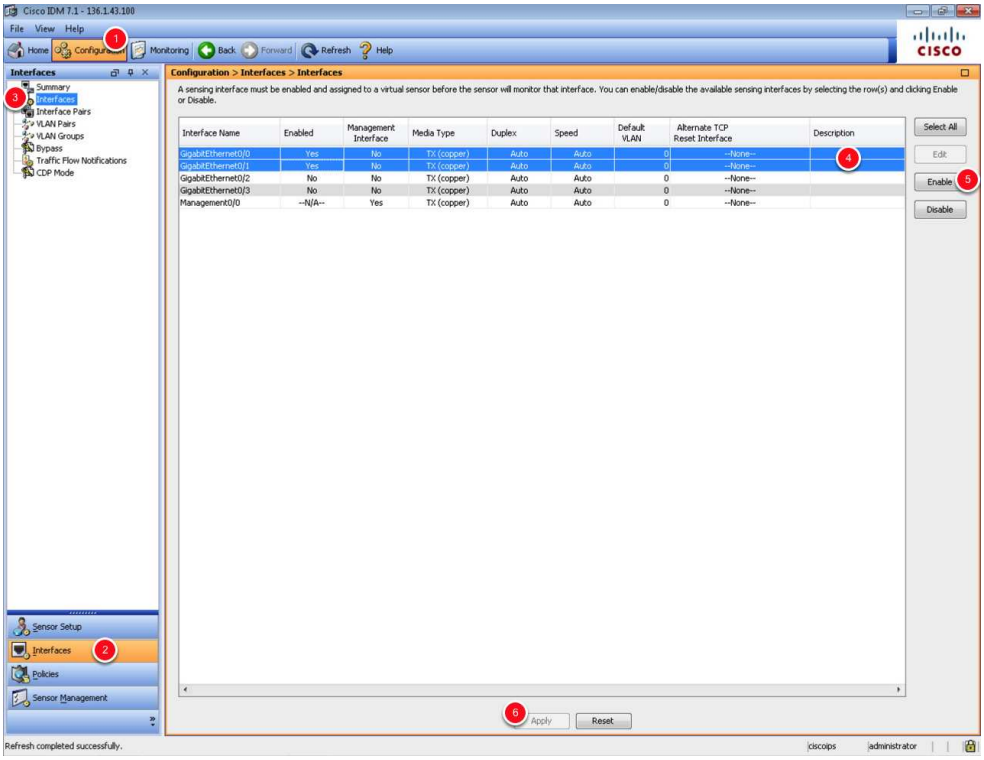
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 60100 0
sensor(config-sig-sig)# engine string-tcp
sensor(config-sig-sig-str)# event-action deny-attacker-inline
sensor(config-sig-sig-str)# service-ports 23
sensor(config-sig-sig-str)# direction to-service
sensor(config-sig-sig-str)# regex-string cisco
sensor(config-sig-sig-str)# exit
sensor(config-sig-sig)# event-counter
sensor(config-sig-sig-eve)# event-count 3
sensor(config-sig-sig-eve)# specify-alert-interval yes
sensor(config-sig-sig-eve-yes)# alert-interval 180
sensor(config-sig-sig-eve-yes)# exit
sensor(config-sig-sig-eve)# exit
```

https://t.me/learningnets

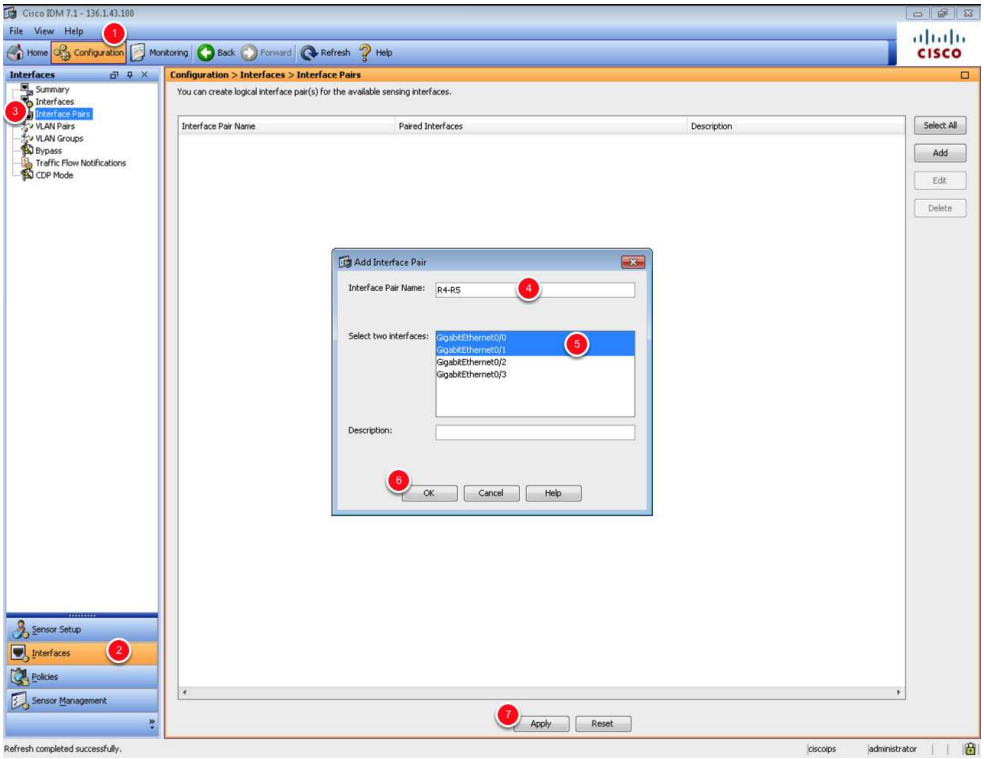
```
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes:?[yes]: yes
!
sensor(config)# service event-action-rules rules0
sensor(config-eve)# general
sensor(config-eve-gen)# global-deny-timeout 600
sensor(config-eve-gen)# exit
sensor(config-eve)# exit
Apply Changes?[yes]:
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

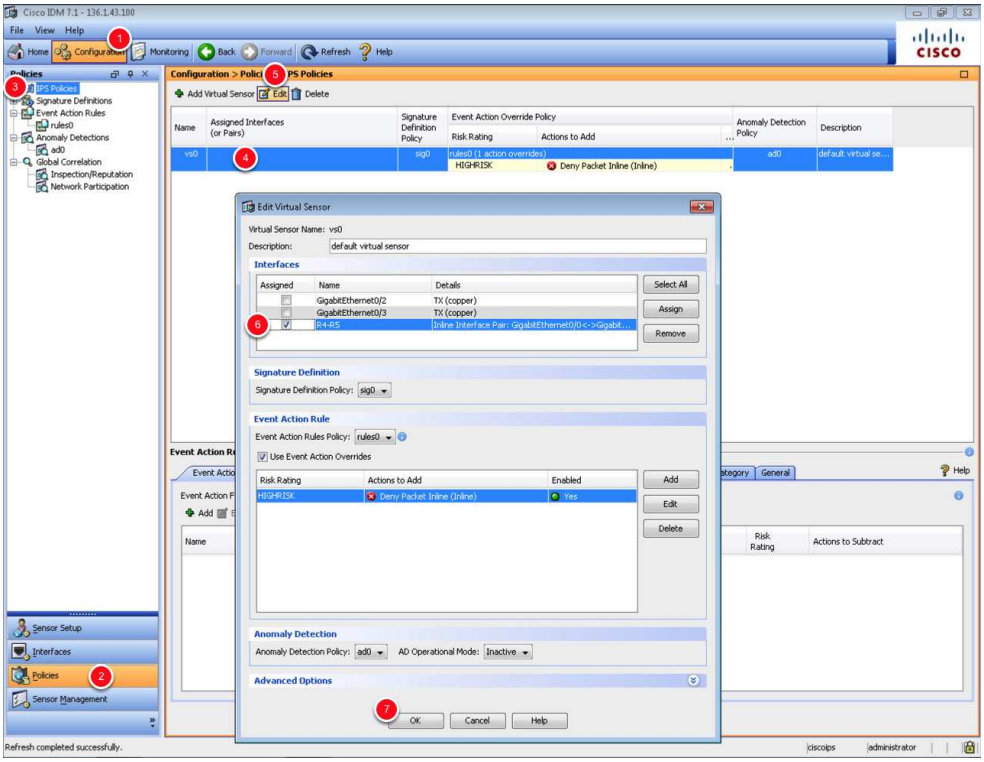


Create the inline interface-pair:

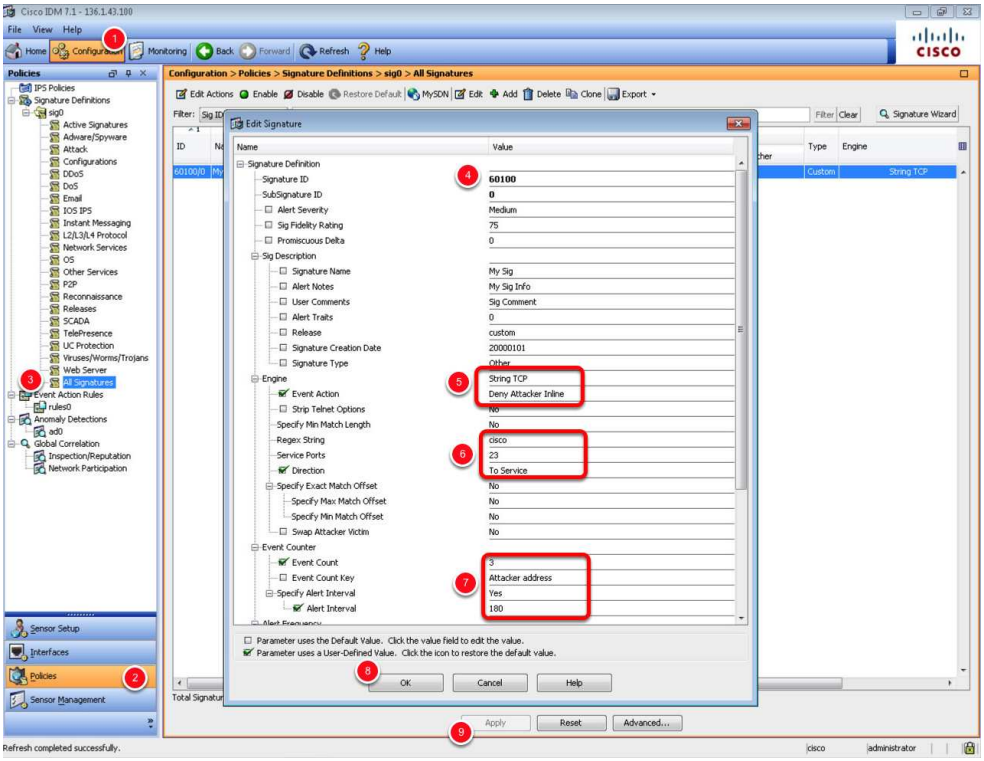


Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets

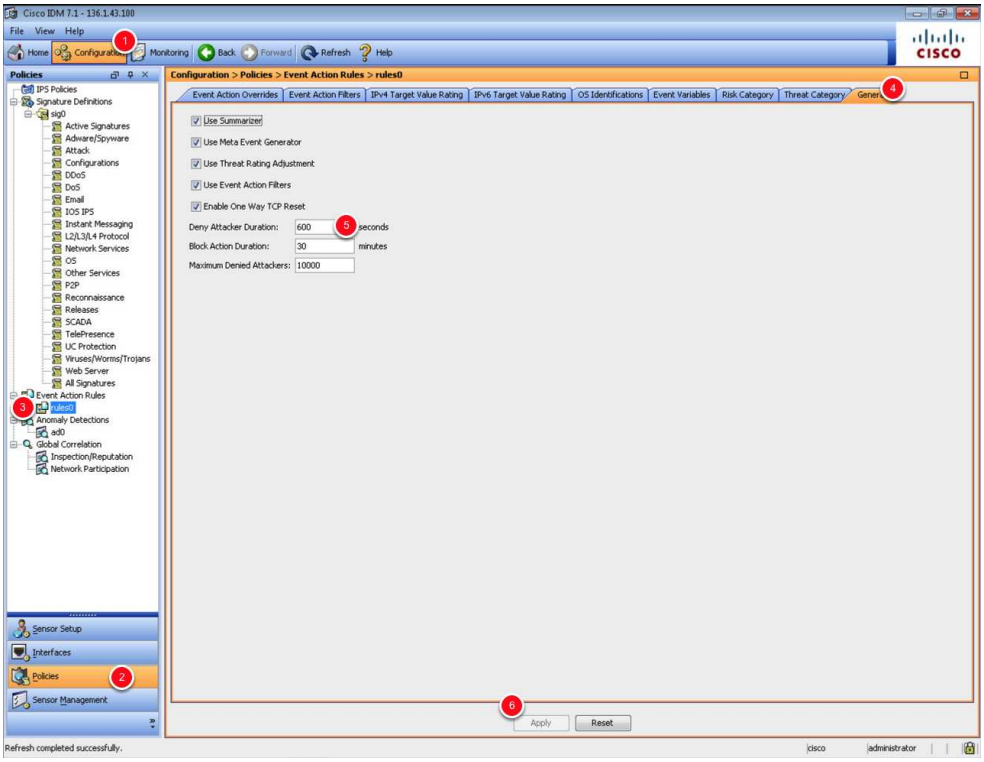


Add the custom signature:



Modify the default value for denied attackers from 3600 seconds to 600 seconds:

https://t.me/learningnets



Verification

Initiate a telnet session from R4 toward R5, authenticate using the password **cisco**, and type the string **cisco** two more times to fire the signature:

```
R4#telnet 136.1.45.5
Trying 136.1.45.5 ... Open


User Access Verification

Password: R5>cisco
Translating "cisco"

% Bad IP address or host name
% Unknown command or computer name, or unable to find computer addressR5>cisc
!
!sensor# show events

evStatus: eventId=1376328122518000129 vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 01:58:04 2003/01/31 01:58:04 UTC denyAttackerStarted:
```

https://t.me/learningnets

```
description: denyAttackerStarted for address: 136.1.45.4
 address: 136.1.45.4
```

Verify that R4 appears in the attacker list and traffic from R4 is dropped; remove R4 from the list to re-establish connectivity:

```
sensor# show statistics denied-attackers
Statistics for Virtual Sensor vs0
   Denied Attackers and hit count for each.
      136.1.45.4 = 94
   Denied Attackers with percent denied and hit count for each.
         Attacker Address    Victim Address    Port    Protocol    Requested Percentage    Actual Percentage    Hit Count

136.1.45.4                                                100                  100                  94             false
!
!R4#ping 136.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.45.5, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!sensor# clear denied-attackers
Warning: Executing this command will delete all addresses from the list of attackers currently being denied by the s
Continue with clear? [yes]: yes
!
!R4#ping 136.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.45.5, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS TCP Reset

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Create a new signature that triggers on the string **cisco** found in telnet sessions.
  - The signature should create an event and issue a bidirectional TCP reset.
  - Monitor R4's Fa0/1 interface; IPS alerts should include the VLAN number.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

When IPS is deployed in promiscuous mode, the **deny** action discussed in the inline mode task does not apply and cannot be performed by the IPS, but the action of **reset tcp connection** can be configured. This will result in a TCP reset being sent to both the attacker and the victim; this action is also valid for the inline mode. IPS will basically spoof the source/destination MAC addresses and source/destination IP

https://t.me/learningnets

addresses and build a TCP RST packet, while also keeping the VLAN tag if the packet was received as tagged; TCP RST is by default sent back on the interface the packet that triggered the action was received on, but (optionally) another monitoring interface can be specified (not the management interface).

Configuration

SW4:

```
monitor session 1 source interface Gi1/0/1
monitor session 1 destination interface Gi1/0/9 encapsulation dot1q ingress dot1q vlan 43
!
interface GigabitEthernet1/0/9
 no shutdown
```

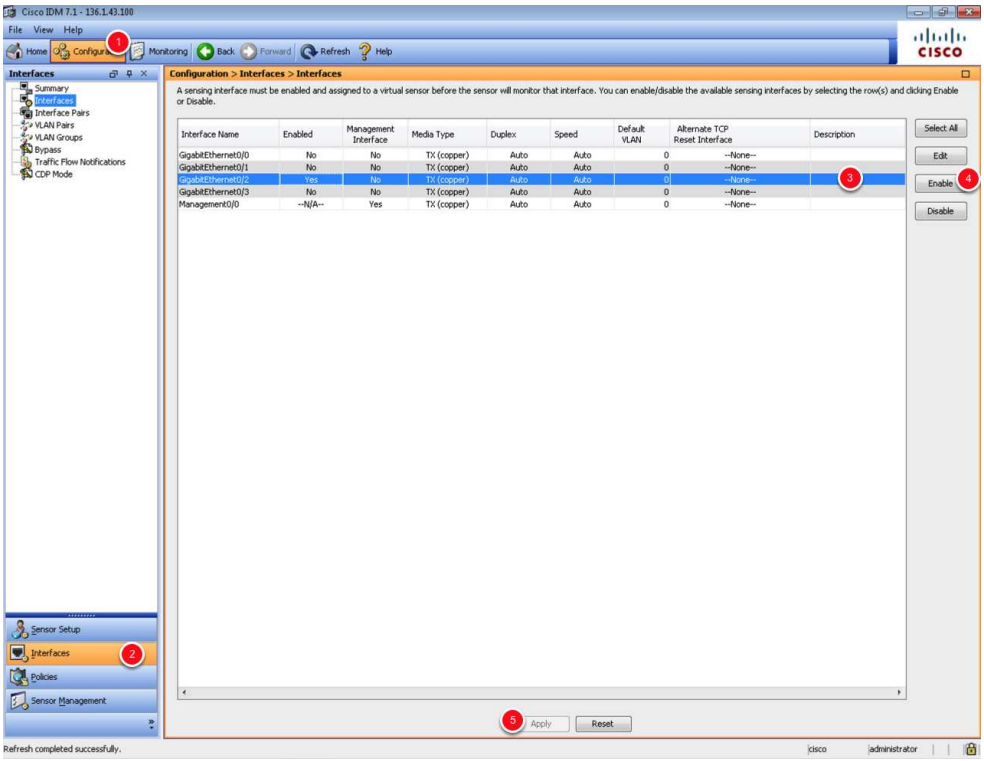
IPS:

```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 60100 0
sensor(config-sig-sig)# engine string-tcp
sensor(config-sig-sig-str)# event-action reset-tcp-connection|produce-alert
sensor(config-sig-sig-str)# service-ports 23
sensor(config-sig-sig-str)# direction to-service
sensor(config-sig-sig-str)# regex-string cisco
sensor(config-sig-sig-str)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes:?[yes]: yes
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/2
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
```

https://t.me/learningnets

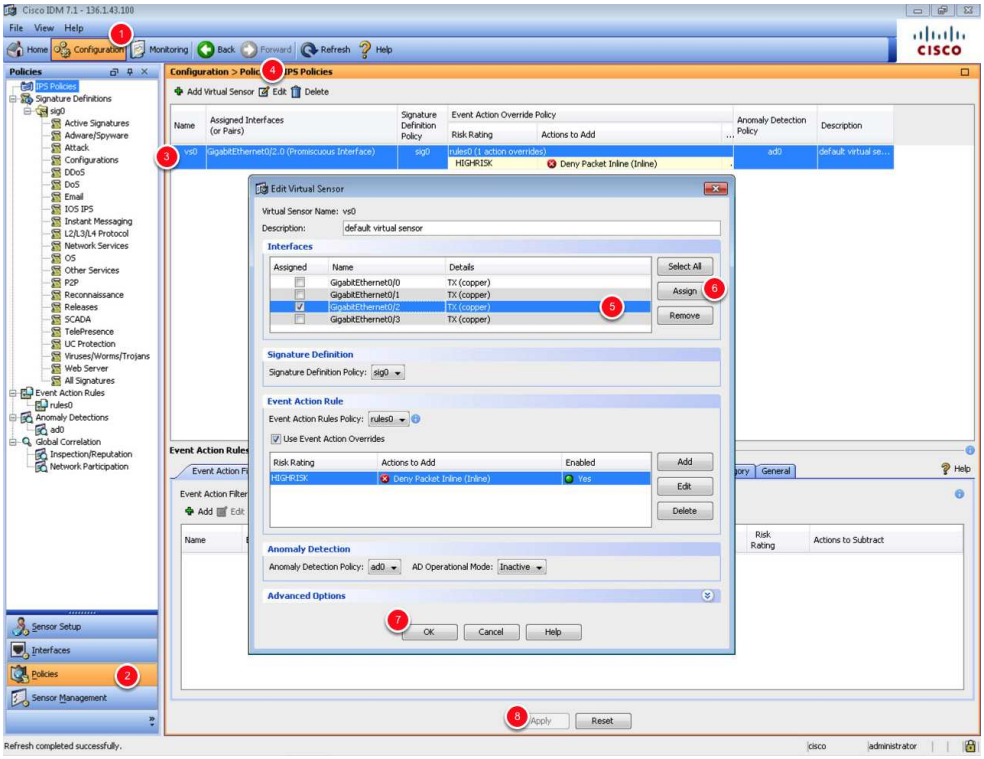
```
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# physical-interface gigabitEthernet0/2
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, enable Gi0/2 interface:

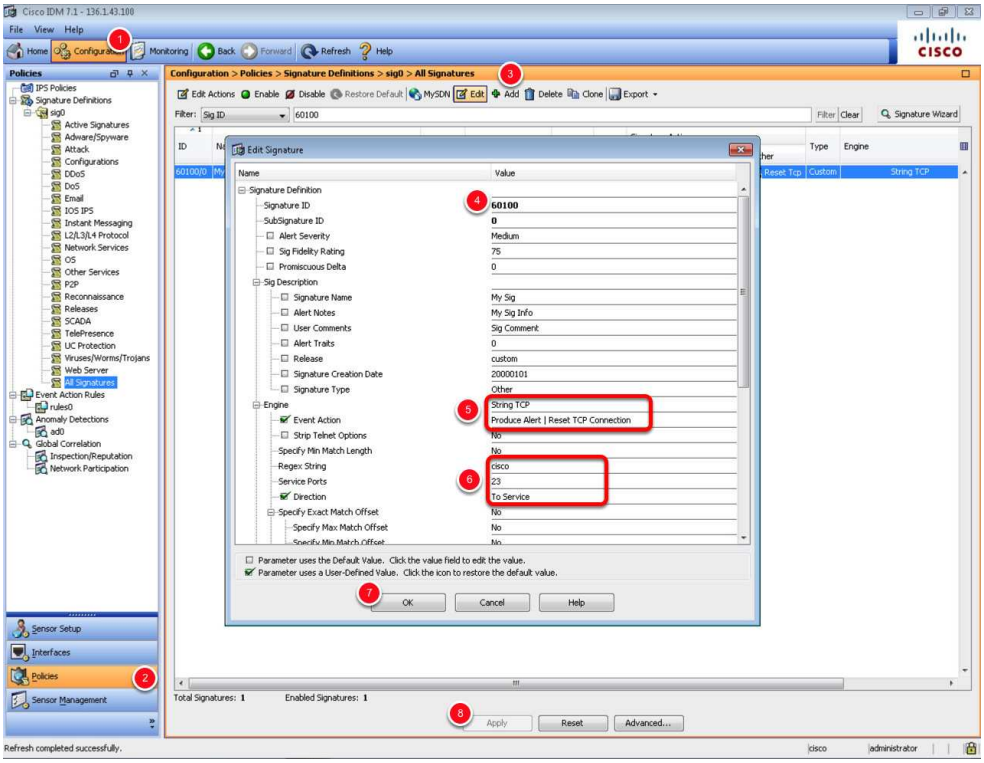


Assign the interface to the virtual sensor:

https://t.me/learningnets



Configure the signature:



Verification

Initiate a telnet session from R4 to SW1 and authenticate with the string of **cisco** to trigger the signature; IPS should reset the connection:

https://t.me/learningnets

```
R4#telnet 136.1.43.9

Trying 136.1.43.9 ... Open



User Access Verification

Password: [Connection to 136.1.43.9 closed by foreign host]
!
!sensor# show events

evIdsAlert: eventId=1376328122518000461 severity=medium vendor=Cisco alarmTraits=32768
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 15:55:56 2003/01/31 15:55:56 UTC
  signature: description=My Sig id=60100 created=20000101 type=other version=custom
    subsigId: 0
    sigDetails: My Sig Info
  interfaceGroup: vs0 vlan: 43
  participants:
    attacker:
      addr: locality=OUT 136.1.43.4
      port: 51691
    target:
      addr: locality=OUT 136.1.43.9
      port: 23
      os: idSource=unknown relevance=relevant type=unknown actions:
resetTcpFlowSent: true
  context:
    fromTarget:
000000  FF FB 01 FF FB 03 FF FD  18 FF FD 1F 0D 0A 0D 0A  ................
000010  55 73 65 72 20 41 63 63  65 73 73 20 56 65 72 69  User Access Veri
000020  66 69 63 61 74 69 6F 6E  0D 0A 0D 0A 50 61 73 73  fication....Pass
000030  77 6F 72 64 3A 20 FF FE  20 FF FE 21 FF FE 18 0D  word: .. ..!....
000040  0A 50 61 73 73 77 6F 72  64 3A 20                 .Password:
    fromAttacker:
000000  FF FD 03 FF FB 20 FF FB  1F FF FB 21 FF FD 01 FF  ..... .....!....
000010  FC 18 FF FA 1F 00 50 00  18 FF F0 FF FC 20 FF FC  ......P...... ..
000020  21 0D 0A 63 69 73 63 6F                           !..cisco
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 66
  threatRatingValue: 46 interface: ge0_2
```

https://t.me/learningnets

```
    protocol: tcp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Blocking and Rate-Limiting

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Monitor R4's Fa0/1 interface; alerts should include the VLAN number.
- Configure the IPS sensor to block any fragmented ICMP Echo Request packets inbound on R6's Fa0/0 interface.
- As soon as there is an ICMP flood detected, configure R4 to alert and rate-limit the respective flow to 25% of the Fa0/0 interface bandwidth.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

In promiscuous mode, the IPS has some additional possible actions (except the TCP RST) to actively prevent a network attack: block the attacker, block the connection, or rate-limit the attacker. These actions are performed by the system

https://t.me/learningnets

component known as Attack Response Controller or ARC. Generally, the IPS communicates with external devices such as routers, firewalls, and switches to implement the blocking action. Communications are performed using telnet or SSH as transport protocols from the management interface of the IPS; IPS emulates a user session and configures devices for attack response actions. To add network devices under IPS control, you must perform the following steps:

- Configure Network Devices.
- Add Device Login Profiles to the IPS.
- Add Blocking Devices to the IPS.
- Configure Blocking Interfaces.
- Configure Signature Actions respectively.

Configure Network Devices

During this step, you configure the routers, firewalls, or catalyst switch with the necessary settings to allow remote access. The IPS could use either Telnet or SSH protocols to access the remote devices. You may need to create a separate user in the devices for the IPS appliance. Based on your device configuration, you may need just one password to access the device remotely and drop into enable mode instantly, or have a separate enable password required.

Add Device Login Profiles to the IPS

This is where you create objects representing device access credentials to the IPS. To add a new profile, go to **Configuration > Sensor Management > Blocking > Device Login Profiles** and click **Add**. Every profile has a login name and password and the enable password. In some situations, you may not need the login name or password, such as when the router is configured to authenticate just using the line password.

Add Blocking Devices to the IPS

At this step, you add the objects representing remote devices to the IPS. Navigate to **Configuration > Sensor Management > Blocking > Blocking Devices** to access this feature. Every blocking device has an IP address, a Device Login Profile associated, a Device Type (Cisco router, PIX/ASA, Catalyst 6500), a list of capabilities (block/rate-limit), and the protocol to use for communicating with the device (Telnet or SSH). Naturally, the protocol selected and the device type should match the remote device settings. Some devices may not support the rate-limiting capability, because this one is typically associated with Cisco routers only. Notice the special setting called Sensor’s NAT Interface. When the IPS sensor applies

https://t.me/learningnets

blocking to a device, it makes sure its own IP address is not affected by the blocking configuration. The IPS automatically discovers its own IP address, but in case when the IPS control interface is behind a NAT device, this is impossible. In such situation you should configure the post-NAT IP address for the particular device manually, and this IP address will always remain unblocked.

If you are using SSH for device communication, make sure you requested and accepted the SSH RSA server keys. Without this step, the IPS will reject the remote host identity. In real life, you are supposed to check the public key received and make sure it belongs to the correct device. This step could be done using IDM by going to **Configuration > Sensor Management > SSH > Known Host Keys**, clicking **Add**, and entering the remote device IP address to retrieve the public key.

Configure Blocking Interfaces

With Cisco routers, the IPS applies blocking configuration to an interface that you select at this step. On the contrary, blocking is global on devices such as PIX/ASA and those do not require blocking interfaces selected. When you navigate to **Configuration > Sensor Management > Blocking > Router Blocking Devices Interfaces**, you may add an interface, specify its name (no spaces), and associate it with a blocking device. Notice that you may specify the direction to apply blocking (inbound or outbound), which actually means the direction to apply an access-list or service policy. In addition to this, you may also specify the device-specific access-lists that should be merged with the access-list constructed by the IPS. For example, if you already had an access-list configured on the interface, you may want its rules to be prepended to the list applied by the IPS. In this case, it is the pre-block ACL. If you want the local access-list to be appended to the IPS ACL, configure the list name as a post-block ACL.

Configure Signature Actions

You may want to tune some of your signatures to apply the blocking or rate-limiting actions in addition to simply reporting an event. Finally, remember that the IPS appliance is not aware of the network topology. Thus, it will apply blocking or rate-limiting actions to all capable devices when responding to an event. Assuming this, you may want to adjust the number of network devices controlled by a given IPS to be in immediate proximity of the IPS controlled zone.

Configuration

SW4:

```
monitor session 1 source interface Gi1/0/1
monitor session 1 destination interface Gi1/0/9 encapsulation dot1q
```

https://t.me/learningnets

```
!
interface GigabitEthernet1/0/9
 no shutdown
```

IPS:

```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/2
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# physical-interface gigabitEthernet0/2
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
!
sensor(config)# service network-access
sensor(config-net)# user-profiles R4
sensor(config-net-use)# password
Enter password[]: cisco
Re-enter password: cisco
sensor(config-net-use)# enable-password
Enter enable-password[]: cisco
Re-enter enable-password: cisco
sensor(config-net-use)# exit
sensor(config-net)# user-profiles R6
sensor(config-net-use)# password
Enter password[]: cisco
Re-enter password: cisco
sensor(config-net-use)# enable-password
Enter enable-password[]: cisco
Re-enter enable-password: cisco
```

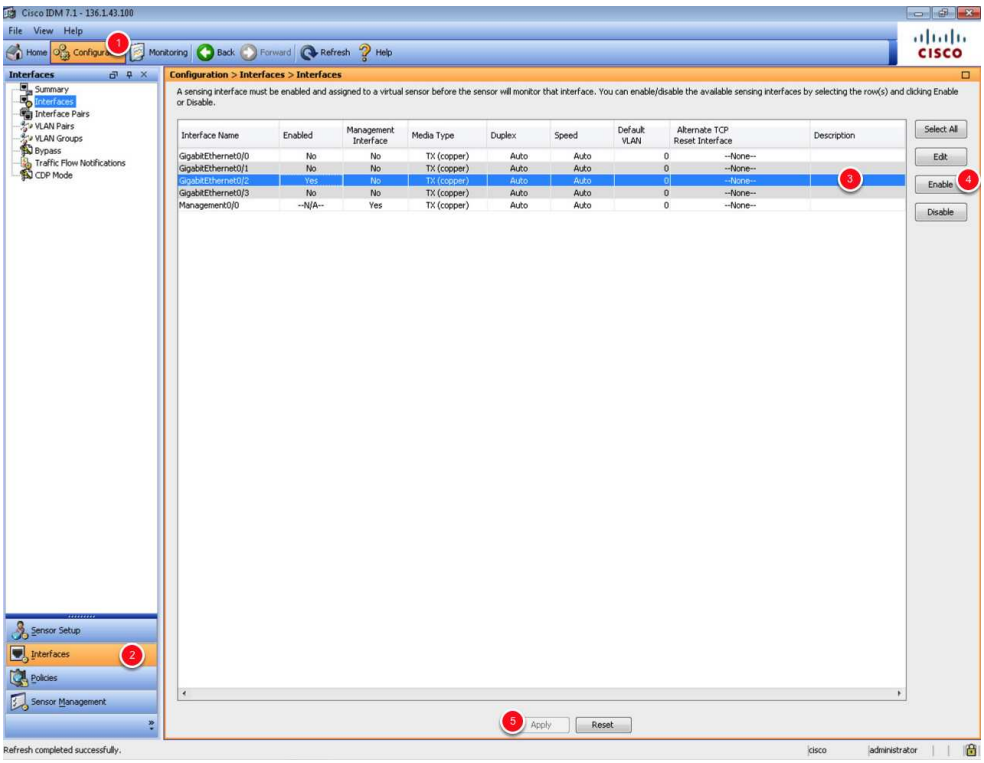
https://t.me/learningnets

```
sensor(config-net-use)# exit
sensor(config-net)# exit
Apply Changes?[yes]:
!
sensor(config)# service network-access
sensor(config-net)# router-devices 150.1.4.4
sensor(config-net-rou)# communication telnet
sensor(config-net-rou)# profile-name R4
sensor(config-net-rou)# response-capabilities rate-limit
sensor(config-net-rou)# block-interfaces FastEthernet0/0 in
sensor(config-net-rou-blo)# exit
sensor(config-net-rou)# exit
sensor(config-net)# router-devices 150.1.6.6
sensor(config-net-rou)# communication telnet
sensor(config-net-rou)# profile-name R6
sensor(config-net-rou)# response-capabilities block
sensor(config-net-rou)# block-interfaces FastEthernet0/0 in
sensor(config-net-rou-blo)# exit
sensor(config-net-rou)# exit
sensor(config-net)# exit
Apply Changes?[yes]:
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 2152 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# retired false
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine flood-host
sensor(config-sig-sig-flo)# event-action produce-alert|request-rate-limit
sensor(config-sig-sig-flo)# event-action-settings
sensor(config-sig-sig-flo-eve)# external-rate-limit-type percentage
sensor(config-sig-sig-flo-eve-per)# external-rate-limit-percentage 25
sensor(config-sig-sig-flo-eve-per)# exit
sensor(config-sig-sig-flo-eve)# exit
sensor(config-sig-sig-flo)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# signatures 2004 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine atomic-ip
sensor(config-sig-sig-ato)# fragment-status want-fragments
sensor(config-sig-sig-ato)# event-action produce-alert|request-block-host
sensor(config-sig-sig-ato)# exit
sensor(config-sig-sig)# exit
```

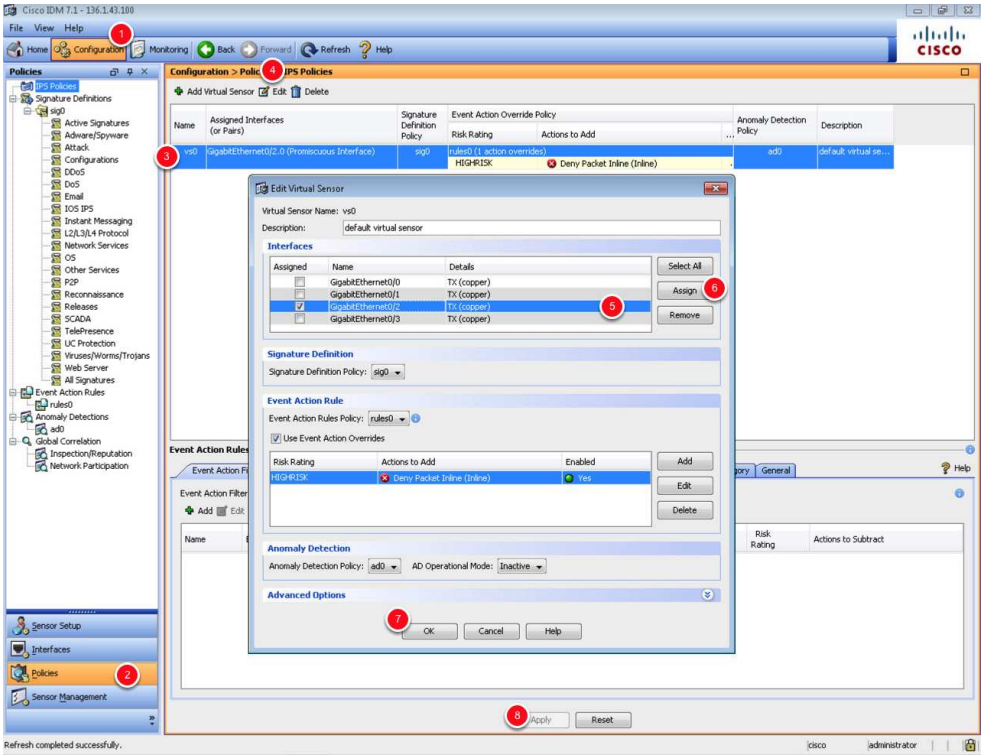
https://t.me/learningnets

```
sensor(config-sig)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, enable Gi0/2 interface:

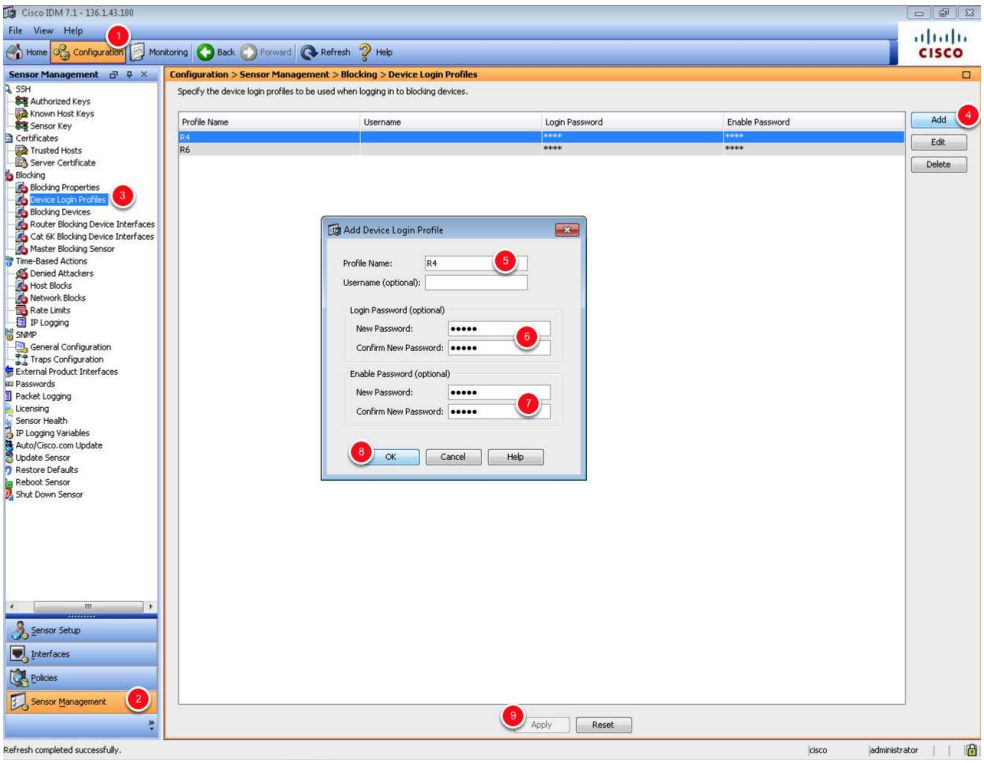


Assign the interface to the virtual sensor:

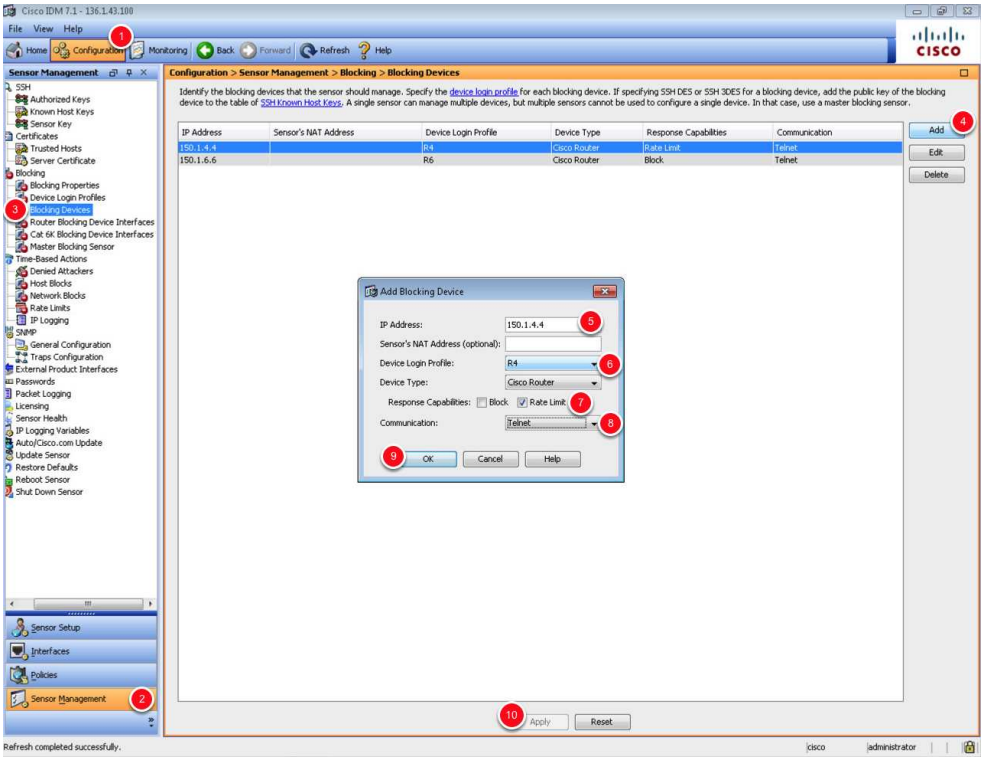


https://t.me/learningnets

Configure login profiles for both R4 and R6:

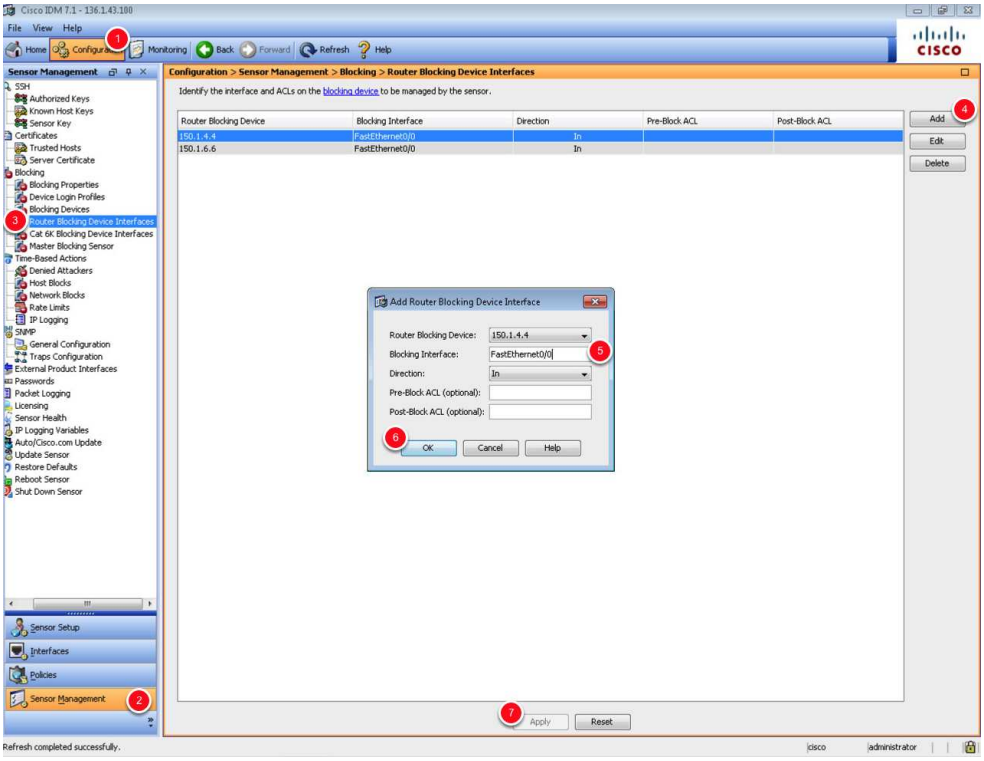


Configure R4 and R6 as blocking devices; use telnet and specify the actions:

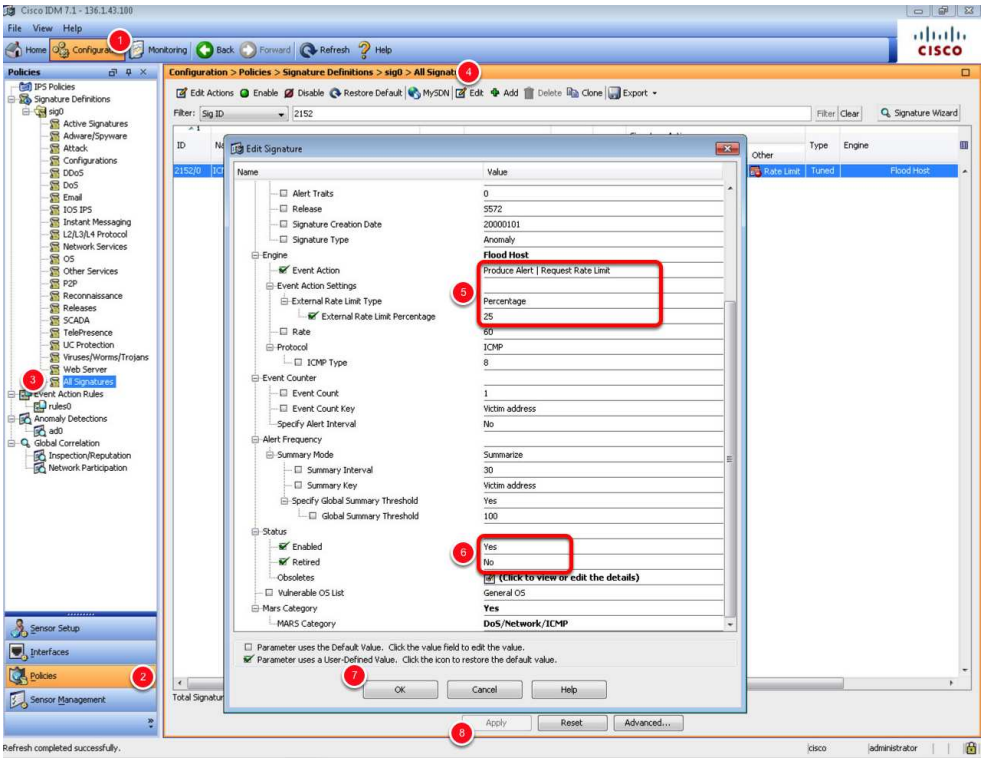


Configure the device interfaces to be used for ACL blocking and rate-limiting, for both R4 and R6:

https://t.me/learningnets

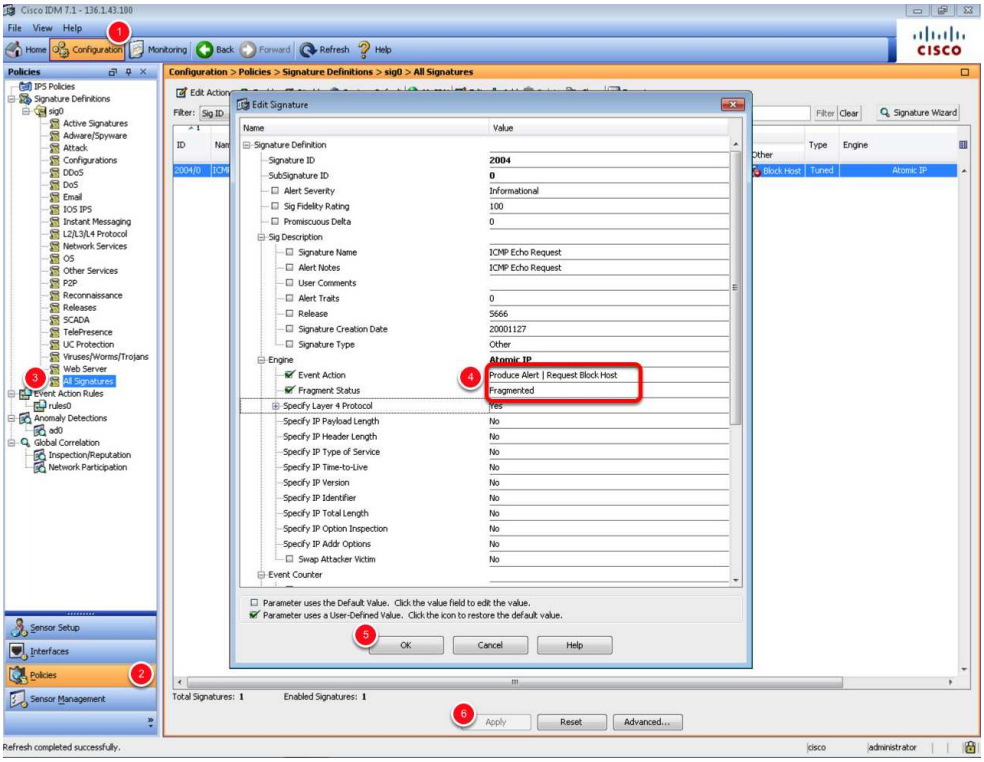


Activate, enable, and configure the ICMP Flood signature:



Enable and configure the ICMP Echo Request signature:

https://t.me/learningnets



Verification

Generate ICMP fragmented packets from R4 to trigger the ACL blocking:

```
R4#ping 150.1.6.6 size 1600

Type escape sequence to abort.
Sending 5, 1600-byte ICMP Echos to 150.1.6.6, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/8 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518000607 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 18:35:33 2003/01/31 18:35:33 UTC
signature: description=ICMP Echo Request id=2004 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request
  interfaceGroup: vs0 vlan: 43
  participants:
    attacker:
      addr: locality=OUT 136.1.43.4
    target:
```

https://t.me/learningnets

```
      addr: locality=OUT 150.1.6.6
      os: idSource=unknown relevance=relevant type=unknown actions:
blockRequested: true
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 15 interface: ge0_2

  protocol: icmp
```

Verify that IPS has an active telnet session to R6 and that ACL was applied:

```
R6#show tcp brief
TCB        Local Address               Foreign Address             (state)
496ABB18  150.1.6.6.23                136.1.43.100.42623          ESTAB
!
!R6#show users
    Line       User       Host(s)              Idle       Location
*  0 con 0                 idle                 00:00:00
514 vty 0                  idle                 00:00:10 136.1.43.100

  Interface    User               Mode         Idle     Peer Address
!
!R6#show running-config interface fastEthernet0/0
Building configuration...

Current configuration : 140 bytes
!
interface FastEthernet0/0
 ip address 136.1.69.6 255.255.255.0 ip access-group IDS_FastEthernet0/0_in_0 in
 duplex auto
 speed auto
!
!R6#show ip access-lists interface fastEthernet0/0
Extended IP access list IDS_FastEthernet0/0_in_1 in
    10 permit ip host 136.1.43.100 any (48 matches) 20 deny ip host 136.1.43.4 any

    30 permit ip any any (21 matches)
```

Generate an ICMP flood attack from R5 to trigger the rate-limiting:

```
R5#ping 150.1.6.6 repeat 200

Type escape sequence to abort.
Sending 200, 100-byte ICMP Echos to 150.1.6.6, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

https://t.me/learningnets

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (200/200), round-trip min/avg/max = 1/2/4 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518000651 severity=medium vendor=Cisco alarmTraits=32768
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 18:53:23 2003/01/31 18:53:23 UTC
signature: description=ICMP Flood id=2152 created=20000101 type=anomaly version=S572
    subsigId: 0
  interfaceGroup: vs0 vlan: 43
  participants:
    attacker:
      addr: locality=OUT 136.1.45.5
    target:
      addr: locality=OUT 150.1.6.6
      os: idSource=unknown relevance=relevant type=unknown actions:
rateLimitRequested: true
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 85
  threatRatingValue: 65 interface: ge0_2

  protocol: icmp
```

Verify that IPS has an active telnet session to R4 and that rate-limiting was applied:

```
R4#show tcp brief
TCB        Local Address               Foreign Address             (state)
4AA76150  www.google.com.23           136.1.43.100.46391          ESTAB
!
!R4#show users
    Line       User       Host(s)              Idle       Location
*  0 con 0                 idle                 00:00:00
514 vty 0                  idle                 00:00:17 136.1.43.100

  Interface    User               Mode         Idle     Peer Address
!
!R4#show running-config interface fastEthernet0/0
Building configuration...

Current configuration : 137 bytes
!
interface FastEthernet0/0
```

https://t.me/learningnets

```
 ip address 136.1.45.4 255.255.255.0
 duplex auto
 speed auto service-policy input IDS_RL_POLICY_MAP_1
!
!R4#show policy-map interface fastEthernet0/0
 FastEthernet0/0

  Service-policy input: IDS_RL_POLICY_MAP_1

    Class-map: IDS_RL_CLASS_MAP_icmp-xxBx-8-25_1 (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps Match: access-group name IDS_RL_ACL_icmp-xxBx-8-25_1
        0 packets, 0 bytes
        5 minute rate 0 bps police:
cir 25 %
cir 25000000 bps, bc 781250 bytes
        conformed 0 packets, 0 bytes; actions:
          transmit
        exceeded 0 packets, 0 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps

    Class-map: class-default (match-any)
      55 packets, 4070 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
!
!R4#show ip access-lists IDS_RL_ACL_icmp-xxBx-8-25_1
Extended IP access list IDS_RL_ACL_icmp-xxBx-8-25_1 10 permit icmp any host 150.1.6.6 echo
```

Verify on the IPS that both blocking and rate-limiting appear as active:

```
sensor# show statistics network-access
Current Configuration
   LogAllBlockEventsAndSensors = true
   EnableNvramWrite = false
   EnableAclLogging = false
   AllowSensorBlock = false
   BlockMaxEntries = 250
   MaxDeviceInterfaces = 250
   NetDevice
      Type = Cisco
      IP = 150.1.4.4
      NATAddr = 0.0.0.0
```

https://t.me/learningnets

```
      Communications = telnet
      ResponseCapabilities = rateLimit
      BlockInterface
         InterfaceName = FastEthernet0/0
         InterfaceDirection = in
   NetDevice
      Type = Cisco
      IP = 150.1.6.6
      NATAddr = 0.0.0.0
      Communications = telnet
      ResponseCapabilities = block
      BlockInterface
         InterfaceName = FastEthernet0/0
         InterfaceDirection = in
State
   BlockEnable = true
   NetDevice
      IP = 150.1.4.4
      AclSupport = uses Named ACLs
      Version = 15.1
      State = Active
   NetDevice
      IP = 150.1.6.6
      AclSupport = uses Named ACLs
      Version = 15.1
      State = Active BlockedAddr
      Host IP = 136.1.43.4
         Vlan = 43
         ActualIp =
         BlockMinutes = 30
         MinutesRemaining = 3 RateLimit
      Protocol = icmp Rate = 25
DestAddr = 150.1.6.6

      Data = echo-request
      RateLimitMinutes = 30
      MinutesRemaining = 21
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Event Summarization

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Configure signatures as follows:
  - Tune the ICMP echo request signature so that it generates an alarm for every match.
  - Every ICMP echo-reply packet should be reported only if the number of packets within 30 seconds does not exceed 15; if exceeded, present only the maximally condensed report every 30 seconds.
- Configure IPS to allow communication between R4 and R5 using its Gi0/0 and Gi0/1 interfaces as per the diagram.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

https://t.me/learningnets

Event summarization is an important feature of IPS signature alerting, because it allows reducing the load of information and protecting against IPS evasion tools, such as Stick. These tools attempt to create a huge load of non-relevant IPS events, overloading the sensor and hiding the actual attack within the decoy alerts. To resist this type of attack, Cisco IPS implements event counting and summarization, effectively performing two-level summarization.

Event Counting

The first summarization level is known as event counting. Its purpose is to count repetitive events, and fire an alarm only when the certain threshold is reached. An event is defined as repetitive if it triggers the same signature and parameters specified by the **event-count-key** match. Event count key is essentially a mask that selects fields for counting unique events from the following list: attacker IP, attacker port, victim IP, victim port. If the key is **Axxx**, it would select just the attacker IP address to define repeating events; if two attackers generate ICMP echo events, those two will be counted separately—no matter how many targets they ping, there will be just two series of events counted. If the key is **xxBx**, it selects the victim IP address, and the series of events will be generated based on unique victim addresses. For example, if ten attackers generate packets to three victims, only three separate series of counting will occur, for every victim. The other three keys are **AxBx** , **Axxb**, and **AaBb**, meaning attacker IP/victim IP, attacker IP/victim port, and attacker/victim IP/ports, with the latter generating the most separate event series.

Every time a signature fires, the IPS starts counting the events until the value specified by the variable **event-count** is reached; after this, it generates an aggregated alarm, reporting the total number of repetitive events for every unique sequence tracked. By default, **event-count=1**, so every signature match generates an event. After this, the event counter is reset to zero. In addition to unconstrained counting, it is possible to define a time-interval over which to perform counting. With the interval defined, the event counter for a given signature is reset to zero every alert-interval, which is a signature attribute. This allows for triggering an alarm only if the rate of repetitive events exceeds the threshold. For example, you may only want to trigger an alarm if there are 100 ICMP echo packets sent within an interval of 20 seconds. You set event-count to 100 and alert-interval to 20 seconds. There will be no event triggered unless the rate rises above 100/20 hits/second.

Event Summarization

Following the event counting procedure is event summarization, which is the second layer of aggregation. It applies fixed time intervals and summarizes information for events within this interval into a single report. For every signature, you may select a

https://t.me/learningnets

different alarm summarization mode. The first and most simple mode, called Fire-All, generates an alarm for every event produced by event counting.

The next mode (also the default) is known as summarize. When this mode is configured for a signature, the special parameter known as summary-interval is used to perform aggregation. The IPS counts every summary-interval second for the signature and detects the first alert during the interval; this first alert is reported immediately. All subsequent hits within the interval are delayed and reported only when the summary-interval expires. If there are events in subsequent intervals, they will have only one summary event generated until there is an interval with no events at all. After this, the IPS will go back to generating an alert for the first event and summary for the interval.

The final report for every summary-interval is condensed based on the summary-key setting for the signature. The concept of key is widely used for aggregation purposes. Here it defines the fields that could be aggregated and ignored among different events detected. Effectively, the key is the mask that selects relevant fields from attacker/victim IP address/ports pairs and stores them while performing summarization or counting. Here are the key values and the diagram to illustrate their effect when performing event summarization (summary-keys are pre-configured for every signature and should not be modified unless that is really needed):

- **Axxx**: Events are summarized using the attacker's IP address as the key. A summary report will be generated for every attacker's IP address, and all other information, including the attacker's port and victim's IP address and port, will be summarized and ignored.
- **Axxb **: The attacker's IP and the victim ports are the key. For every combination of an attacker IP address and victim port, a summary report is generated.
- **AxBx**: The summary events are created for every attacker/victim address pair detected during the summary-interval.
- **AaBb**: The summary report will contain an entry for every combination of attacker/victim IP/port addreses.
- **xxBx**: Creates an entry in the summary report for every victim. This is the opposite of Axxx.

The next summarization mode is known as Fire-Once. During the summary-interval, there will be one alert generated for every unique event based on the summary-key to define uniqueness. For example, if the key is **Axxx**, there will be one alarm generated for every attacker IP address; if the key is **Axxb**, there will be one alarm generated for every first matching packet with unique attacker source and victim port address. The most intense alert summarization mode is known as global-

https://t.me/learningnets

summarize. This mode will report a summary at the end of every summary-interval, with just the general information about the attack, such as the number of packets matched, and no information stored about the attackers or victims.

The alarm summarization modes described above where fixed; in other words, their behavior does not change with time. However, it is possible to configure a signature to switch between the summarization modes, based on the number of alarms generated. This allows for dynamic response to the attack behavior, damping the amount of events aggressively when there is a flood of attacks and allowing for a more detailed view when the load is low. This is known as variable summarization. It is implemented as follows: For every summary mode selected, you may configure a threshold level of alarms that should be reached during the current summary-interval to change the summary mode. Changing the summarization mode is called upgrading the summarization mode in IPS term. It is possible to do the following upgrade sequences:

- Fire-All -> Summarize -> Global-Summarize
- Fire-Once -> Global-Summarize
- Summarize -> Global-Summarize

The upgrades are performed based on the two thresholds set for a signature: **summary-threshold** and **global-summary-threshold**. For example, assume that you have a signature configured as follows:

```
Summary Mode: Fire-All
Summary-Interval: 300 seconds
Summary-Threshold: 100
Global-Summary-Threshold: 200
```

Let’s assume that during the first summary-interval there were 50 events. That means that the next interval would not use any upgrades and continue using Fire-All mode. Now assume that the second interval produced 120 events. That means that the third interval would upgrade and use the Summarize method to produce alarms. If there were more than 200 events during the third interval, the summarization mode would change to global-summarize. If during some interval the number of events fell below both thresholds, the signature would return to normal summarization mode.

Configuration

```
SW3:
interface GigabitEthernet1/0/8
 switchport mode access
```

https://t.me/learningnets

```
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

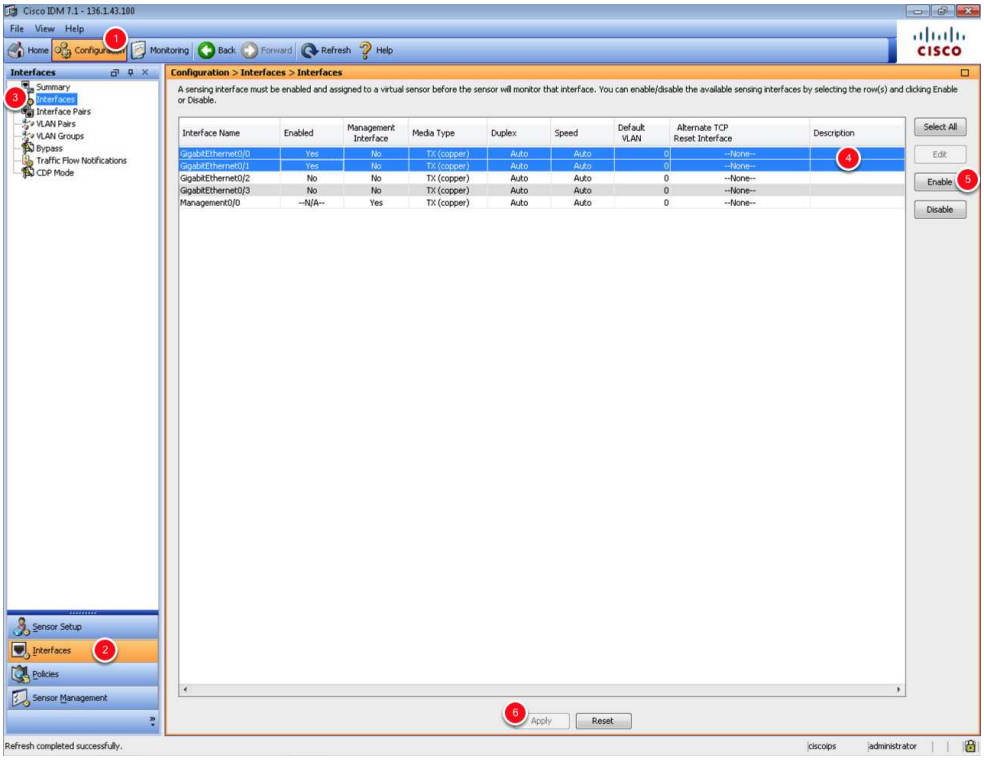
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
!
sensor(config)# service signature-definition sig0
```

https://t.me/learningnets

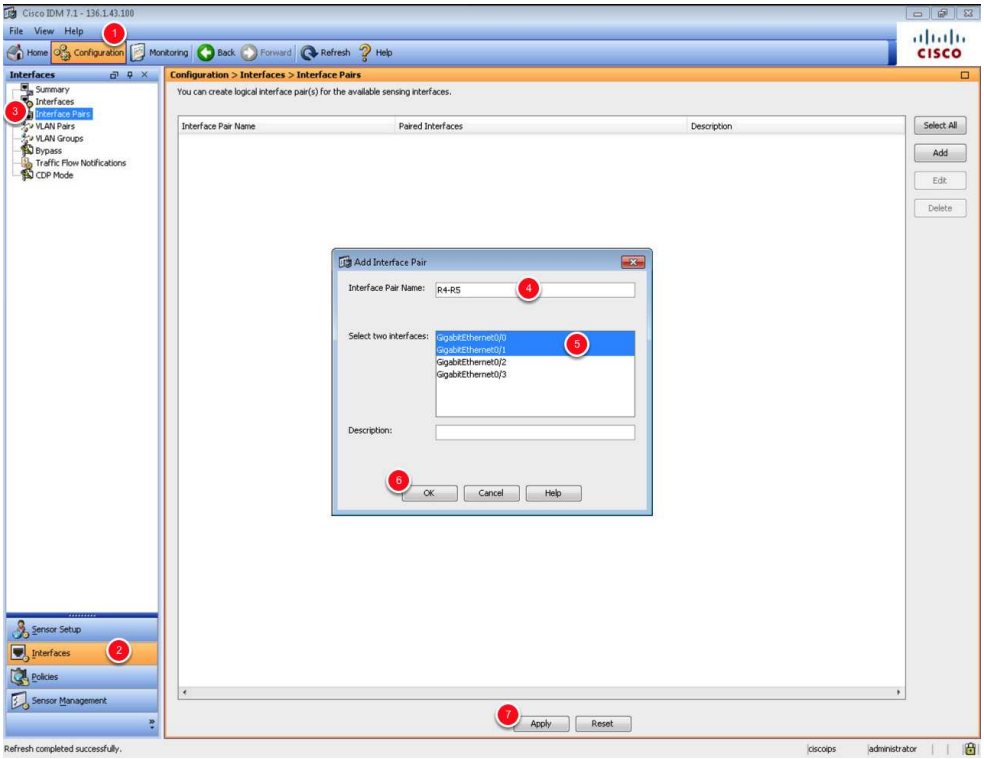
```
sensor(config-sig)# signatures 2004 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine atomic-ip
sensor(config-sig-sig-ato)# event-action produce-alert
sensor(config-sig-sig-ato)# exit
sensor(config-sig-sig)# alert-frequency
sensor(config-sig-sig-ale)# summary-mode fire-all
sensor(config-sig-sig-ale-fir)# exit
sensor(config-sig-sig-ale)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# signatures 2000 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine atomic-ip
sensor(config-sig-sig-ato)# event-action produce-alert
sensor(config-sig-sig-ato)# exit
sensor(config-sig-sig)# alert-frequency
sensor(config-sig-sig-ale)# summary-mode fire-all
sensor(config-sig-sig-ale-fir)# specify-summary-threshold yes
sensor(config-sig-sig-ale-fir-yes)# summary-interval 30
sensor(config-sig-sig-ale-fir-yes)# summary-threshold 15
sensor(config-sig-sig-ale-fir-yes)# exit
sensor(config-sig-sig-ale-fir)# exit
sensor(config-sig-sig-ale)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

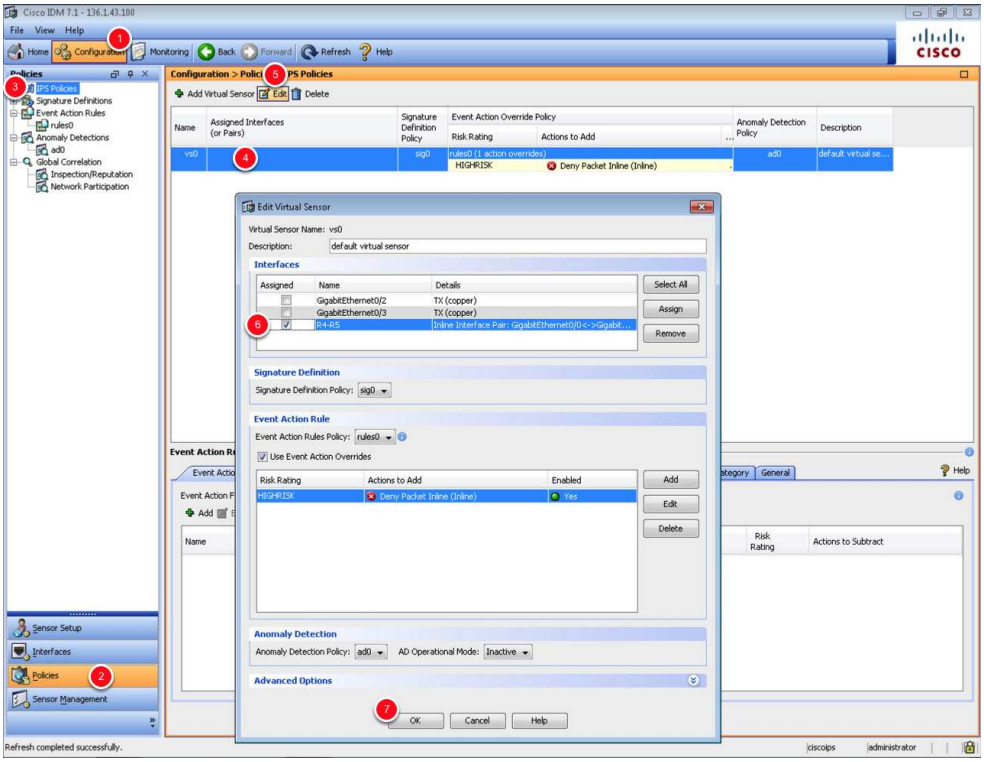


Create the inline interface-pair:

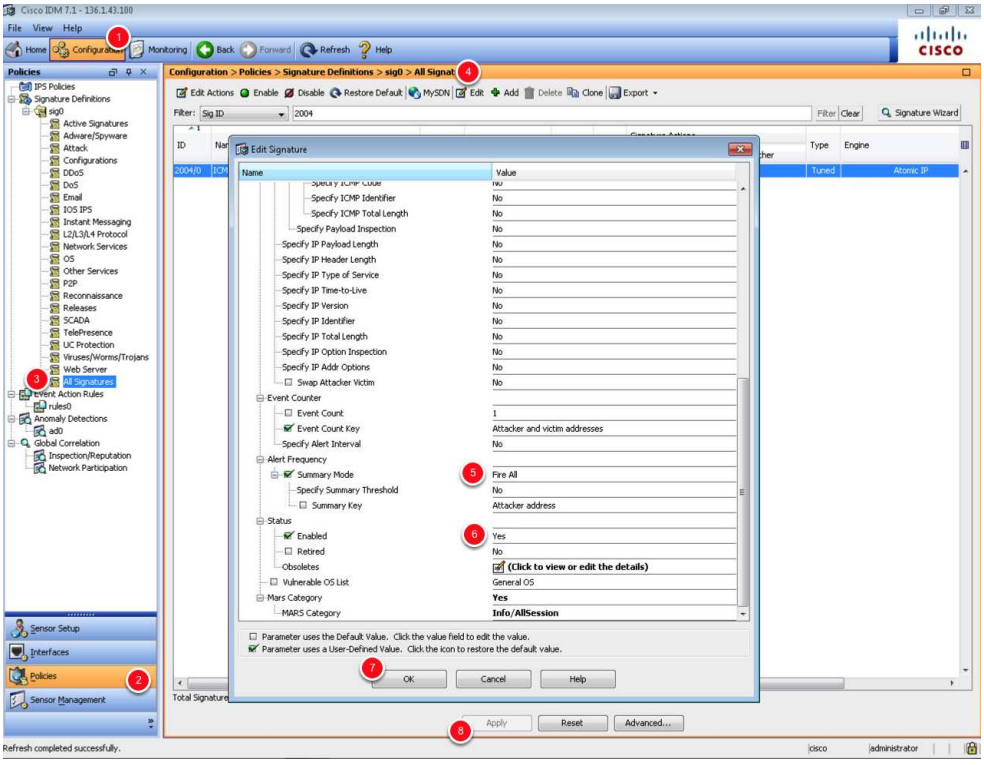


Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets

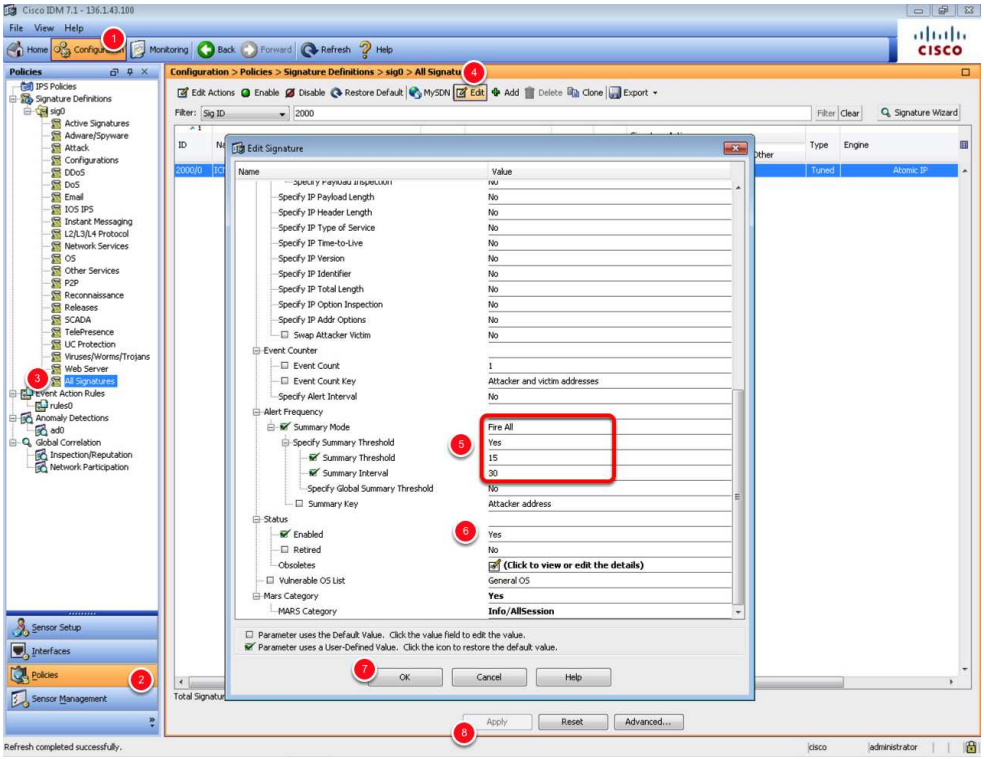


Enable and configure ICMP Echo Request signature:



Enable and configure ICMP Echo Reply signature:

https://t.me/learningnets



Verification

Generate a regular ping from R5 to R4; IPS should log five events for both ICMP Echo Request and ICMP Echo Reply:

```
R5#ping 150.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518000838 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 22:48:28 2003/01/31 22:48:28 UTC
signature: description=ICMP Echo Request id=2004 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request
  interfaceGroup: vs0
  vlan: 0
  participants: attacker:
```

https://t.me/learningnets

```
addr: locality=OUT 136.1.45.5
 target:
addr: locality=OUT 150.1.4.4
      os: idSource=unknown relevance=relevant type=unknown
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 35
  interface: ge0_1
  protocol: icmp


evIdsAlert: eventId=1376328122518000839 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 22:48:28 2003/01/31 22:48:28 UTC
signature: description=ICMP Echo Reply id=2000 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Reply
  interfaceGroup: vs0
  vlan: 0
  participants:  attacker:
addr: locality=OUT 150.1.4.4
target:
addr: locality=OUT 136.1.45.5

      os: idSource=unknown relevance=relevant type=unknown
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 35
  interface: ge0_0
  protocol: icmp
```

Generate an extended ping from R5 to R4, issuing 20 packets; IPS should log 20 events for both ICMP Echo Request and ICMP Echo Reply:

```
R5#ping 150.1.4.4 repeat 20

Type escape sequence to abort.
Sending 20, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!!!!!!!!!!!!!!!!

Success rate is 100 percent (20/20), round-trip min/avg/max = 1/3/12 ms
```

Wait 30 seconds and generate the same extended ping again; this time IPS should log a summary event for the ICMP Echo Reply (note that victim address shows up

https://t.me/learningnets

as unknown, because of the default summary functional mode):

```
R5#ping 150.1.4.4 repeat 20

Type escape sequence to abort.
Sending 20, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (20/20), round-trip min/avg/max = 1/3/12 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518001094 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 23:10:06 2003/01/31 23:10:06 UTC
signature: description=ICMP Echo Reply id=2000 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Reply
  interfaceGroup: vs0
  vlan: 0
  participants:  attacker:
 addr: locality=OUT 150.1.4.4
 target:
 addr: locality=OUT 0.0.0.0
      os: idSource=unknown relevance=unknown type=unknown
summary: final=true initialAlert=0 summaryType=Regular 20
 alertDetails: Regular Summary: 20 events this interval ;
  riskRatingValue: targetValueRating=medium 25
  threatRatingValue: 25
  interface: ge0_0
  protocol: icmp


evIdsAlert: eventId=1376328122518001095 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/01/31 23:10:16 2003/01/31 23:10:16 UTC
  signature: description=ICMP Echo Request id=2004 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request
  interfaceGroup: vs0
  vlan: 0
  participants:
```

https://t.me/learningnets

```
attacker:
 addr: locality=OUT 136.1.45.5
target:
 addr: locality=OUT 150.1.4.4

      os: idSource=unknown relevance=relevant type=unknown
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 35
  interface: ge0_1
  protocol: icmp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Event Processing

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Configure the IPS sensor so that all ICMP echo request packets sent to R6's Loopback0 subnet generate a host block request; if sourced from R5's Loopback0 interface, the block request should be suppressed.
- Make sure the source/destination IP addresses for R5/R6 Loopback0 subnets appear as meaningful names **R5_LO0** and **R6_LO0**.
- Configure IPS to allow communcation between R4 and R5 using its Gi0/0 and Gi0/1 interfaces according to the diagram.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

When a signature generates an event, the event goes through a series of event

https://t.me/learningnets

processing rules. These rules define whether to actually trigger an event and what action to associate with it. Here are the key concepts used in event processing:

- Event Variables: allows for macro variables to be used with event processing
- Target Value Rating (TVR): defines the importance of a particular resource
- Event Action Override: changes the action for a specific event
- Event Action Filters: damps specific events

Event Variables

These are normal variables that can be used in the event output—that is, the output you see in the IPS logs. For example, you may want to give meaningful names to your network ranges, such as TECHNICAL to 10.0.0.0/24 and SALES to 10.0.1.0/24. You could define these macro variables by navigating to **Configuration > Policies > Event Action Rules > rules0 > Event Variables**; all variable are of type address.

Target Value Rating

This value specifies the importance of a particular network resource and is used when calculating a Risk Rating (RR) for a particular event. You assign TVR to an IP address range, using the values Mission Critical (200), High (150), Medium (100), and Low (75). You may use the variables defined previously when working with TVR configuration. We will see how TVR is used later for RR computations.

Event Action Override

This feature allows for changing the default signature action (for example, changing from produce event to deny attacker inline) if the event Risk Rating (RR) falls within a certain range. This allows for accounting of user-defined TVR assigned to corporate assets and changing the event action if the event appears to be more dangerous than it looked originally. Here is the formula for RR calculations:

RR = (ASR*TVR*SFR)/10000 + ARR – PD, where:

- **ASR** is Attack Severity Rating, taken from the Severity field for signature configuration with the following possible values: 25 (Information), 50 (Low), 75 (Medium), and 100 (High). This is configured by Cisco by default for all signatures.
- **TVR** is Target Value Rating, discussed previously, ranging from 75 to 200.
- **SFR** is Signature Fidelity Rating, which defines the accuracy of the signature, or, the risk of a false positive. This field is assigned to every signature and takes values from 0 to 100. You should not change the default values assigned by Cisco unless

https://t.me/learningnets

absolutely necessary.

- **ARR** is Attack Relevance Rating. This value is calculated by IPS automatically, for a given event and victim IP address. The IPS correlates to the operating system passively detected for the victim system and the list of vulnerable OSs configured for the signature. Based on this, the relevance rating is calculated. For example, an IIS attack is probably not relevant for a UNIX server.
- **PD** is Promiscuous Delta; this value is set for every signature and specifies the number of points that should be deducted from RR if the IPS is in promiscuous monitoring mode. The idea is that promiscuous monitoring is less reliable than inline, and thus the overall risk might be overestimated. The range for PD is from 0 to 30.

Every event action override rule defines a range of RR values and the action to produce when the event's risk rating falls within this range. For example, you may configure it so that 90 < Risk Rating < 100 results in deny packet inline (the default Cisco rule). Defining an actual range requires understanding of the RR formula and manipulation of TVRs set for network resources. To create an action override rule, navigate to **Configuration > Policies > Even Action Rules > rules0 > Event Action Overrides** and click **Add**. You must define the action and the RR range.

As mentioned previously, IPS automatically calculates ARR based on the OS type detected for the victim. The detection is based on passive OS fingerprintin or POSFP procedure. In addition to automatic detection, you may provide OS information manually, by navigating to **Configuration > Policies > Event Action Rules > rules0 > OS Identification** and adding OS to IP address block association manually. You may aso enable passive OS fingerprinting under the same tab. As stated previously, every signature has a list of vulnerable OSs, and this list is used to calculate the ARR.

Event Action Filtering

The idea of action filtering is to subtract the action associated with an event if the event matches certain criteria. Among the criteria are signature ID, risk rating range, and attacker/victim IP addresses/ports. Every filter also has the action to subtract from the matching event. For example, you may want to delete the **deny-attacker-inline** action for all events with RR in range 30-50. Event filtering applies after event overriding, so you have a change to decline the decision made at the previous step. To create a new Event Action Filter, navigate to **Configuration > Policies > Event Action Rules > rules0 > Event Action Filters** from within the IDM menu.

Configuration

https://t.me/learningnets

Normally, ICMP Echo is an informational signature with high fidelity (value of 100). Thus, if we look at the formula for RR, we already have some values defined: RR = 25*TVR*100/10000 + ARR = TVR/4 + ARR. Now it is apparent that the RR could be at maximum 200/4+ARR=50+ARR. If we set the RR range to 50-60 and set the TVR to Mission-Critical, this will catch the event. In addition, the new event action rule should define action **host block request**. Next, we should create an event filtering rule that matches the source of R5's Loopback0 and RR range 50-60 for ICMP Echo Request signature; configure this rule to subtract the block host action.

SW3:

```
interface GigabitEthernet1/0/8
 switchport mode access
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

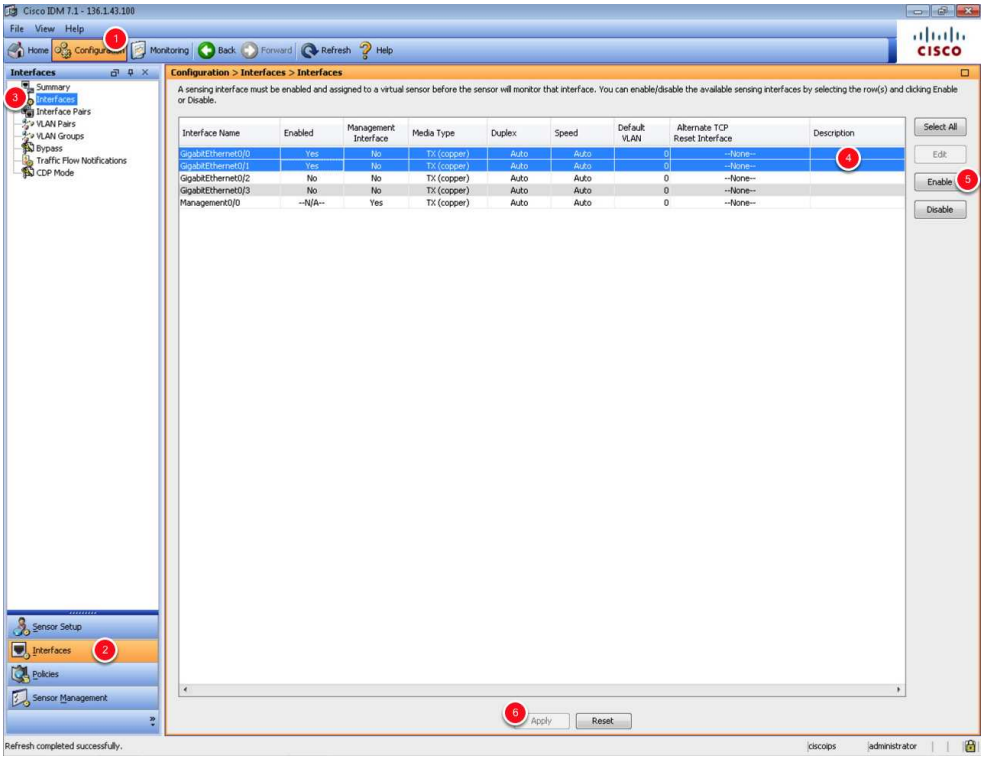
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
```

https://t.me/learningnets

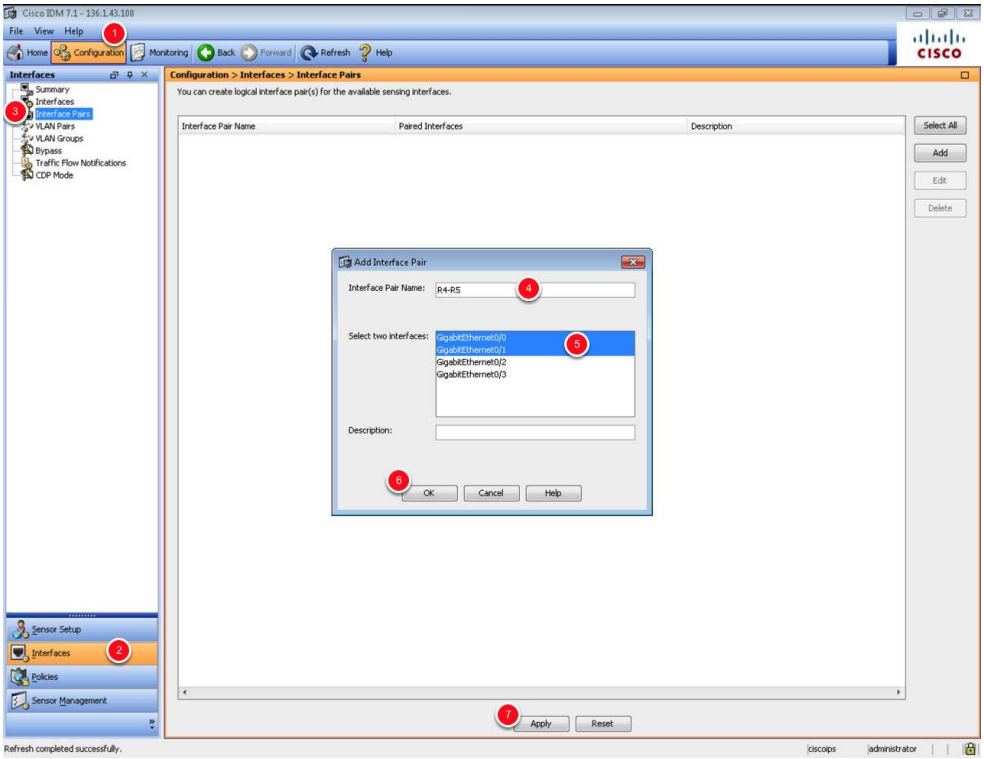
```
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 2004 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes?[yes]:
!
sensor(config)# service event-action-rules rules0
sensor(config-eve)# variables R5_LO0 address 150.1.5.5
sensor(config-eve)# variables R6_LO0 address 150.1.6.6
sensor(config-eve)# target-value mission-critical target-address 150.1.6.6
sensor(config-eve)# overrides request-block-host
sensor(config-eve-ove)# risk-rating-range 50-60
sensor(config-eve-ove)# exit
sensor(config-eve)# filters insert R5 begin
sensor(config-eve-fil)# attacker-address-range 150.1.5.5-150.1.5.5
sensor(config-eve-fil)# actions-to-remove request-block-host
sensor(config-eve-fil)# exit
sensor(config-eve)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

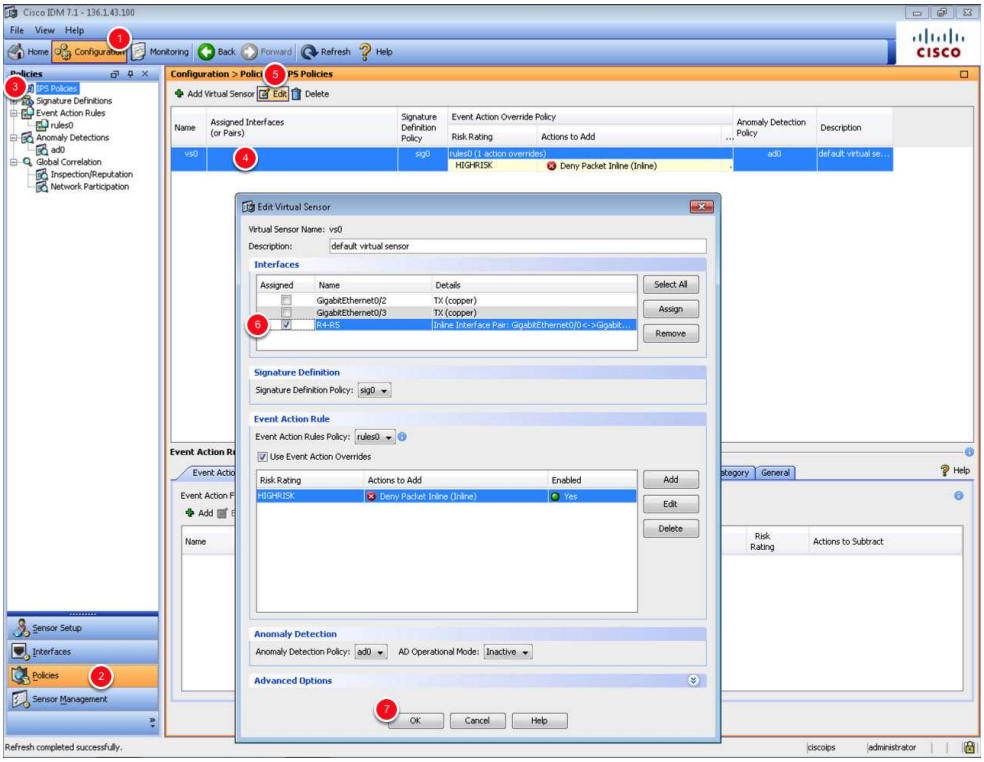


Create the inline interface-pair:

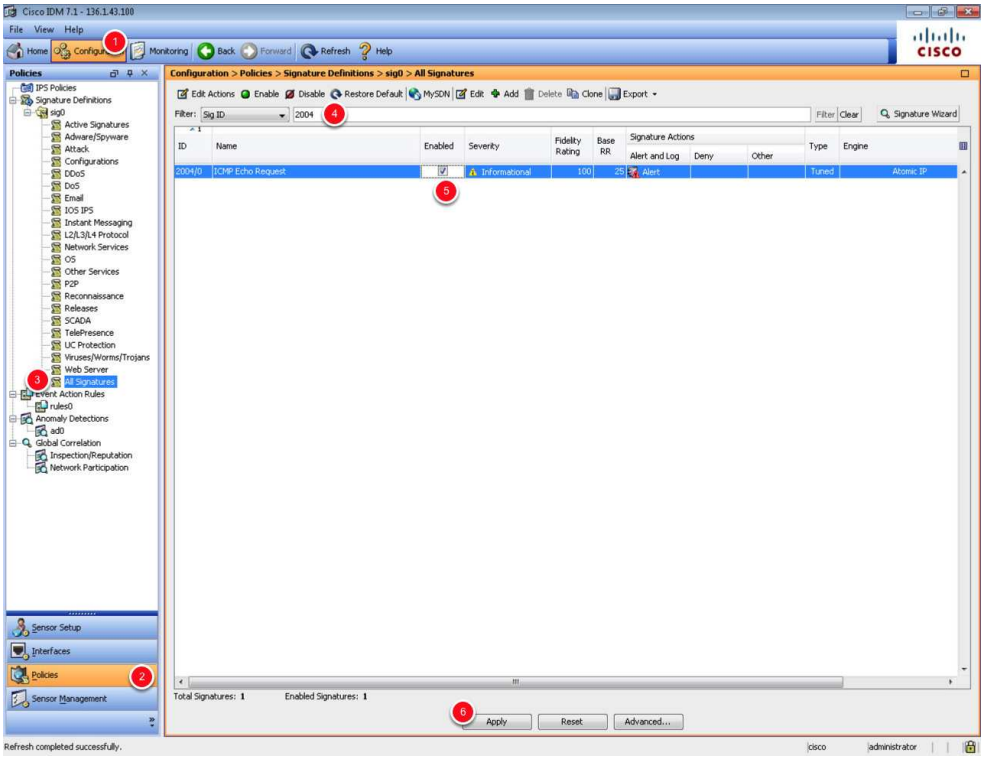


Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets

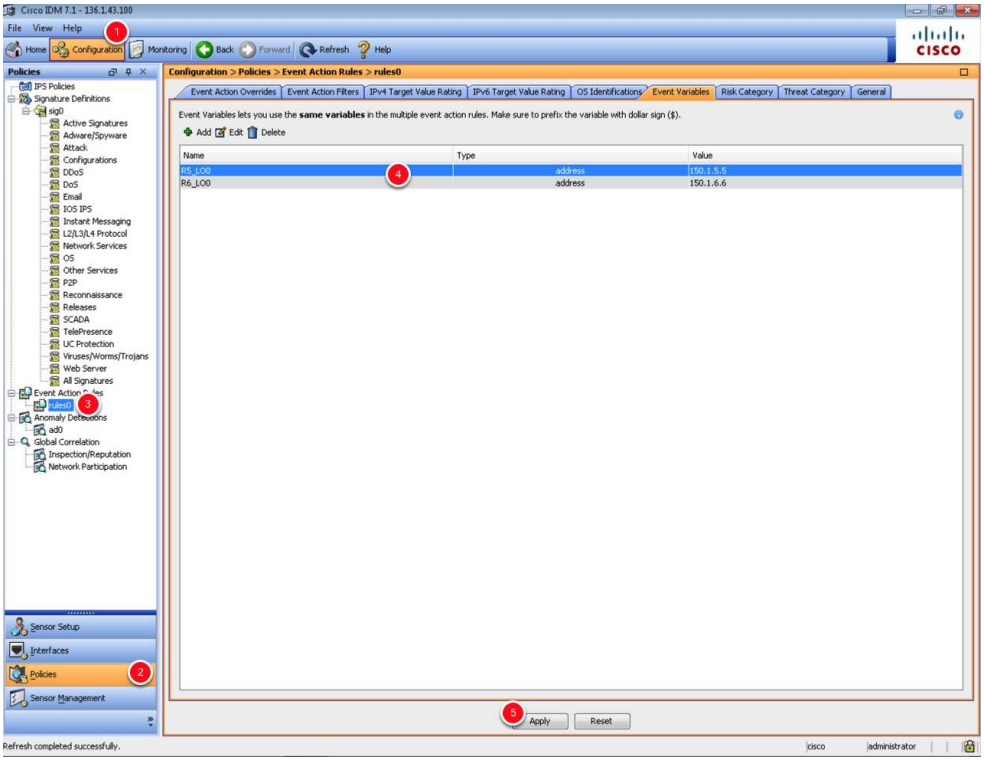


Enable the ICMP Echo Request signature:

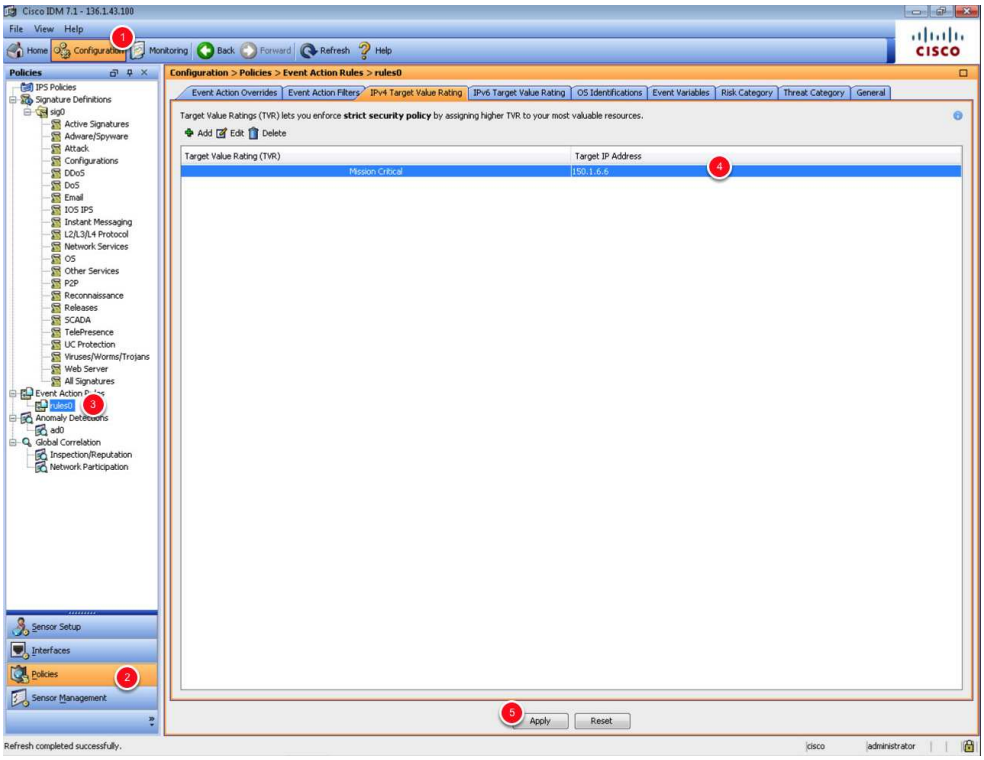


Define R5 and R6 Loopback0 addresses as event variables:

https://t.me/learningnets

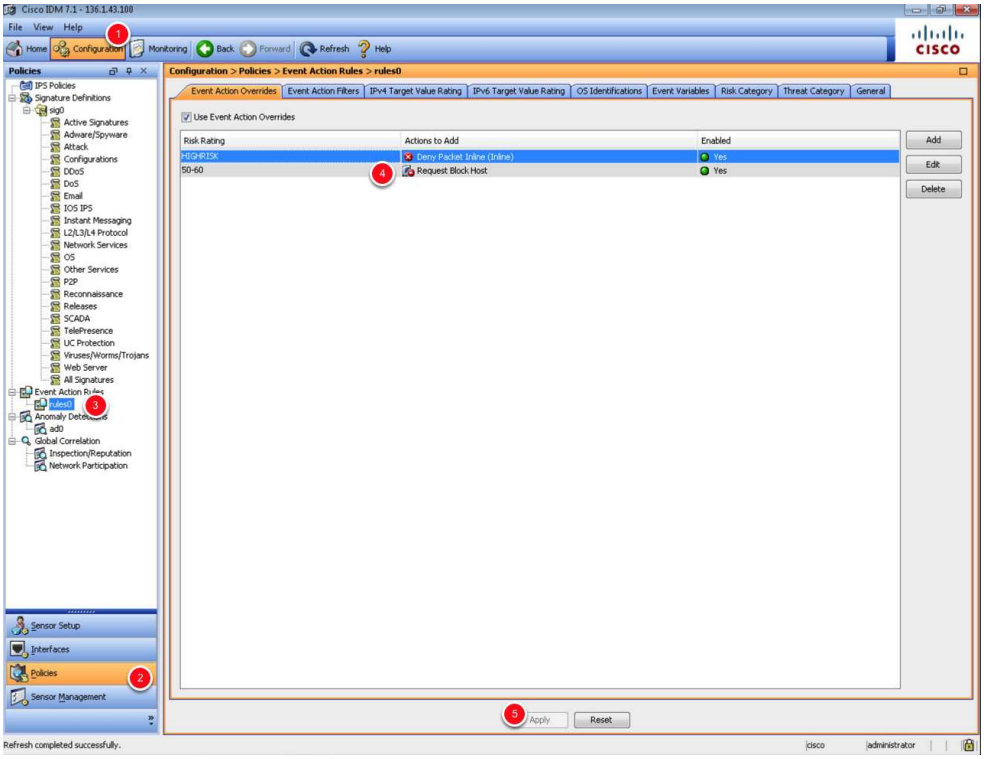


Assign R6 Loopback0 a TVR of Mission Critical:

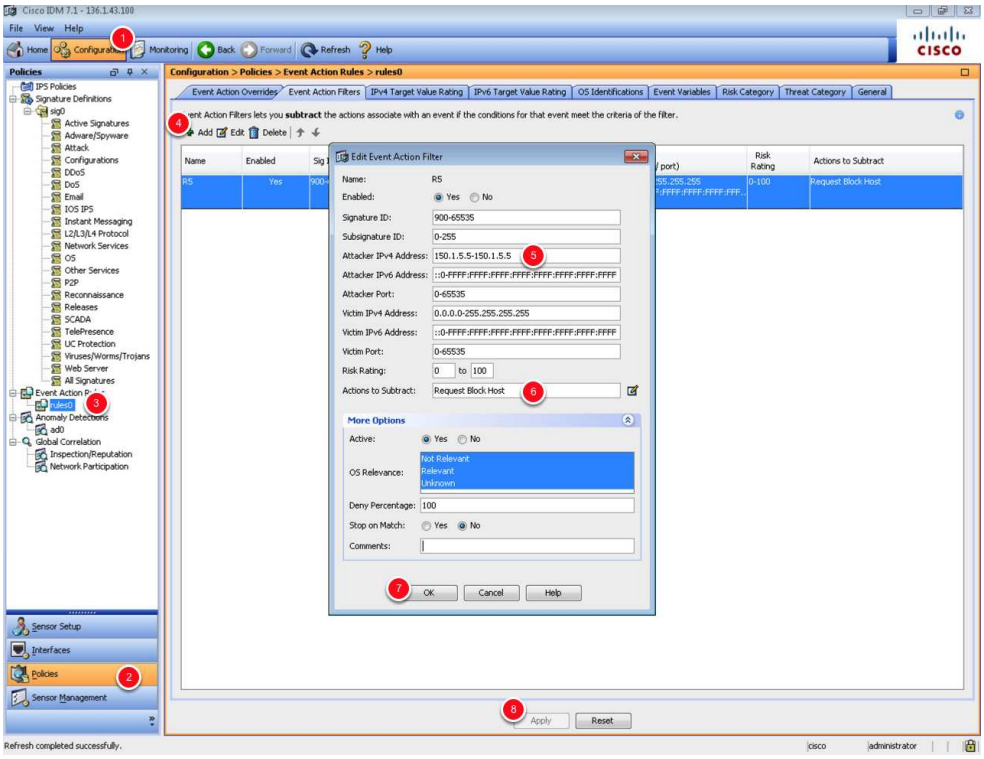


Add an event action override of request block host for an event with risk rating between 50-60:

https://t.me/learningnets



Configure an event action filter to remove the block host action if the attacker is R5's Loopback0 address:



Verification

Generate ICMP Echo Request packets from R5; along with the signature events, an error messages should be logged because there are no blocking devices configured

https://t.me/learningnets

but the action has been overriden.

```
R5#ping 150.1.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.6.6, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/8 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518001197 severity=informational vendor=Cisco alarmTraits=32768
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 00:19:31 2003/02/01 00:19:31 UTC
  signature: description=ICMP Echo Request id=2004 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request
  interfaceGroup: vs0
  vlan: 0
  participants:
    attacker:
      addr: locality=OUT 136.1.45.5
    target:  addr: locality=R6_LO0 150.1.6.6
      os: idSource=unknown relevance=relevant type=unknown
  actions:  blockRequested: true
 riskRatingValue: attackRelevanceRating=relevant targetValueRating=mission-critical 60
  threatRatingValue: 40
  interface: ge0_1
  protocol: icmp

 evShunRqst: eventId=1376328122518001198 vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 00:19:31 2003/02/01 00:19:31 UTC
  shunInfo:
    host: connectionShun=false
      srcAddr: 136.1.45.5
      vlan: 0
      destAddr: 150.1.6.6
      protocol: numericType=1 other
    timeoutMinutes: 30
  evAlertRef: hostId=sensor 1376328122518001197
```

https://t.me/learningnets

```
evStatus: eventId=1376328122518001199 vendor=Cisco
  originator:
    hostId: sensor
    appName: nac
    appInstanceId: 1050
  time: 2003/02/01 00:19:31 2003/02/01 00:19:31 UTC shunEntryAdded:
    description: Block Host
    shunInfo:
      host:
        srcAddr: 136.1.45.5
        srcPort: 0
        destAddr: 150.1.6.6
        destPort: 0
        protocol: numericType=1
        vlan: 0
        interface:
      timeoutMinutes: 30
    evAlertRef: hostId=sensor 1376328122518001197


evError: eventId=1376328122518001200 severity=error vendor=Cisco
  originator:
    hostId: sensor
    appName: nac
    appInstanceId: 1189
  time: 2003/02/01 00:19:31 2003/02/01 00:19:31 UTC
errorMessage: name=errSystemError Unable to execute a host block [136.1.45.5] because blocking is not configured
```

Generate ICMP Echo Request packets again from R5, but sourced from its Loopback0 address. This time the host-block action does not appear; it is filtered:

```
R5#ping 150.1.6.6 source loopback0 repeat 1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 150.1.6.6, timeout is 2 seconds:
Packet sent with a source address of 150.1.5.5 !
Success rate is 100 percent (1/1), round-trip min/avg/max = 4/4/4 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518001202 severity=informational vendor=Cisco alarmTraits=32768
  originator:
```

https://t.me/learningnets

```
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 00:20:37 2003/02/01 00:20:37 UTC
signature: description=ICMP Echo Request id=2004 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request
  interfaceGroup: vs0
  vlan: 0
  participants:
    attacker:
      addr: locality=R5_LO0 150.1.5.5
    target:  addr: locality=R6_LO0 150.1.6.6
      os: idSource=unknown relevance=relevant type=unknown
riskRatingValue: attackRelevanceRating=relevant targetValueRating=mission-critical 60

  threatRatingValue: 60
  interface: ge0_1
  protocol: icmp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Application Inspection and Control

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Configure IPS to allow communication between R4 and R5 using its Gi0/0 and Gi0/1 interfaces according to the diagram.
- Any use of the **HTTP CONNECT** method and the **FTP DELE** command on the segment between R4 and R5 should be only logged.
- Make the inspection as precise as possible and use the most specific engines.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

Overview

HTTP is probably one of the most popular protocols used in modern networks, and this is why it is such a popular target for exploitation. The primary goal of attackers is exploiting HTTP application vulnerabilities or using HTTP to hide malicious

https://t.me/learningnets

activity. FTP is the other popular protocol, which has a long history of security issues. Since version 5.x, Cisco IPS supports two advanced application inspection and control engines for HTTP and FTP protocols. These engines support deep protocol inspection, allowing for policy enforcement at Level 7 of the OSI model, using the properties and commands associated with FTP and HTTP protocols. Specifically, here is the list of the features supported with HTTP AIC:

- Detection of non-HTTP traffic tunneling through port 80.
- Ensuring RFC compliance of HTTP methods.
- Filtering traffic based on MIME types.
- Control of permitted traffic via user-defined policies, .e.g. based on regular expression.

Using a destination port of 80 for traffic that is not related to HTTP is known as tunneling traffic through port 80. You may see some instant messenger applications using this technique or some proxy programs. The IPS may check the HTTP requests and compare their format with the standard requirements. This also allows for detecting of protocol abuses or malformed messages. Filtering on MIME types allows for blocking certain types of traffic, such as video or audio or gzip-compressed content, which might be used to hide a virus. You can regulate which HTTP methods are allowed based on your security policy. For instance, your security policy may prohibit the use of **TRACE** or **HEAD** methods, or the **CONNECT** method that may be used for tunneling.

As for FTP, the IPS provides a large set of signatures to match almost any legitimate FTP command, or an unrecognized command. This allows you to enforce security policy based on the commands allowed for FTP.

If you want to configure an AIC signature, you may find the list by running the IDM and going to **Configuration > Policies > Signature Definitions > sig0 > All Signatures** and searching based on the Engine **AIC HTTP** or **AIC FTP**. Select the signature you need from the list and enable/tune it if needed. Notice that by default the signatures are configured to produce an alert and deny connection inline, so you may want to change this behavior. The HTTP signatures are grouped as follows:

- HTTP Request Method signatures
- HTTP Content-Type signatures (based on the MIME types used in HTTP response)
- HTTP Transfer Encoding signatures

As for FTP, there are signatures for FTP commands. Before tuning any of the above-mentioned signatures, you may navigate to **Configuration > Policies > Signature Definitions > sig0 > All Signatures** and click the **Miscellaneous** tab in the **Advanced** section. Under this tab you should set **Application Policy > HTTP Policy > Enable HTTP**

https://t.me/learningnets

and **Application Policy > Enable FTP** to make the respective engines active. Under the same tab, you may also edit the AIC web Ports attribute if you want to change the ports that the HTTP AIC engine is using. FTP engine always uses port 21 for inspection.

Configuration

SW3:

```
interface GigabitEthernet1/0/8
 switchport mode access
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

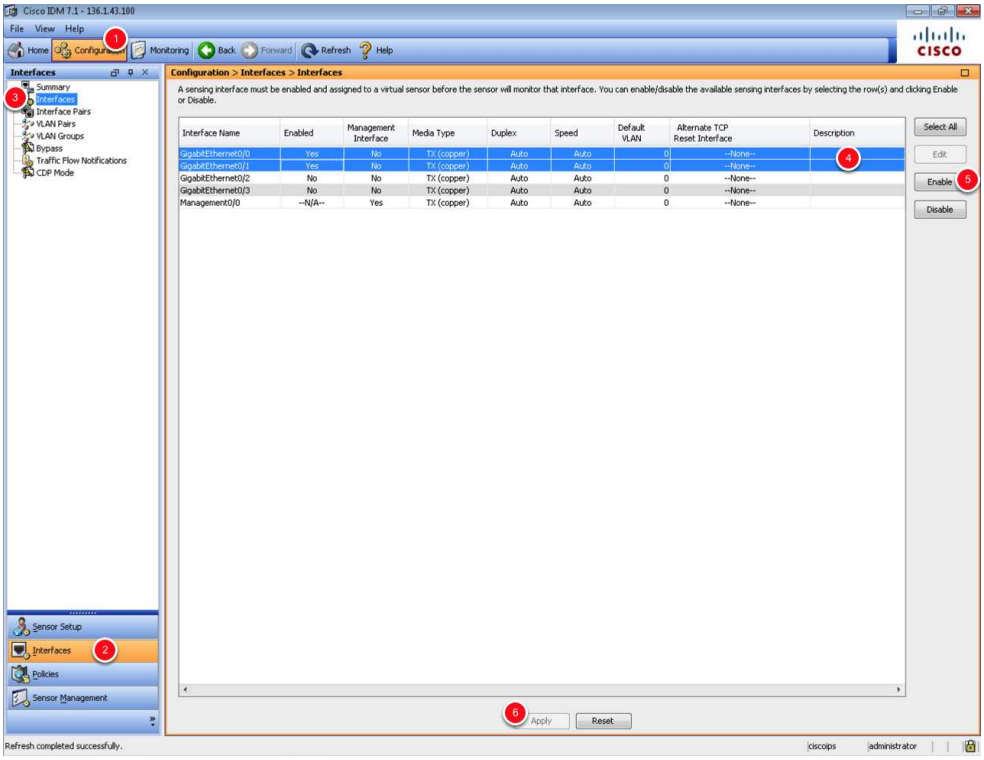
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
```

https://t.me/learningnets

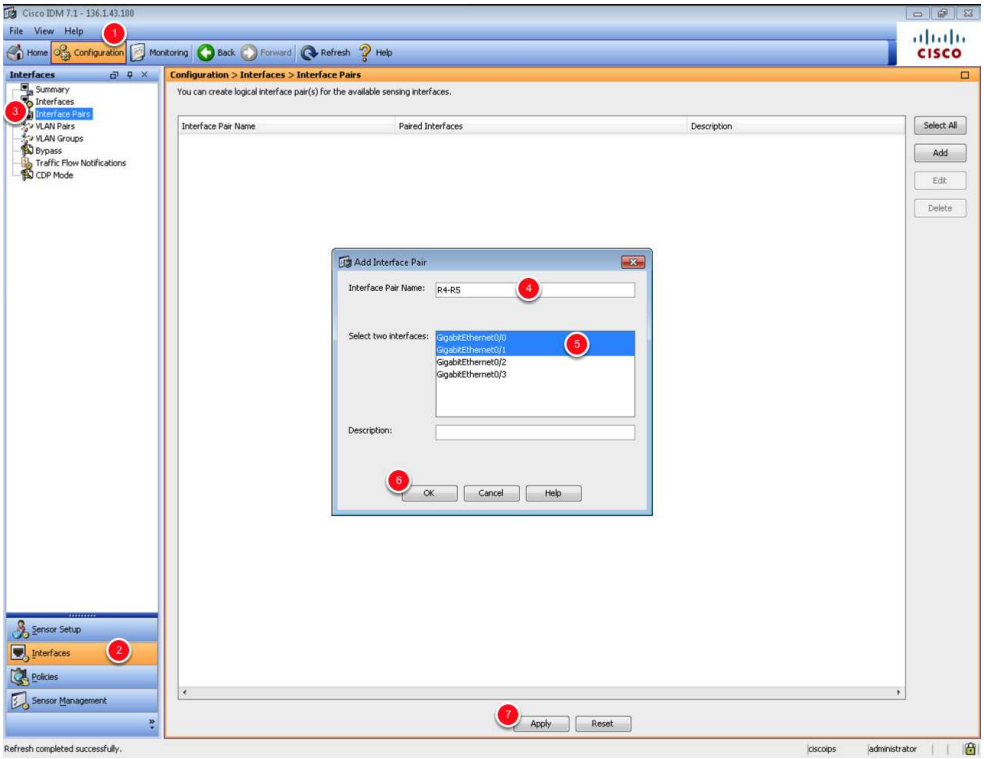
```
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
!
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 12907 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# retired false
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine application-policy-enforcement-ftp
sensor(config-sig-sig-app)# event-action produce-alert
sensor(config-sig-sig-app)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# signatures 12678 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# retired false
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine application-policy-enforcement-http
sensor(config-sig-sig-app)# event-action produce-alert
sensor(config-sig-sig-app)# exit
sensor(config-sig-sig)# exit
sensor(config-sig)# application-policy
sensor(config-sig-app)# http-policy
sensor(config-sig-app-htt)# http-enable true
sensor(config-sig-app-htt)# exit
sensor(config-sig-app)# ftp-enable true
sensor(config-sig-app)# exit
sensor(config-sig)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

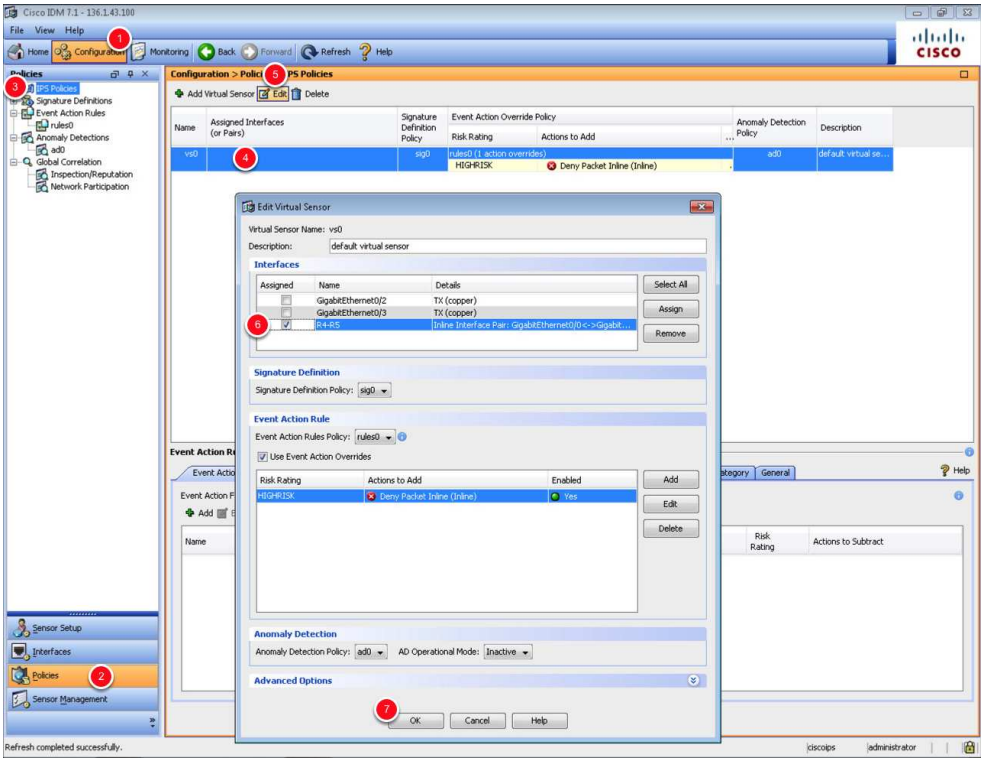


Create the inline interface-pair:

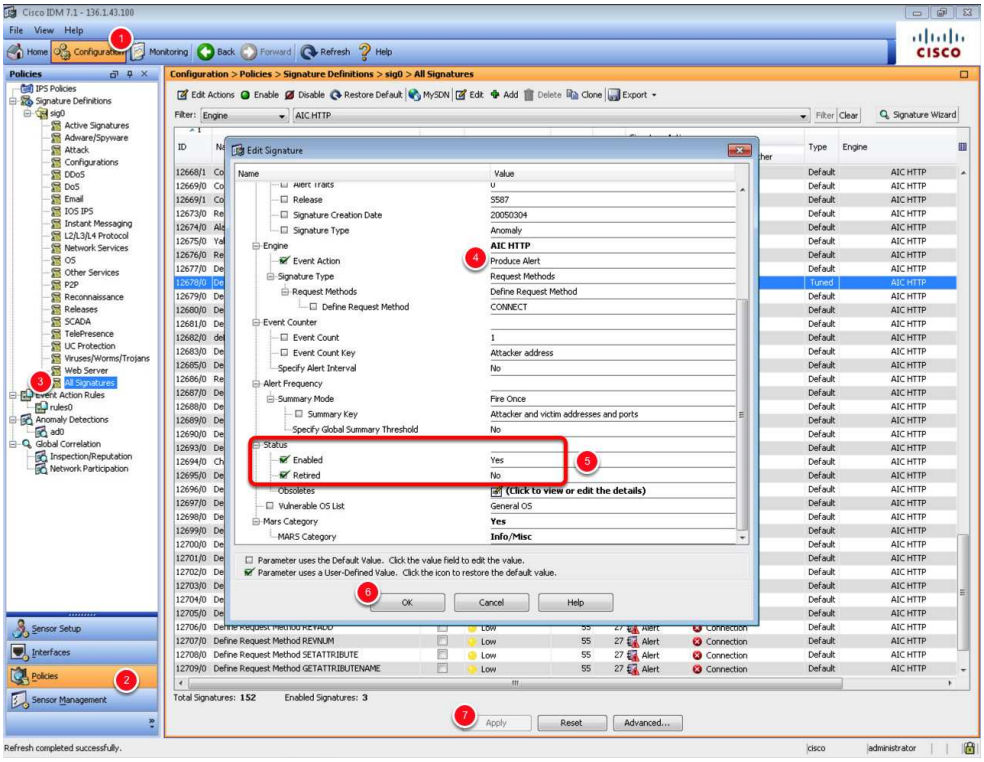


Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets

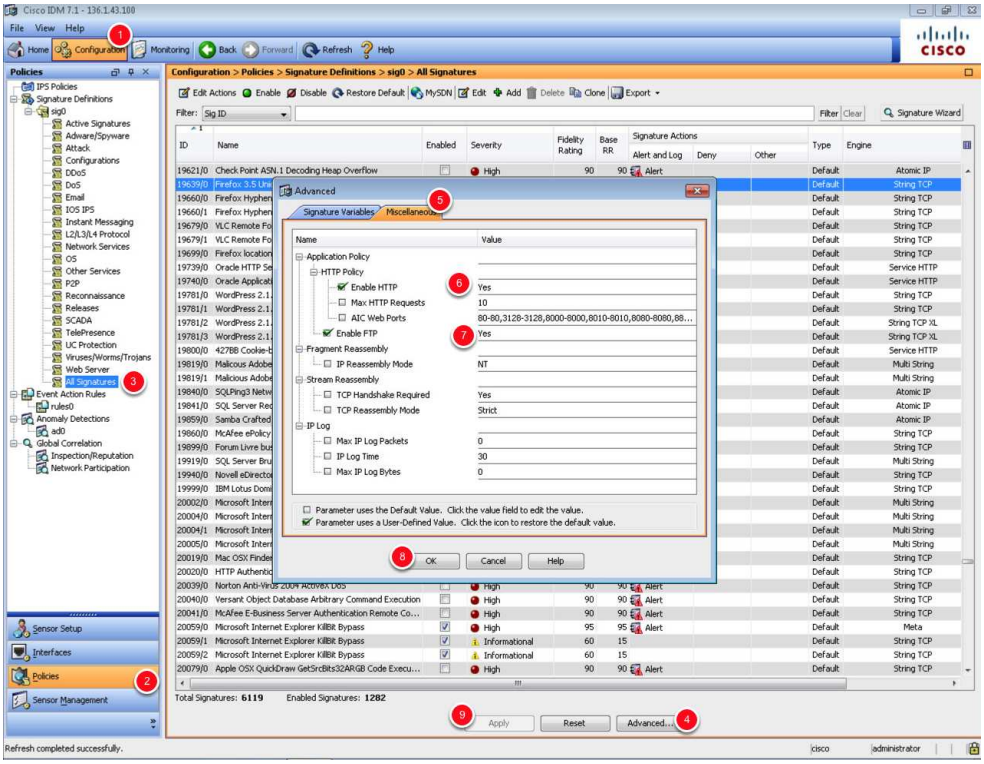


Activate and enable the HTTP signature for the `CONNECT` method and remove the `deny connection inline` action (perform the same for the signature matching the `FTP DELE` command):



Enable the HTTP and FTP AIC inspection:

https://t.me/learningnets



Verification

Initiate a telnet session from R5 to R4 on the HTTP port and simulate a `CONNECT` request; the connection will be closed by R4, not by the IPS, which will only log:

```
R5#telnet 150.1.4.4 80
Trying 150.1.4.4, 80 ... OpenCONNECT 150.1.4.4
HTTP/1.1 400 Bad Request
Date: Sat, 02 Nov 2013 20:31:21 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.4.4 closed by foreign host]
!
!sensor# show events


evIdsAlert: eventId=1376328122518001280 severity=low vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 01:32:27 2003/02/01 01:32:27 UTC
  signature: description=Define Request Method CONNECT id=12678 created=20050304 type=anomaly version=S587
    subsigId: 0
```

https://t.me/learningnets

```
sigDetails: Define Request Method CONNECT

  interfaceGroup: vs0
  vlan: 0
  participants:
    attacker:
      addr: locality=OUT 136.1.45.5
      port: 46333
    target:
      addr: locality=OUT 150.1.4.4
      port: 80
      os: idSource=unknown relevance=relevant type=unknown
  context:
    fromAttacker:
000000  31                                                1
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 37
  threatRatingValue: 37
  interface: ge0_1
  protocol: tcp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Meta Engine

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Configure IPS to allow communication between R4 and R5 using its Gi0/0 and Gi0/1 interfaces according to the diagram.
- Configure the IPS to generate an additional alarm when there is a group of ICMP Echo Requests of 10 packets generated from R5 toward R4's Loopback0 interface within 3 seconds, followed by a telnet session from the same IP that fails to authenticate at least two times in 60 seconds.
- These two events should occur sequentially within at least 120 seconds.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command **`erase current-config`** .

Overview

The Meta engine monitors sequences of the events generated by other IPS

https://t.me/learningnets

signatures and not the actual network patterns, hence the name. A signature based on this engine could be configured to trigger an event when there is a series of other events seen—for example, if signature ID 2000 is fired after signature ID 2004 within a 10-second interval. Every Meta signature has an ordered list of other IPS signatures associated with it. The Meta signature fires when all the listed signatures are triggered within a time interval. You may set whether the order of triggering the listed signatures is important. Additionally, a time interval known as Meta Reset Interval is provided to control the amount of time required to trigger all the listed signatures.

It is important to understand the concept of the Meta Key, which is similar to the keys we have seen for event summarization/counting. The Meta Key defines the mask that selects consecutive events, effectively storing them in the persistent storage for the duration of the Meta Reset Interval. For example, if the mask is **Axxx** , meaning attacker address, then only the events generated by the same source IP address are considered to be at the sequence. You may change the key to **xxBx**, for instance, if you want to track events toward the same victim address and you don't care for the source address.

In this scenario, we start by creating two custom signatures. The first one is cloned from ICMP Echo Request signature (ID 2004) and tuned to have an event-count of 10 and an alert-interval of 3 seconds. After this, a new custom **String.TCP** signature is created to match the regular expression **Bad Passwords**, which is reported by an IOS router not configured for AAA when the authentication fails three times. We set event count to 2 and alert-interval to 60 for this signature. After the two custom signatures have been created, we create new Meta signature that lists those two custom signatures in the order required with the time-interval for all events to be triggered set to 120 seconds.

Configuration

All signatures are configured only from GUI for this task.

```
SW3:
interface GigabitEthernet1/0/8
 switchport mode access
 switchport access vlan 45
 spanning-tree portfast
 no shutdown
!
interface GigabitEthernet1/0/9
 switchport mode access
 switchport access vlan 54
 spanning-tree portfast
```

https://t.me/learningnets

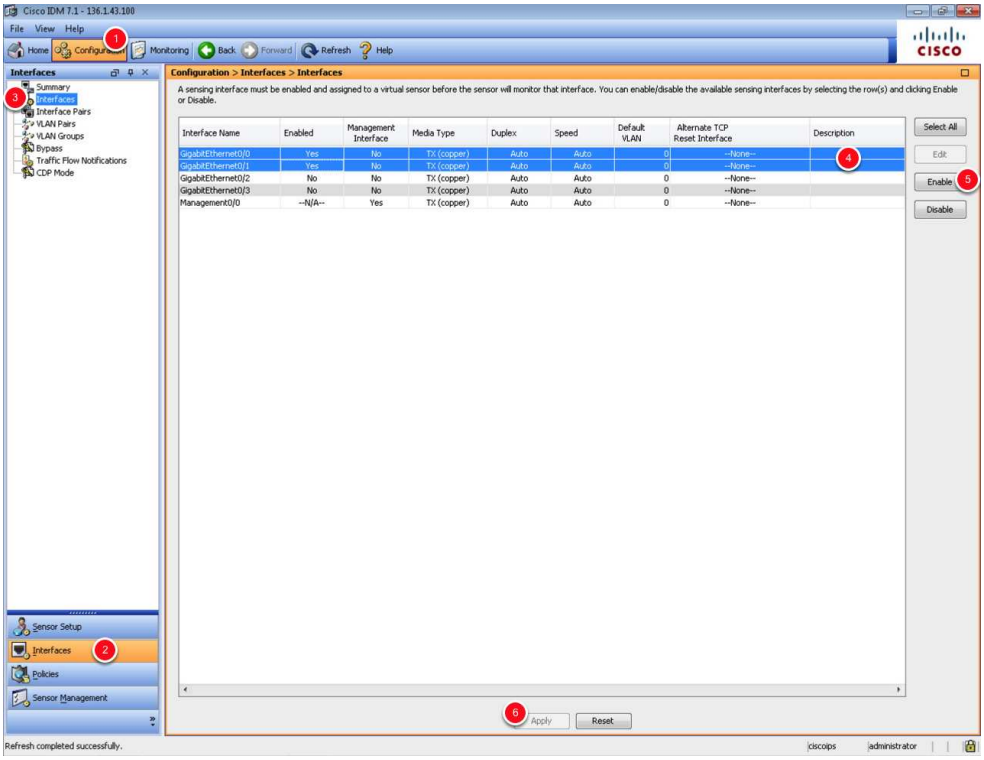
```
 no shutdown
!
interface GigabitEthernet1/0/2
 switchport access vlan 54
```

IPS:

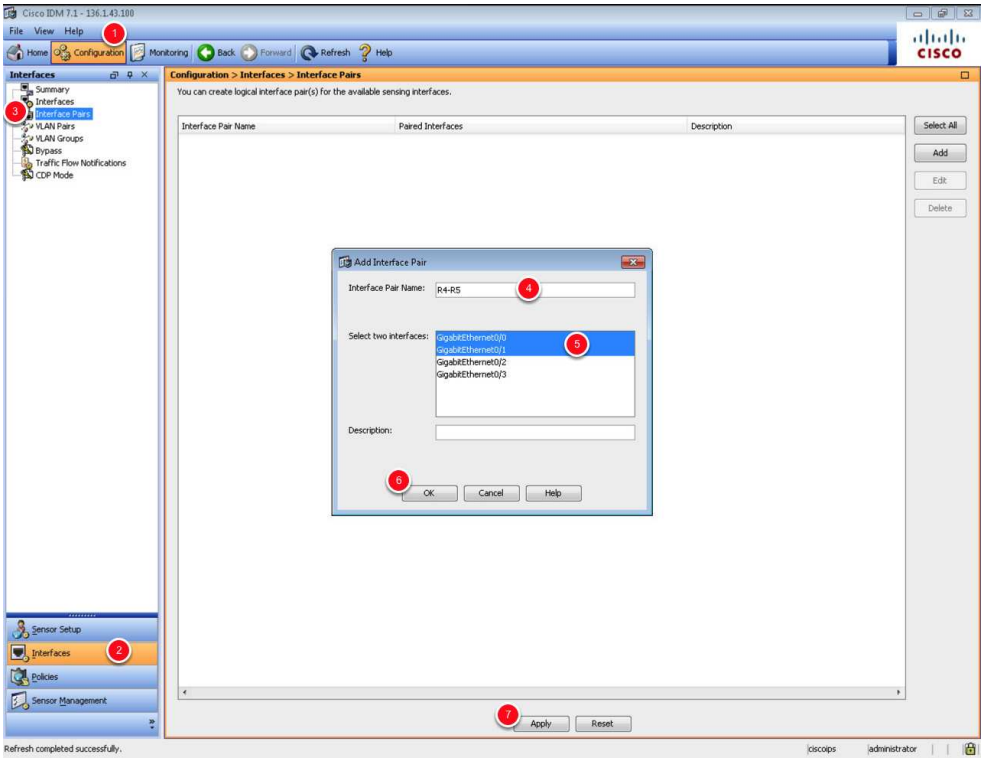
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/0
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# inline-interfaces R4-R5
sensor(config-int-inl)# interface1 gigabitEthernet0/0
sensor(config-int-inl)# interface2 gigabitEthernet0/1
sensor(config-int-inl)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# logical-interface R4-R5
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first enable the two interfaces:

https://t.me/learningnets

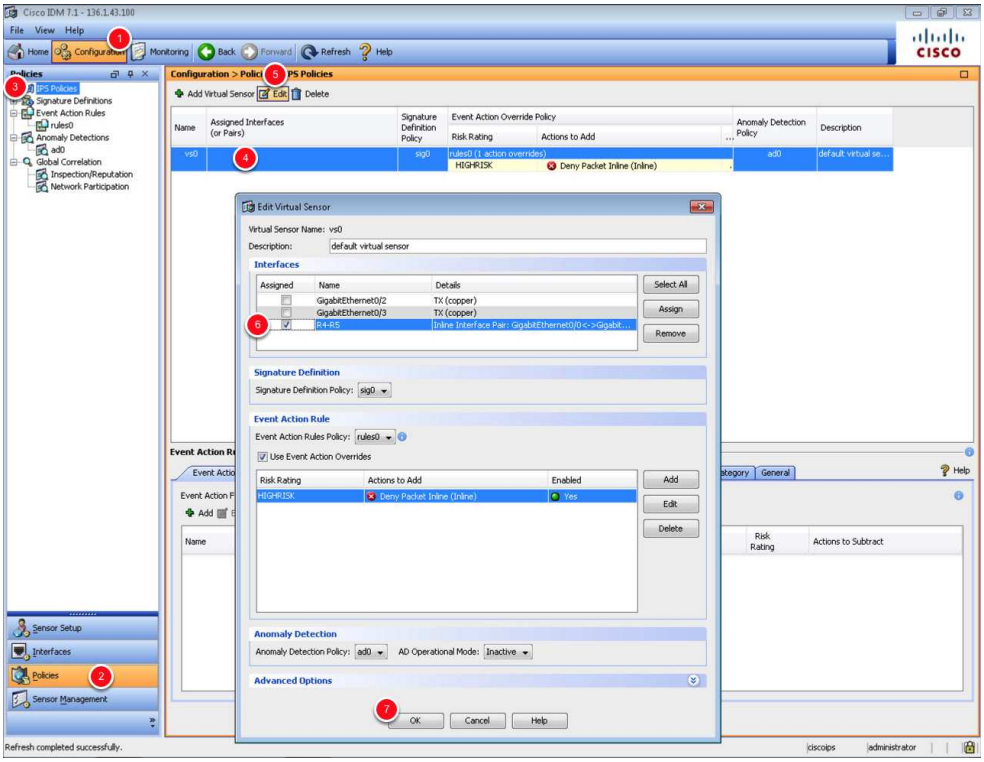


Create the inline interface-pair:

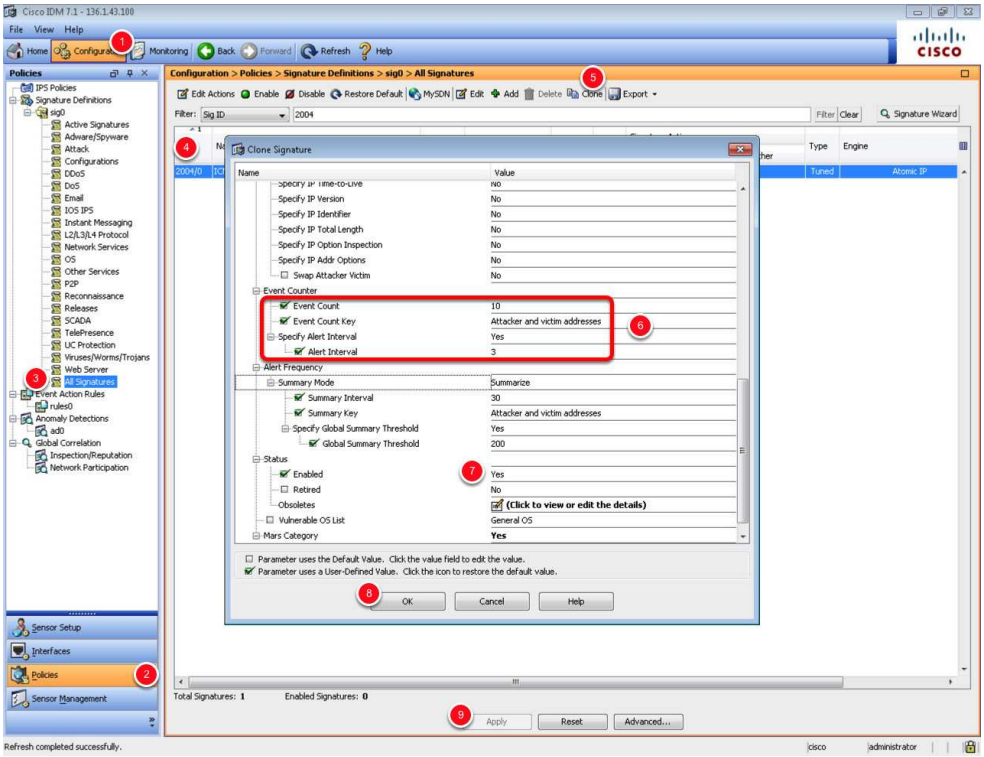


Assign the interface-pair to the default virtual sensor:

https://t.me/learningnets

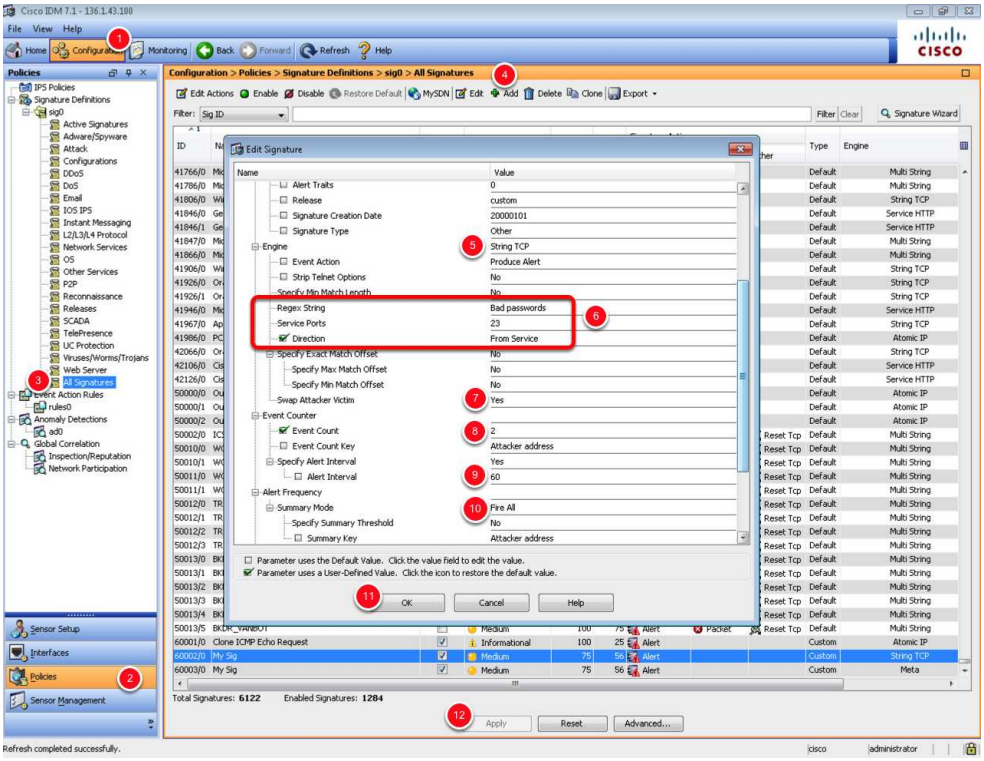


Clone the ICMP ECho Request signature, enable it, and configure it:

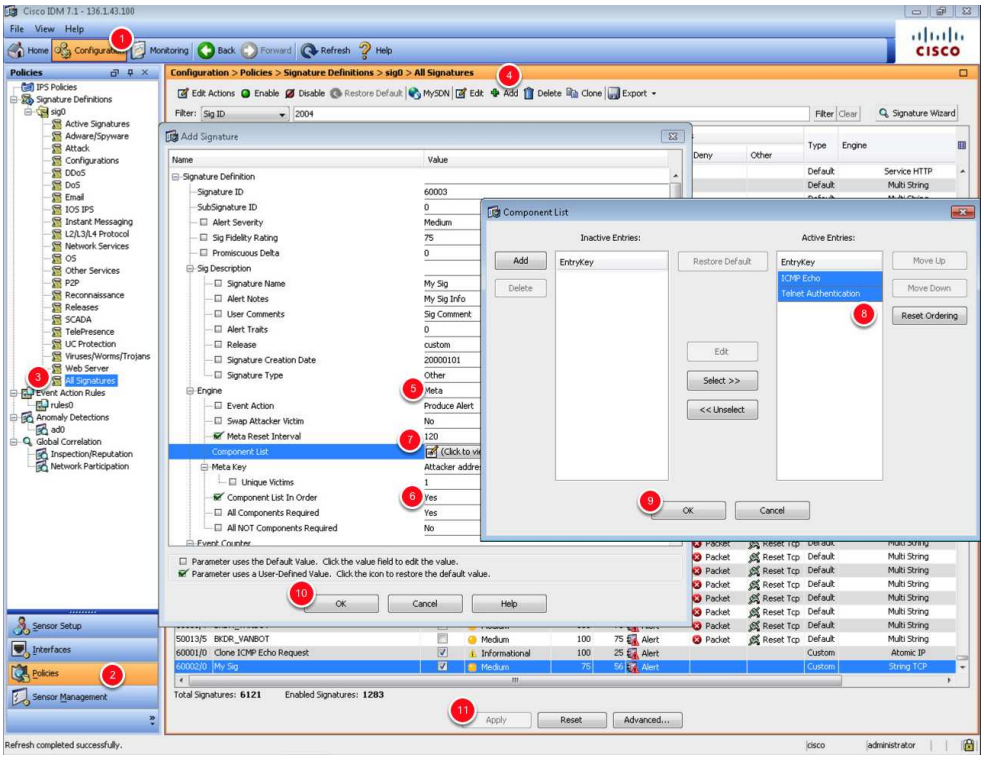


Create a custom **String.TCP**. Notice the event counting settings, the regular expression, and the service port. Ensure that the direction is set to **From Service** because the router will send the warning message in response to the client. Additionally, make sure that **Summary Mode** is set to **Fire All** and **swap attacker victim** is enabled:

https://t.me/learningnets



Finally, create the Meta signature, set the Engine to Meta, and set Meta Reset Interval to 120 seconds. Click the Component List and add the previously configured signatures to the list, in the correct order:



Verification

https://t.me/learningnets

Initiate traffic to trigger the first signature, and note the attacker address:

```
R5#ping 150.1.4.4 repeat 10

Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 1/2/4 ms
!
!sensor# show events

evIdsAlert: eventId=1376328122518001348 severity=informational vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 02:40:20 2003/02/01 02:40:20 UTC
signature: description=Clone ICMP Echo Request id=60001 created=20001127 type=other version=S666
    subsigId: 0
    sigDetails: ICMP Echo Request
  interfaceGroup: vs0
  vlan: 0
  participants:  attacker:
addr: locality=OUT 136.1.45.5

    target:
      addr: locality=OUT 150.1.4.4
      os: idSource=unknown relevance=relevant type=unknown
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 35
  threatRatingValue: 35
  interface: ge0_1
  protocol: icmp
```

Initiate traffic to trigger the second signature, which will also fire the meta signature, and note that the attacker is the same as a result of swapping attacker with victim on the second signature:

```
R5#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

Password:
Password:
Password:  % Bad passwords
```

https://t.me/learningnets

```
[Connection to 150.1.4.4 closed by foreign host]

!
!R5#telnet 150.1.4.4
Trying 150.1.4.4 ... Open



User Access Verification

Password:
Password:
Password:  % Bad passwords
[Connection to 150.1.4.4 closed by foreign host]
!
! sensor# show events

evIdsAlert: eventId=1376328122518001367 severity=medium vendor=Cisco alarmTraits=32768
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 02:50:45 2003/02/01 02:50:45 UTC
signature: description=My Sig id=60002 created=20000101 type=other version=custom
    subsigId: 0
    sigDetails: My Sig Info
  interfaceGroup: vs0
  vlan: 0
  participants:  attacker:
addr: locality=OUT 136.1.45.5
      port: 19228
    target:
      addr: locality=OUT 150.1.4.4
      port: 23
      os: idSource=unknown relevance=relevant type=unknown
  context:
    fromTarget:
000000  FF FB 01 FF FB 03 FF FD  18 FF FD 1F 0D 0A 0D 0A  ................
000010  55 73 65 72 20 41 63 63  65 73 73 20 56 65 72 69  User Access Veri
000020  66 69 63 61 74 69 6F 6E  0D 0A 0D 0A 50 61 73 73  fication....Pass
000030  77 6F 72 64 3A 20 FF FE  20 FF FD 21 FF FA 21 00  word: .. ..!..!.
000040  FF F0 FF FE 18 0D 0A 50  61 73 73 77 6F 72 64 3A  .......Password:
000050  20 0D 0A 50 61 73 73 77  6F 72 64 3A 20 0D 0A 25   ..Password: ..%
000060  20 42 61 64 20 70 61 73  73 77 6F 72 64 73 0D 0A   Bad passwords..
    fromAttacker:
000000  FF FD 03 FF FB 20 FF FB  1F FF FB 21 FF FD 01 FF  ..... .....!....
000010  FC 18 FF FA 1F 00 50 00  18 FF F0 FF FC 20 67 66  ......P...... gf
```

https://t.me/learningnets

```
000020  0D 0A 65 77 0D 0A 66 64  0D 0A                    ..ew..fd..
  riskRatingValue: attackRelevanceRating=relevant targetValueRating=medium 66
  threatRatingValue: 66
  interface: ge0_0
  protocol: tcp


evIdsAlert: eventId=1376328122518001368 severity=medium vendor=Cisco
  originator:
    hostId: sensor
    appName: sensorApp
    appInstanceId: 1206
  time: 2003/02/01 02:50:45 2003/02/01 02:50:45 UTC
signature: description=My Sig id=60003 created=20000101 type=other version=custom
    subsigId: 0
    sigDetails: My Sig Info
  interfaceGroup: vs0
  vlan: 0
  participants:  attacker:
addr: locality=OUT 136.1.45.5
alertDetails: Component Signature List: 60001.0 60002.0  ;

  riskRatingValue: targetValueRating=medium 56
  threatRatingValue: 36
  interface: ge0_0
  protocol: tcp
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 3: Intrusion Prevention and Content Security

IPS Anomaly Detection

Task

- Configure the management IP address to be 136.1.43.100/24, and set the default gateway to 136.1.43.9; ensure that HTTP management access is allowed only from VLAN 37.
- Configure anomaly detection as follows:
  - Add VLANs 32 and 36 and Loopback0 addresses of all routers to the **Internal Zone**
  - Add the /16 subnets for 136.1.0.0/16 and 150.1.0.0/16 to the **Illegal Zone**.
  - Change the KB rotation interval to 12 hours and set the Scanner Threshold for TCP port 445 to 500 in the **Internal Zone**.
  - When the histogram threshold is violated, scanners should be identified as worms for no longer than 120 seconds.
  - A user from the IP 172.16.20.100 often scans the network for new vulnerabilities; these scans should not be interpreted as attacks.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 3 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA or IPS.
>
> To restore the IPS to factory defaults, erase the running configuration using the command `erase current-config` .

https://t.me/learningnets

Overview

Anomaly Detection, or AD, was introduced in IPS 6.0 code. It allows for detection of network worms based on traffic anomalies introduced by worms activity. This idea contrasts with the regular signature-based detection technique, where an alert is generated upon matching a pattern in network traffic. Anomaly Detection uses the fact that some network worms (such as Slammer, Blaster, Code Red) actively employ network scanning to find new victims to infect. Other network worms, such as those that propagate via email or passively learn their targets by sniffing on network traffic, are not detected. Only active scanning behavior is detectable by the AD policy. However, this allows for discovery of potentially unknown viruses based solely on their behavior.

In short, AD detection is based on learning the network traffic patterns and detecting any deviations from the learned parameters. The learned traffic profile is known as Knowledge Database, or KB. The learning phase starts when you initially boot the IPS sensor and continues for 24 hours by default. When the learning has completed, the IPS moves into Detect mode and starts looking for potential scanners. A scan is detected as an IP address range sweep, which generates a lot of traffic in one direction, from attacker to victims. Commonly, the amount of responding traffic is significantly less than the number of scanning requests, because many hosts might be down. During Detect mode, the IPS continues to learn the network traffic patterns, and every 24 hours by default the KB is rotated and replaced with the new learned information. You may change the learning intervals and settings by navigating to **Configuration > Policies > Anomaly Detections > ad0 > Learning Accept Mode**; you may also change whether to accept the new KB every periodic interval automatically, or perform the replacement manually.

To decrease the amount of false positives, the IPS partitions the network into three zones, based on the destination IP addresses:

- **Internal Zone** is empty by default; it should be configured by the administrator to cover the IP address ranges used by the network. This should include only the allocated and used IP addresses.
- **Illegal Zone** is empty by default; it should be configured by the administrator and should cover the destination IP addresses that should not appear in legitimate traffic. This includes the portions of unallocated internal address space, loopback and link-local subnets, and so on. Configuring this zone is very important because it allows for reducing the risk of false positives; a worm would scan all addresses and eventually start sending probes to the illegal zone.

https://t.me/learningnets

- **External Zone** contains all IP addresses by default; all destination IP addresses are classified as belonging to this zone, unless they match the **Illegal** or **Internal** zone ranges.

The knowledge database has the following hierarchical format:

```
KB Name (KBs are rotated)
   Zone Name (Internal, External, Illegal)
     Protocol Name (TCP, UDP, Other)
        Service (Port Number/Protocol Number)
            Scanner Threshold: Value
            Histogram
                High: Number of Src IPs
                Medium: Number of Src IPs
                Low: Number of Src IPs
```

IPS collects information on every protocol and port (service) actively used in the network. For every service (destination port number or IP protocol number for "Other" protocols), the database contains the learned maximum scanner threshold: maximum number of hosts scanned in this zone, on this service. The histogram attribute describes aggregated scanning behavior, which is the maximum number of different source IP addresses scanning High amount of the zone hosts, the maximum number of different source IP addresses scanning Medium amount of hosts, and the same value for the Low amount of hosts. The High, Medium, and Low quantities are defined internally in the IPS and are relative to the total amount of active IP addresses in the zone. The concept of histogram allows for more granular profile of traffic flows toward the particular zone. All these parameters are learned by the IPS automatically, and you should only change them if you have a very good understanding of your network traffic patterns.

The detection of a worm-infected host happens as follows:

- The host is identified as scanner toward the particular zone, when the amount of the unidirectional requests exceeds the configured or discovered scanner threshold. However, this does not yet classify the source of the scans as a worm-infected host.
- When the IPS detects the histogram thresholds violation, it assumes that the amount of infected hosts has increased, and there is potentially a worm on the network. A histogram threshold is violated when the amount of different source IP addresses scanning High, Medium, or Low amount of hosts in the zone exceeds the value in the histogram. Notice that this event does not reveal the actual scanner IP addresses, because this is an aggregate event.

https://t.me/learningnets

- After the histogram violation has been detected for a zone/service, the IPS lowers the scanning threshold for this service to the histogram threshold value that has been violated. This automatic adjustment of the scanner threshold allows for discovery of actual scanners, which are now classified as worms. This adjustment lasts for the number of seconds specified in the Worm Timeout interval. You may change this timeout by navigating to **Configuration > Policies > Anomaly Detections > ad0 > Operation Settings**; under the same tab, you may configure the lists of source/destination IP addresses that will not trigger an alarm when identified as scanners or worm-infected hosts. This might be the case if there are network workstations that need running network scanning software on periodic basis.

IPS will only declare scanners as worm-infected when the histogram thresholds have been exceeded. If the histogram thresholds are not crossed, any host exceeding the scanner threshold is reported as simply a scanner. To allow for events reporting, the IPS supports a group of signatures, in range 13000-13008, where every signature has two sub-signatures: 0 and 1. The first sub-signature reports the attacker as a scanner toward the particular zone/protocol (three zones, three protocols, resulting in 9 signatures). The second sub-signature reports the attacker as a worm-infected host, and only triggers when the histogram has been violated for the particular zone. Basic AD configuration consists of setting the zones and possibly tuning some thresholds. However, AD service is designed to be operating in automatic mode, with minimal tuning.

Configuration

IPS:

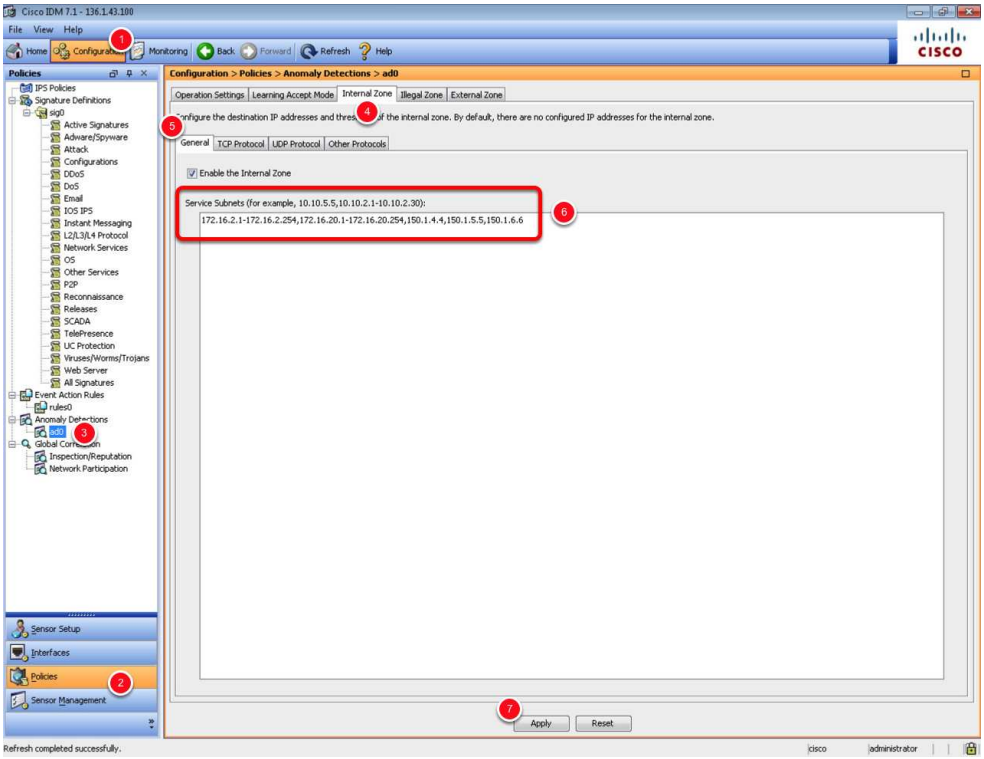
```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.43.100/24,136.1.43.9
sensor(config-hos-net)# access-list 136.1.37.0/24
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
Warning: DNS or HTTP proxy is required for global correlation inspection and reputation filtering, but no DNS or pro
!
sensor(config)# service anomaly-detection ad0
sensor(config-ano)# internal-zone
sensor(config-ano-int)# ip-address-range 172.16.2.1-172.16.2.254,172.16.20.1-172.16.20.254,150.1.4.4,150.1 5.5,150.1
sensor(config-ano-int)# tcp
sensor(config-ano-int-tcp)# dst-port 445
sensor(config-ano-int-tcp-dst)# enabled true
sensor(config-ano-int-tcp-dst)# override-scanner-settings yes
```

https://t.me/learningnets

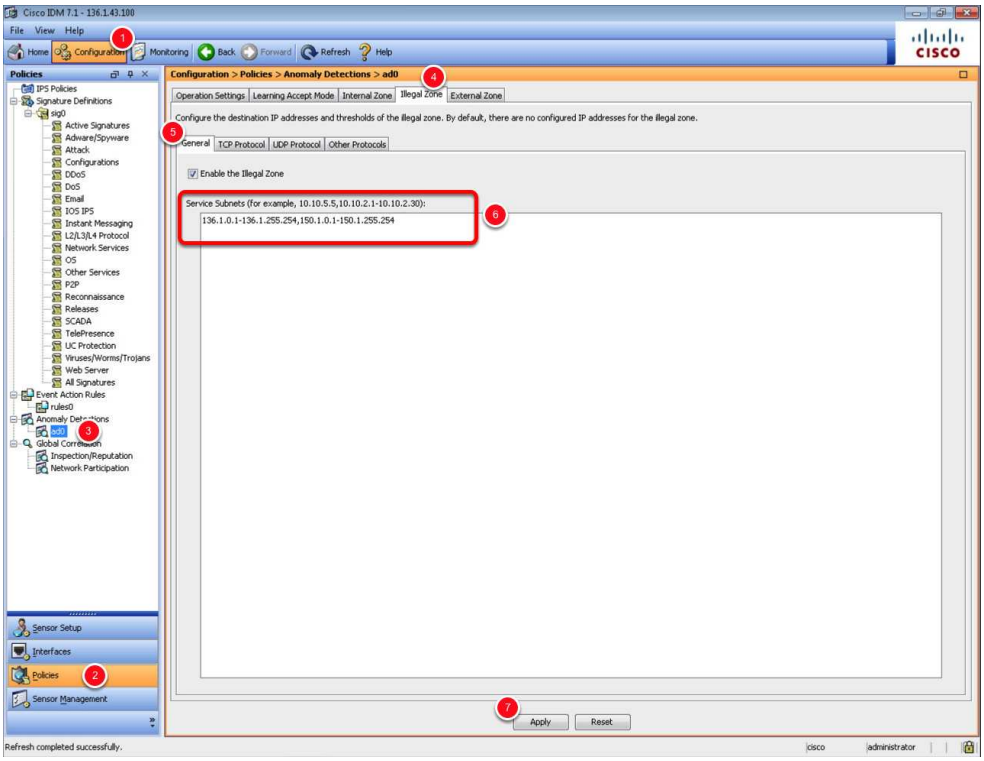
```
sensor(config-ano-int-tcp-dst-yes)# scanner-threshold 500
sensor(config-ano-int-tcp-dst-yes)# exit
sensor(config-ano-int-tcp-dst)# exit
sensor(config-ano-int-tcp)# exit
sensor(config-ano-int)# exit
sensor(config-ano)# illegal-zone
sensor(config-ano-ill)# ip-address-range 136.1.0.1-136.1.255.254,150.1.0.1-150.1.255.254
sensor(config-ano-ill)# exit
sensor(config-ano)# learning-accept-mode auto
sensor(config-ano-aut)# schedule periodic-schedule
sensor(config-ano-aut-per)# interval 12
sensor(config-ano-aut-per)# exit
sensor(config-ano-aut)# exit
sensor(config-ano)# worm-timeout 120
sensor(config-ano)# ignore
sensor(config-ano-ign)# source-ip-address-range 172.16.20.100
sensor(config-ano-ign)# exit
sensor(config-ano)# exit
Apply Changes?[yes]:
!
sensor(config)# service analysis-engine
sensor(config-ana)# virtual-sensor vs0
sensor(config-ana-vir)# anomaly-detection
sensor(config-ana-vir-ano)# operational-mode learn
sensor(config-ana-vir-ano)# exit
sensor(config-ana-vir)# exit
sensor(config-ana)# exit
Apply Changes?[yes]:
```

To configure the same settings from GUI using IDM, first configure the **Internal Zone** :

https://t.me/learningnets

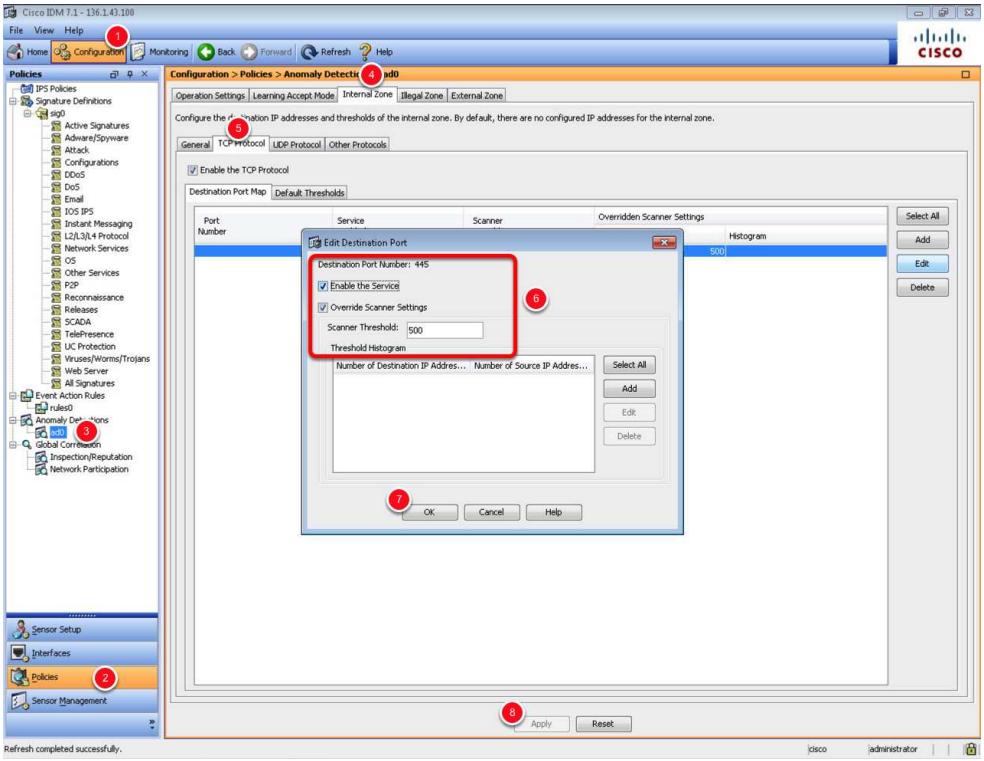


Configure the **Illegal Zone**:

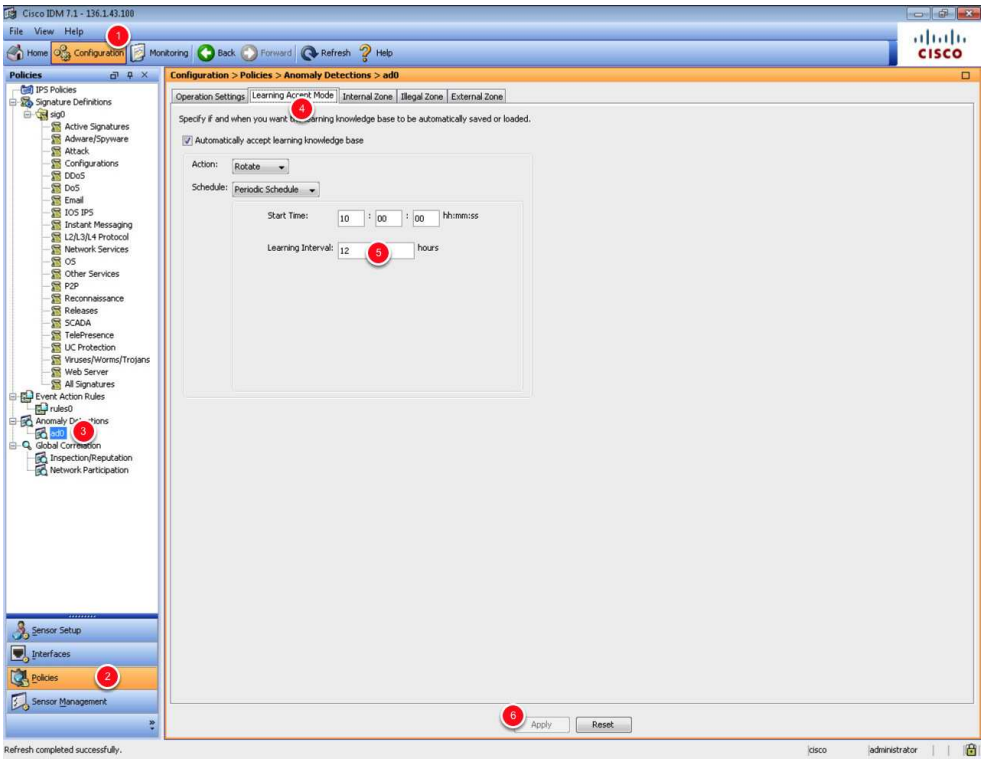


For the **Internal Zone**, add a new TCP service and modify the scanner threshold:

https://t.me/learningnets

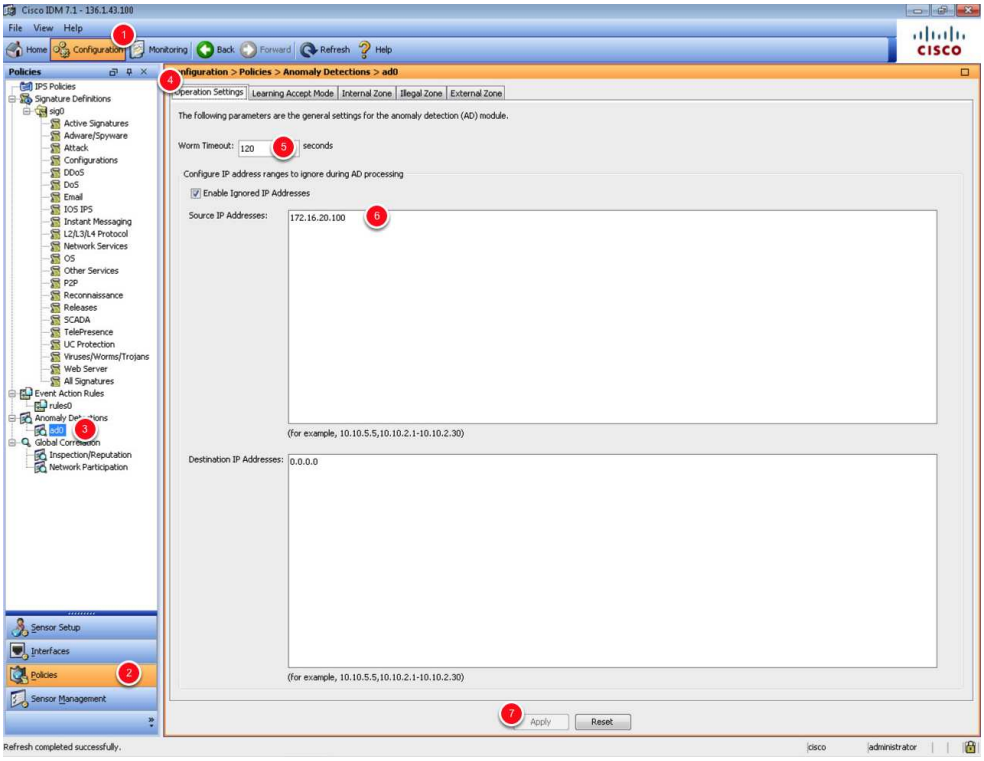


Configure the KB rotation interval value to 12 hours:



Set the worm timeout to 120 seconds and configure the IP address of 172.16.20.100 so that it is ignored by AD processing:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

Section 4 Introduction

> **This section utilizes Firewall Services.**

Lab Topics

Included in the CCIE Security Lab Exam blueprint are the following high-level Identity Management topics:

- AAA Framework
  - AAA Model
  - RADIUS Protocol and Attributes
  - TACACS+ Protocol and Attributes
  - LDAP and AD Integration
- Device Administration
  - Authentication
  - Command and Shell Authorization
  - Command Accounting
- Network Access (TrustSec Model)
  - MAB
  - 802.1X
  - Web Authentication
  - RADIUS VSAs
  - Proxy Authentication
- Cisco Identity Services Engine (ISE)
  - Profiling Configuration (Probes)
  - Guest Services
  - Posture Assessment
  - Client Provisioning (CPP)

These topics are further explained in the following checklist.

https://t.me/learningnets

Technology	Mastered?
Understanding the AAA Framework	
Understanding the RADIUS Protocol	
Understanding RADIUS Attributes (Cisco AV-PAIRS)	
Understanding the TACACS+ Protocol	
Understanding TACACS+ Attributes	
Comparison of RADIUS and TACACS+	
Configuring Basic LDAP Support	
Overview of Cisco Secure ACS	
How to Navigate Cisco Secure ACS	
Cisco Secure ACS – Network Settings Parameters	
Cisco Secure ACS – User Settings Parameters	
Cisco Secure ACS – Group Settings Parameters	
Cisco Secure ACS – Shared Profiles Components (802.1X, NAF, NAR, Command Author, Downloadable ACL, etc.)	
Cisco Secure ACS – Shell Command Authorization Sets Using Both Per-Group Setup and Shared Profiles	

https://t.me/learningnets

Technology	Mastered?
Cisco Secure ACS – System Configuration Parameters	
Enabling AAA on a Router for vty Lines	
Enabling AAA on a Switch for vty Lines	
Enabling AAA on a Router for HTTP	
Enabling AAA on the Cisco ASA Security Appliance for Telnet and SSH Protocols	
Using Default vs. Named Method Lists	
Complex Command Authorization and Privilege Levels, and Relevant Cisco Secure ACS Profiles	
Proxy Service Authentication and Authorization on the Cisco ASA Security Appliance for Pass-Through Traffic (FTP, Telnet, and HTTP), and Relevant Cisco ISE Profiles	
Using Virtual Telnet on the Cisco ASA Security Appliance	
Using Virtual HTTP on the Cisco ASA Security Appliance	
Downloadable ACLs	
AAA 802.1X Authentication Using RADIUS on a Switch	
NAC-L2-802.1X on a Switch	

https://t.me/learningnets

Technology	Mastered?
NAC-L2-IP on a Switch	
Troubleshooting Failed AAA Authentication or Authorization	
Troubleshooting Using Cisco Secure ACS Logs	
Cisco Identity Services Engine Configuration and initialization	
ISE authZ result handling	
ISE Profiling Configuration (Probes)	
ISE Guest Services	
ISE Posture Assessment	
ISE Client Provisioning (CPP)	
ISE Configuring AD Integration/Identity Sources	
ISE support for 802.1x	
ISE MAB support	
ISE Web Auth support	
ISE definition and support for VSAs	
Support for MAB in Cisco IOS	

https://t.me/learningnets

Technology	Mastered?
Support for Web Auth in Cisco IOS	
Using the test aaa Command on the Router, Switch, or Cisco ASA Security Appliance	
Understanding and Interpreting the debug radius Command	
Understanding and Interpreting the debug tacacs+ Command	
Understanding and Interpreting the debug aaa authentication Command	
Understanding and Interpreting the debug aaa authorization Command	
Understanding and Interpreting the debug aaa accounting Command	

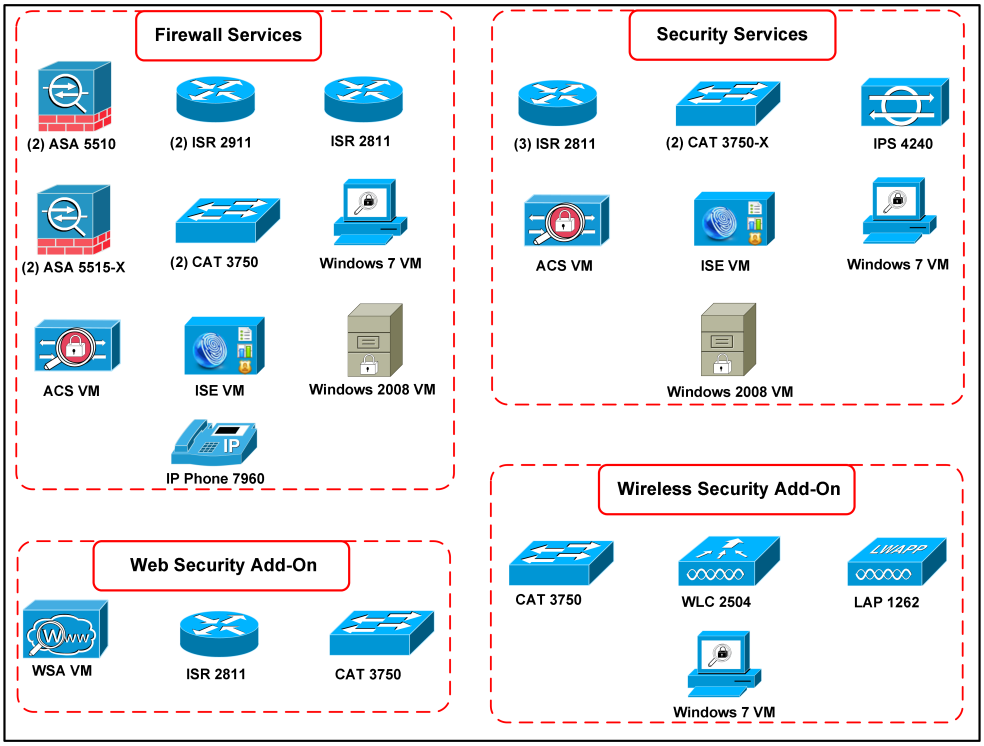
Configurations

Initial configuration files are provided at the top of this page. The initial configuration in most cases is the configuration that the previous lab ended with. Specific instructions will appear at the beginning of a lab if it is necessary to load or modify the existing configurations.

Section 4 General Lab Topology

The INE CCIE Security Racks are divided into modules as shown in the figure below. This section, Section 4, utilizes Firewall Services and the Wireless Security Add-0n.

https://t.me/learningnets



Topology diagrams for each task can be found by clicking the **Resources** button.

Section 4 General Guidelines

The following should be considered when working with VMware images in this section:

- ISE and ACS are VMware appliances configured as non-persistent, so if they are rebooted from the Rack Rental Control Panel, all configured changes from both CLI and GUI are lost.
- TEST-SRV and TEST-PC are running on VMware as non-persistent images, so if they are rebooted from the Rack Rental Control Panel, all configuration changes are lost.
- You should only reboot/reset the VMware machines as a last resort, if machines become unresponsive and there is no other way of recovering them.
- ACS and ISE are initialized and pre-configured; tasks are created so that there is no need to change any CLI parameters that might bring it into an unresponsive state.
- ACS might require a license to be installed first when GUI is accessed; the license file resides in the **Lab Files -> Cisco ACS License** folder, which can be found on the desktop of both TEST-SRV and TEST-PC.
- There is no need to change any IP addressing settings on TEST-SRV or TEST-PC.

The following IP addressing information is used for the VM images:

https://t.me/learningnets

Device	IP Address	Default Gateway	DNS Server	NTP Server
ACS	172.16.1.100	172.16.1.1	172.16.10.100	172.16.1.1
ISE	172.16.3.100	172.16.3.1	172.16.10.100	172.16.3.1
TEST-SRV-A	172.16.10.100	172.16.10.1	172.16.10.100	172.16.10.100
TEST-PC-A	DHCP	DHCP	DHCP	N/A

Section 4 ISE and ACS Guidelines

Before starting any ACS- or ISE-related tasks, connect to the console/CLI and ensure that all services are running. To verify it, as well as the output you should be looking for, you need to ensure that all services are in a **running** state:

```
ACS1/admin# show application status acs

ACS role: PRIMARY
Process 'database' running
Process 'management' running
Process 'runtime' running
Process 'view-database' running
Process 'view-jobmanager' running
Process 'view-alertmanager' running
Process 'view-collector' running
Process 'view-logprocessor' running

ISE2/admin# show application status ise
ISE Database listener is running
, PID: 4114
ISE Database is running, number of processes: 27 ISE Application Server is running
, PID: 5664 ISE M&T Session Database is running
, PID: 3774 ISE M&T Log Collector is running
, PID: 6000 ISE M&T Log Processor is running
, PID: 6109 ISE M&T Alert Process is running
, PID: 5901
% WARNING: ISE DISK SIZE NOT LARGE ENOUGH FOR PRODUCTION USE
```

https://t.me/learningnets

```
% RECOMMENDED DISK SIZE: 200 GB, CURRENT DISK SIZE: 64 GB
```

The following output is not correct because some ACS/ISE services are not running, which results in GUI not being available or ACS/ISE not functioning properly:

```
ACS1/admin# show application status acs

ACS role: PRIMARY

Process 'database'                running Process 'management' running (HTTP is nonresponsive)
Process 'runtime'                 running
Process 'view-database'           running
Process 'view-jobmanager'         running
Process 'view-alertmanager'       running
Process 'view-collector'          running Process 'view-logprocessor' not monitored

ISE1/admin# show application status ise

ISE Database listener is running, PID: 4114
ISE Database is running, number of processes: 27 ISE Application Server is still initializing
.
ISE M&T Session Database is running, PID: 3774
ISE M&T Log Collector is running, PID: 6000
ISE M&T Log Processor is running, PID: 6109
ISE M&T Alert Process is running, PID: 5901
% WARNING: ISE DISK SIZE NOT LARGE ENOUGH FOR PRODUCTION USE
% RECOMMENDED DISK SIZE: 200 GB, CURRENT DISK SIZE: 64 GB
```

The above states for some ACS/ISE services are accessible only if the ACS/ISE VMware has been rebooted, because it might take as long as 15-30 minutes for it to properly initialize. However, if after 30 minutes not all services are in a **running** state, there are two options:

- Reset the VMware machine and wait for it to initialize.
- Manually stop and start ACS/ISE services using the **`application stop acs`**, **`application start acs`**, **`application stop ise`**, **`application start ise`** commands.

The following is an example of performing the stop/start function for ACS/ISE services:

```
ACS1/admin# application stop acs

Stopping ACS.
Stopping Management and View.....................
```

https://t.me/learningnets

```
Stopping Runtime.....................
Stopping Database.....
Cleanup....
!
!ACS1/admin# application start acs

Starting ACS ....

To verify that ACS processes are running, use the
'show application status acs' command.

ISE1/admin# application stop ise

Stopping ISE Monitoring & Troubleshooting Log Processor...
Stopping ISE Monitoring & Troubleshooting Log Collector...
Stopping ISE Monitoring & Troubleshooting Alert Process...
Stopping ISE Application Server...
Stopping ISE Monitoring & Troubleshooting Session Database...
Stopping ISE Database processes...
!
!ISE1/admin# application start ise



Starting ISE Database processes...
Starting ISE Monitoring & Troubleshooting Session Database...
Starting ISE Application Server...
Starting ISE Monitoring & Troubleshooting Alert Process...
Starting ISE Monitoring & Troubleshooting Log Collector...
Starting ISE Monitoring & Troubleshooting Log Processor...
Note: ISE Processes are initializing. Use 'show application status ise'
  CLI to verify all processes are in running state.
```

Additionally, we recommend that you use Internet Explorer or Chrome for configuring ACS and ISE via the graphical user interface. Although Firefox is available on TEST-SRV-A and TEST-PC-A, it is known for having incompatibility issues with either ACS or ISE.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license

https://t.me/learningnets

- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124
- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA Local Authentication

Task

- Configure ASA3 so that telnet sessions from VLAN37 subnet are authenticated using the password **CCIELAB**.
- Configure ASA3 so that HTTP sessions from VLAN27 are authenticated using the local username **HTTP-USER** with the password **HTTP-PASS**.
- Configure ASA3 so that console sessions are authenticated locally.
- Ensure that ASA3 requires the password **SECRET** to allow users configuration mode access.
- Ensure that user accounts failing authentication for three consecutive times are locked out, but the username **NO-LOCK** with the password **NO-LOCK** should not be affected by this setting.
- Configure the following set of users:
    - Username **NO-ASDM** with password **NO-ASDM** is not allowed to connect via ASDM to the ASA.
    - Username **ASDM-RO** with password **ASDM-RO** is allowed to connect via ASDM but not configure the ASA.
    - Username **CONSOLE** without a password is allowed to manage the ASA only via the console.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

Overview

By default, without any AAA-related configurations applied, the ASA behaves as follows for management sessions:

https://t.me/learningnets

- Does not require users to authenticate at the console, so it works in **no login mode**, just like IOS.
- Requires users to authenticate for telnet/SSH/HTTP, so it works in **login mode** asking for password authentication, just like IOS.
- If telnet sessions are allowed, ASA asks for the default password **cisco**, which can be changed with the command `passwd <string>` and sets the user into privilege level 1; to revert back to the default password, use the command `no passwd`. To obtain privilege level 15, use the `enable` command with a blank password, which is the default value but can be changed with the command `enable password <string>`. To revert back to the default, use the command `enable password`.
- If SSH is activated and sessions are allowed, it will still not work because SSH requires username/password-based authentication, so it needs AAA to be configured.
- If HTTP is activated and sessions are allowed, although ASDM will ask for username/password authentication, it uses enable password authentication; it will work using blank as the username and the enable password, which also defaults to blank.

Console access and telnet server are active by default, but SSH and HTTP services are not and must be enabled. By default, the ASA does not allow any telnet/SSH/HTTP sessions for management unless manually allowed from specific hosts/subnets and terminated on specified ASA interfaces, which is opposite of the IOS behavior, which requires you to manually restrict the sources from which management sessions are allowed. Also note that telnet sessions are not allowed to be terminated on the interfaces with the lowest security-level unless protected by an IPsec session (for example, if your ASA has 4 interfaces with security levels of 20, 20, 30, and 40, no clear-text telnet sessions are allowed on either interfaces with the security-level of 20).

Compared with IOS, ASA has no `aaa new-model` command, and AAA commands are active and available by default. With the local user database, the ASA supports only authentication and exec/command authorization for management sessions; command accounting is only supported using a remote TACACS server (RADIUS does not support command accounting). To enable authentication for console/telnet/ssh/http management sessions using the local username database, the following commands must be configured:

https://t.me/learningnets

```
aaa authentication serial console LOCAL
aaa authentication telnet console LOCAL
aaa authentication ssh console LOCAL
aaaa authentication http console LOCAL
```

To add users to the local database, one of the following commands can be used, the first one allowing users to have a blank password. The privilege-level is optional; if not configured, a default level of 2 is assigned to the user, which is different from IOS where a user has a default level of 1 because ASDM access requires a minimum level of 2.

```
username <string> nopassword privilege <0-15>
username <string> password <string> privilege <0-15>
```

When users are authenticated to console/telnet/ssh, they will be always be placed into privilege level 1, identified by the `>` character, regardless of the privilege level configured for the user, which is different from IOS. From here, there are three options to access the privilege exec mode:

- If the **`aaa authentication enable console LOCAL`** command has not been configured, use the **`enable`** command and authenticate using the global enable password to be assigned privilege level 15, regardless of the configured privilege level or service-type attribute configured at the user level.
- If the **`aaa authentication enable console LOCAL`** command has been configured, use the **`enable`** command and authenticate using the user's password to be assigned to the privilege-level assigned to the user level.
- Use the **`login`** command and authenticate again using the username/password combination to be assigned to the privilege level configured at the user level; if the user has a configured privilege level of 1, session remains at privilege level 1 and one of the first two options must be used to have access to privilege exec mode.

All commands on the ASA have a default privilege level assigned, just like on IOS. However, without command authorization enabled using the command **`aaa authorization command LOCAL`** , as long as the user's session has been placed into a privilege level higher than 1, it will be allowed to perform any commands on the ASA; basically, the ASA uses the privilege level to restrict user access to commands only if command authorization is enabled, which is different from IOS.

If the **`aaa authentication enable console LOCAL`** command is configured, users must use their own user password instead of the global enable password when using the **`enable`**

https://t.me/learningnets

command. Additionally, using the `enable` password will no longer assign privilege level 15 to the session, but the privilege level configured at the user level.

If the `aaa authorization exec LOCAL` command is configured, the ASA will check the **service-type** attribute of the user to determine which management sessions the user is allowed to use and does not have a direct correlation with the user-configured privilege level. To configure the service-type for a user, the following command sequence must be used:

```
username <string> password <string> privilege <0-15>
username <string> attributes
 service-type <admin|nas-prompt|remote-access>
```

If not specified, the default service-type is admin. The following access is granted to the user based on the configured value; in all cases the user can be used for authenticating on remote-access VPNs:

- **admin** allows all management sessions (console/telnet/ssh/HTTP).
- **nas-prompt** allows all management sessions except HTTP/ASDM but requires command authorization to be configured (if `aaa authentication enable console LOCAL` is configured, users will not be able to access privilege exec mode through the `enable` command); note that ASDM session will be successful but allow only monitoring tasks, the configuration tab being disabled.
- **remote-access** does not allow any management sessions except console access (if `aaa authentication enable console LOCAL` is configured, users will not be able to access privilege exec mode through the `enable` command).

The service-type behavior depends on the configured enable authentication method, because if the global enable password is used, ASA does not check the privilege-level or the service-type attributes of the user; essentially, management sessions appear as being from the **enable_15** user and not from the username typed at the authentication prompt.

Configuration

```
ASA3:

passwd CCIELAB
enable password SECRET
!
telnet 136.1.37.0 255.255.255.0 VLAN37
!
```

https://t.me/learningnets

```
http 136.1.27.0 255.255.255.0 VLAN17
http server enable
!
interface GigabitEthernet0/0
 security-level 10
!
aaa authentication serial console LOCAL
aaa authentication http console LOCAL
aaa authorization exec LOCAL
aaa authorization command LOCAL
aaa local authentication attempts max-fail 3
!
username HTTP-USER password HTTP-PASS privilege 15
username NO-LOCK password NO-LOCK privilege 15
username NO-ASDM password NO-ASDM privilege 1
!
username ASDM-RO password ASDM-RO
username ASDM-RO attributes
 service-type nas-prompt
!
username CONSOLE nopassword
username CONSOLE attributes
 service-type remote-access
```

Verification

Verify that telnet and HTTP services are active, as configured:

```
ASA3# show asp table socket


Protocol  Socket    Local Address               Foreign Address         State TCP       000098cf
136.1.37.17:23              0.0.0.0:* LISTEN
SSL       00076dff 136.1.17.17:443             0.0.0.0:* LISTEN
```

Initiate a telnet session from VLAN 37 and authenticate using the configured telnet and enable passwords; verify that the session shows as established:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open


User Access Verification
```

https://t.me/learningnets

```
Password:
Type help or '?' for a list of available commands.
ASA3> show curpriv
Username : enable_1Current privilege level : 1
Current Mode/s : P_UNPRASA3> enable
Password: ******
ASA3# show curpriv
Username : enable_15Current privilege level : 15
Current Mode/s : P_PRIV
!
!ASA3# show asp table socket


Protocol  Socket    Local Address               Foreign Address         State
TCP       000098cf  136.1.37.17:23              0.0.0.0:*               LISTEN
SSL       00076dff  136.1.17.17:443             0.0.0.0:*               LISTEN
TCP       00354038  136.1.37.17:23              136.1.37.3:58982        ESTAB
```

If the telnet session was initiated without changing the security-level of the interface that restricts cleat-text telnet sessions, and considering that logging was activated on the ASA, telnet will not be successful and following message will be logged:

```
%ASA-4-402117: IPSEC: Received a non-IPSec packet (protocol= TCP) from 136.1.37.3 to 136.1.37.17.
```

This can also be verified by viewing the packets dropped in the ASP path:

```
ASA3# show asp drop


Frame drop: Telnet not permitted on least secure interface (telnet-not-permitted)     4
```

Verify that console access is authenticated and allowed for user **CONSOLE**:

```
ASA3# exit

Logoff

Username: CONSOLE
Password: *******
Type help or '?' for a list of available commands.ASA3> enable
```

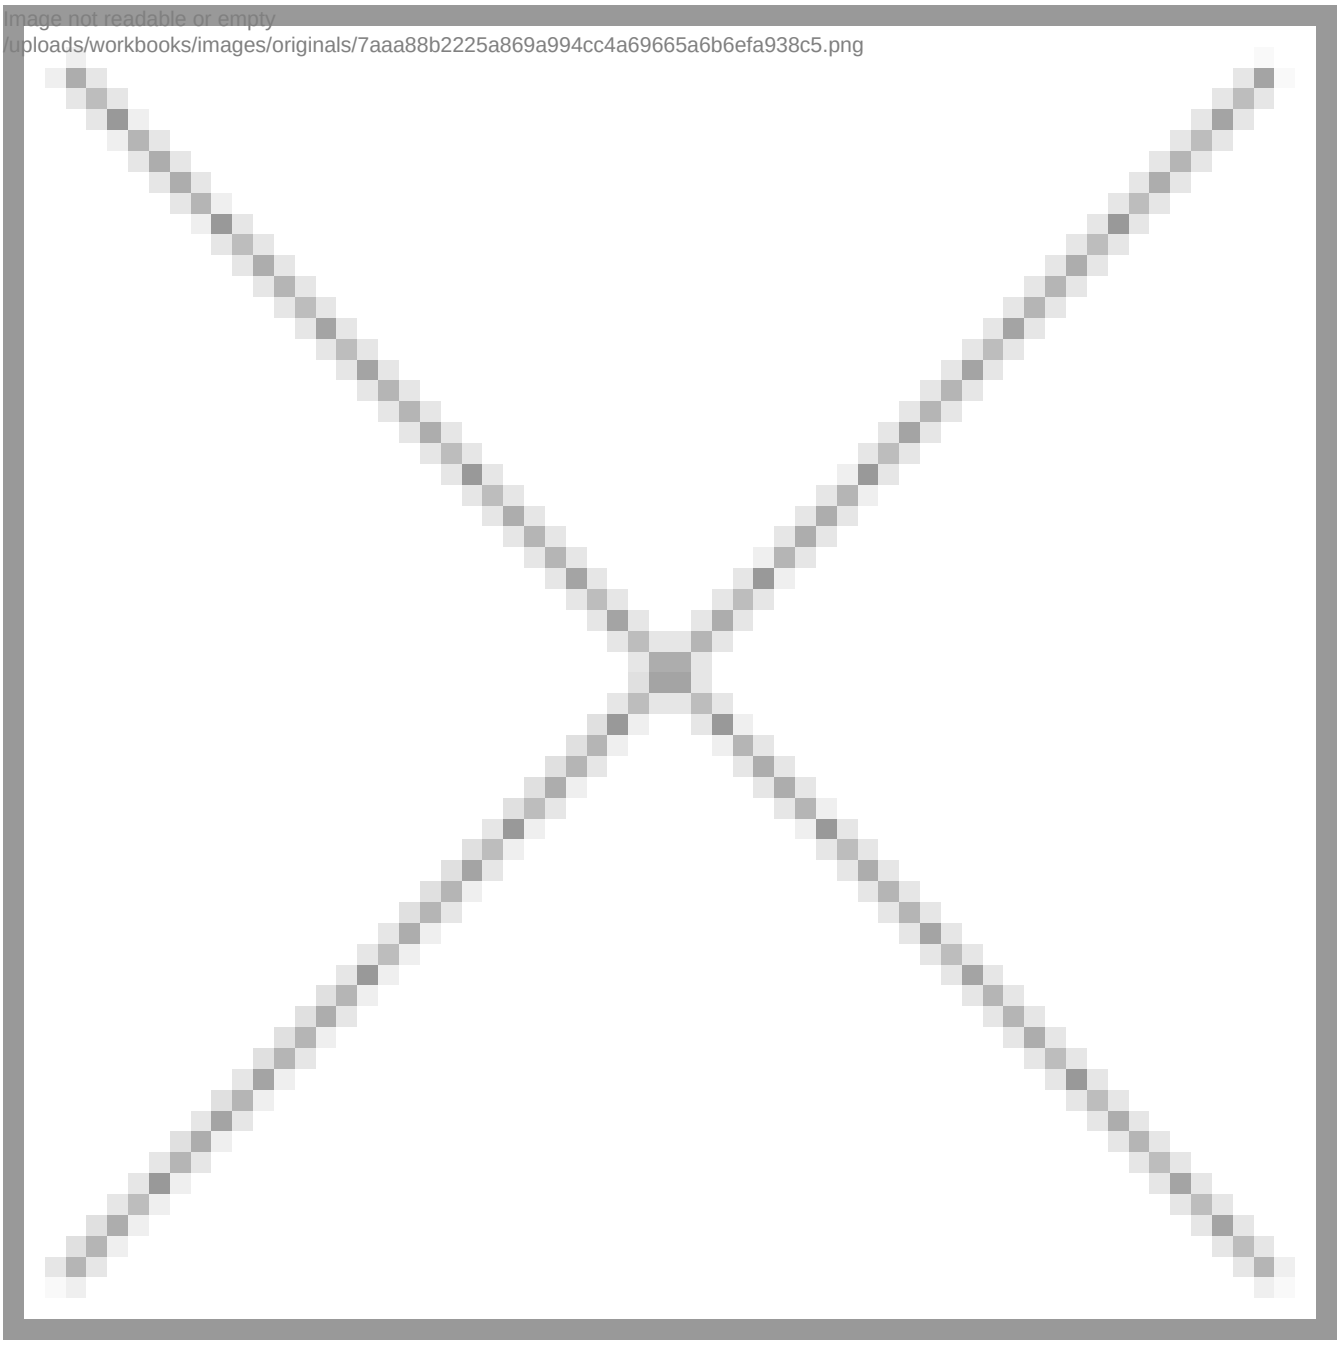
https://t.me/learningnets

```
Password: ******ASA3# show curpriv

Username : enable_15 Current privilege level : 15

Current Mode/s : P_PRIV
```

User **CONSOLE** will not be allowed to connect via ASDM, because it will keep asking for username and password as shown below:



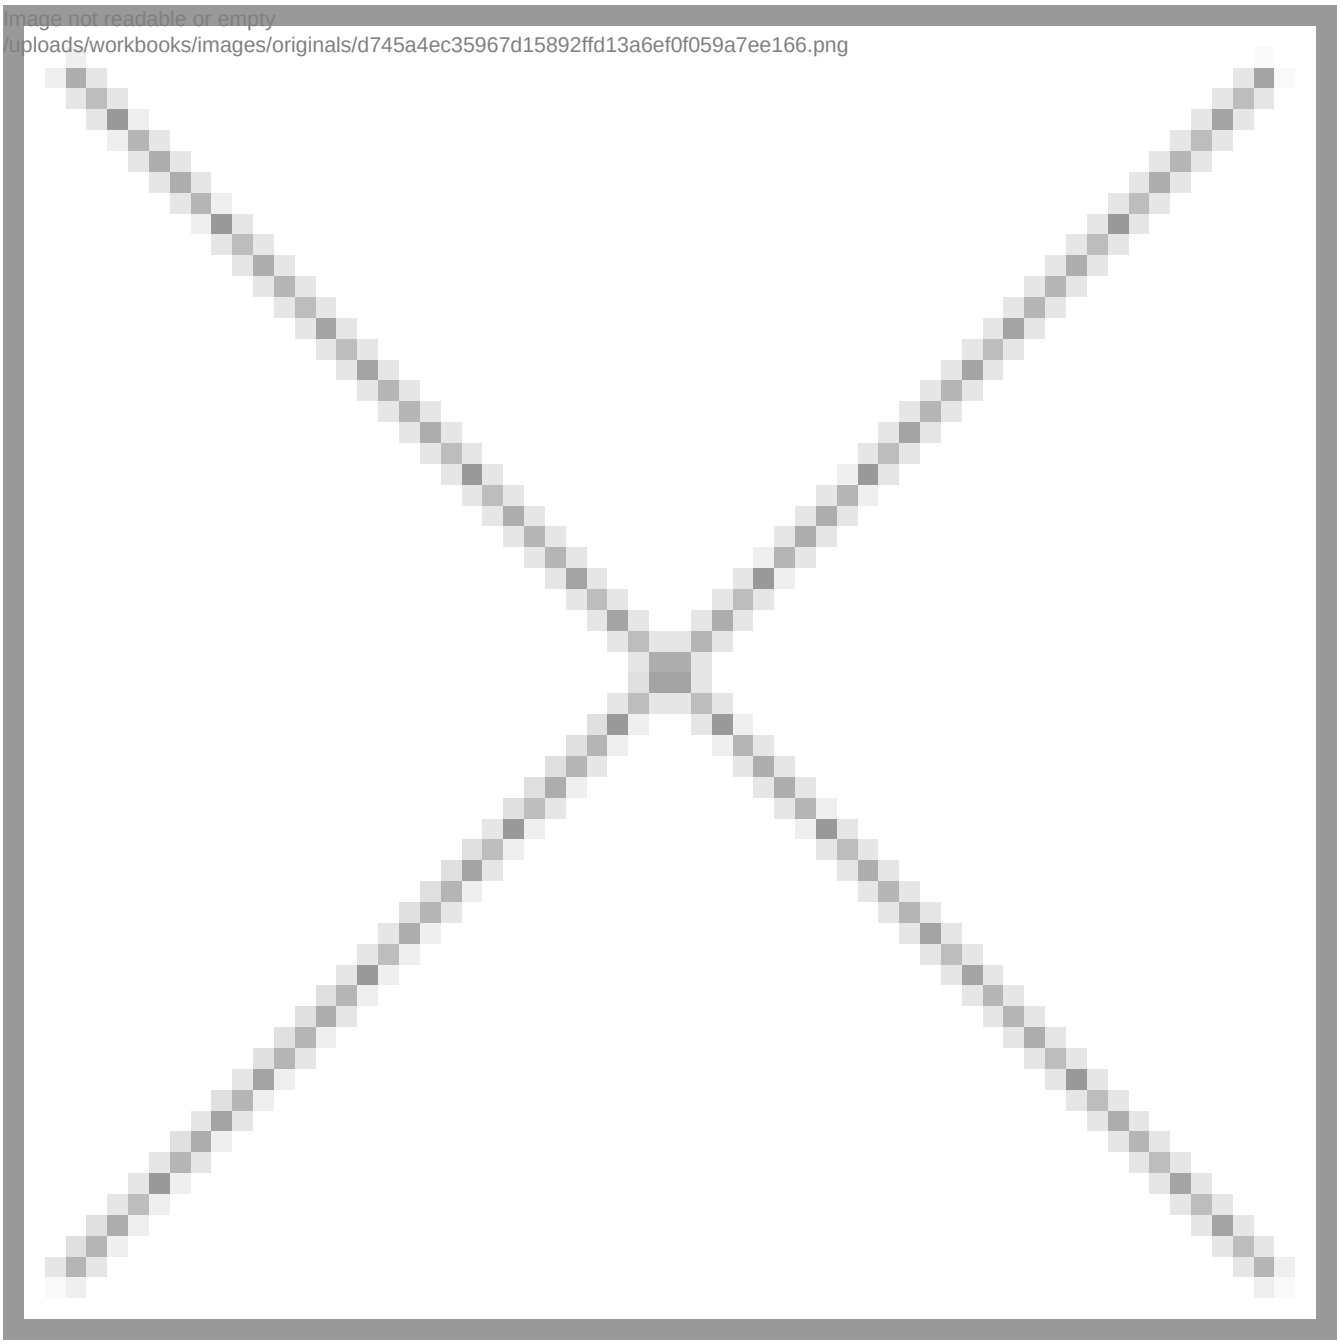
At the CLI, if logging is enabled, the following messages will appear, indicating that the user is not allowed to connect via ASDM based on its service-type attribute:

https://t.me/learningnets

```
%ASA-6-113008: AAA transaction status ACCEPT : user = CONSOLE
%ASA-3-113021: Attempted console login failed user 'CONSOLE' did NOT have appropriate Admin Rights.

%ASA-6-725007: SSL session with client VLAN17:136.1.27.51/19687 terminated.
```

User **NO-ASDM** will not be allowed to connect via ASDM; the following error message appears:



User **HTTP-USER** is allowed to fully manage the ASA via ASDM; user **ASDM-RO** can connect but the configuration tab will be disabled. Fail authentication three consecutive times for the user **NO-ASDM** and verify that the user account is locked; fail authentication three consecutive times for user **NO-LOCK** and verify that the

https://t.me/learningnets

user account is not locked:

```
ASA3# show aaa local user lockout

Lock-time  Failed-attempts      Locked  User13:27:55                3       Y       NO-ASDM
!
!ASA3# clear aaa local user lockout username NO-ASDM
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA Local Command Authorization

Task

Option 1

- Configure ASA3 to allow telnet traffic from VLAN 37 and authenticate sessions using the local user database.
- Using the default enable authentication method, configure ASA3 to:
  - Allow user **PRIV7** access to any **show crypto** type of commands at privilege level 7.
  - Allow user **PRIV8** access to configure nameifs and security-levels and assign IP addresses to interfaces at privilege level 8.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The commands that users **PRIV7** and **PRIV8** require have access have a default privilege level of 15, so it must be modified. To view the default level of a command, use the `show running-config all privilege command <command>` :

```
ASA3# show running-config all privilege command route

privilege show level 15 mode exec command route
```

https://t.me/learningnets

```
privilege show level 15 mode configure command route
privilege clear level 15 mode exec command route
privilege cmd level 15 mode configure command route
privilege clear level 15 mode configure command route
privilege cmd level 15 mode router command router-id
privilege cmd level 15 mode router command routerid
privilege cmd level 15 mode mpf-policy-map-param command router-alert
```

Configuration

ASA3:

```
interface GigabitEthernet0/0
 security-level 10
!
telnet 136.1.37.0 255.255.255.0 VLAN37
!
aaa authentication telnet console LOCAL
aaa authorization command LOCAL
!
username PRIV7 password PRIV7
username PRIV8 password PRIV8
enable password PRIV7 level 7
enable password PRIV8 level 8
!
privilege show level 7 command crypto
privilege level 8 command configure
privilege configure level 8 command interface
privilege configure level 8 command nameif
privilege configure level 8 command security-level
privilege configure level 8 command ip
```

Verification

Initiate a telnet session and authenticate using **PRIV7** user; verify which **show** commands are allowed:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open


User Access Verification
Username: PRIV7
```

https://t.me/learningnets

```
Password: *****
Type help or '?' for a list of available commands.ASA3> show curpriv
Username : PRIV7Current privilege level : 1
Current Mode/s : P_UNPRASA3> show crypto ikev1 sa
         ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
ASA3> enable 7
Password: *****ASA3# show crypto ikev1 sa
There are no IKEv1 SAs
ASA3# show crypto ipsec sa
There are no ipsec sas
ASA3# show crashinfo
         ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
ASA3# configure terminal
     ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
```

Initiate a telnet session and authenticate using the **PRIV8** user; verify the configuration commands that are allowed. Additionally, the **PRIV8** user will have access to all commands with a privilege lower than 8, thus to **show** commands configured for privilege level 7.

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open


User Access Verification

Username: PRIV8
Password: *****
Type help or '?' for a list of available commands.ASA3> show curpriv
Username : PRIV8
Current privilege level : 1
Current Mode/s : P_UNPRASA3> enable 8
Password: *****ASA3# show curpriv
Username : enable_8Current privilege level : 8
Current Mode/s : P_PRIVASA3# configure terminal
ASA3(config)# interface GigabitEthernet0/2
ASA3(config-if)# nameif TEST
ASA3(config-if)# security-level 50
ASA3(config-if)# ip address 10.10.10.17 255.255.255.0
ASA3(config-if)# no shutdown
```

https://t.me/learningnets

```
                    ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
ASA3(config-if)# exit
ASA3(config)# route VLAN37 0 0 136.1.37.3
                ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
ASA3(config)# exit
ASA3# show crypto ikev1 sa
There are no IKEv1 SAs
```

Option 2

- Configure ASA3 to allow telnet traffic from VLAN37 and authenticate sessions using the local user database.
- Without configuring any enable passwords, configure ASA3 to:
  - Allow user **PRIV7** access to any **show crypto** type of commands at privilege level 7.
  - Allow user **PRIV8** access to configure nameifs and security-levels and assign IP addresses to interfaces at privilege level 8.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices.

Overview

Configuration is mostly the same, except we'll be using the user's own passwords to reach the privilege exec level. When the `aaa authentication enable console LOCAL` command is configured, users can no longer specify the privilege level when using the `enable` command; instead, users will be assigned the privilege-level configured at the user level.

https://t.me/learningnets

Configuration

```
ASA3:

interface GigabitEthernet0/0
 security-level 10
!
telnet 136.1.37.0 255.255.255.0 VLAN37
!
aaa authentication telnet console LOCAL
aaa authentication enable console LOCAL
aaa authorization command LOCAL
!
username PRIV7 password PRIV7 privilege 7
username PRIV8 password PRIV8 privilege 8
!
privilege show level 7 command crypto
privilege level 8 command configure
privilege configure level 8 command interface
privilege configure level 8 command nameif
privilege configure level 8 command security-level
privilege configure level 8 command ip
```

Verification

Initiate a telnet session and authenticate using **PRIV7** user; verify which **show** commands are allowed:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open


User Access Verification
Username: PRIV7
Password: *****
Type help or '?' for a list of available commands.ASA3> enable 7
Enabling to privilege levels is not allowed when configured for AAA authentication. Use 'enable' only.
ASA3> enable
Password: *****ASA3# show curpriv
Username : PRIV7Current privilege level : 7
Current Mode/s : P_PRIVASA3# configure terminal
```

https://t.me/learningnets

```
    ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
ASA3# show crypto ikev1 sa
 There are no IKEv1 SAs
```

Initiate a telnet session and authenticate using the **PRIV8** user; verify the configuration commands that are allowed. Additionally, the **PRIV8** user will have access to all commands with a privilege lower than 8, thus to **show** commands configured for privilege level 7:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open


User Access Verification
 Username: PRIV8
Password: *****
Type help or '?' for a list of available commands.ASA3> enable 8
 Enabling to privilege levels is not allowed when configured for AAA authentication. Use 'enable' only.
ASA3> enable

Password: *****ASA3# show curpriv
Username : PRIV8Current privilege level : 8
Current Mode/s : P_PRIVASA3# configure terminal
ASA3(config)# interface GigabitEthernet0/2
ASA3(config-if)# nameif TEST
ASA3(config-if)# security-level 50
ASA3(config-if)# ip address 10.10.10.17 255.255.255.0
ASA3(config-if)# no shutdown
                   ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA RADIUS Authentication

Task

- Configure ASA3 to authenticate telnet sessions from VLAN 37 and console sessions through ACS using RADIUS protocol.
- ACS is reachable at the IP address of 172.16.1.100.
- Configure the following users on ACS:

Username	Password	Enable Password	Privilege Level
ADMIN	adminpass	adminenable	15
PRIV4	priv4pass	priv4enable	4
CON	conpass	conenable	15

- Configure ACS so that user **ADMIN** can manage ASA3 through all possible methods, and user **CON** can only authenticate at the console.
- Ensure that user **PRIV4** is allowed to issue the command `show route` .

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of

https://t.me/learningnets

reloading the devices. This command, however, is not supported by the ASA.

Overview

Telnet and console access must be enabled for RADIUS-based authentication, as well as access to privileged exec mode through the `enable` command, to be assigned the privilege-level from the RADIUS server. To allow user **PRIV4**, issue the `show route` command at privilege level 4; because RADIUS does not support command authorization, we must configure local command authorization and lower the privilege level of the command.

For user **CON** to be allowed only console access, the RADIUS server must assign to the user a service-type value of 5. For ASA to take this value into consideration, exec authorization must be configured as well through the RADIUS server.

Note that the configured enable passwords are not actually usable by the task; this is also why the enable password is different from the login password. If you recall from ACS 3.x and 4.x, the enable password configured at the user level was actually called **TACACS+ Enable Password**, because it can only be used by TACACS, not by RADIUS. When you log in on IOS or ASA via telnet/SSH/console and the device is configured for TACACS+ authentication, it will use a TACACS+ **Service** of **Login**, so ACS will know to use the user's password. When you use the `enable` command on an IOS or ASA and the device is configured for TACACS+ authentication, it will use a TACACS+ **Service** of **Enable**, so ACS will know to use the user's enable password. This different functionality between the **login** and **enable** actions is not supported by RADIUS IETF attributes and Cisco did not use VSA's to implement it, either. For both **login** and **enable** actions, when using RADIUS, IOS and ASA will use the RADIUS attribute **NAS-Port-Type** with a value of **virtual**. IOS uses the RADIUS attribute **User-Name** with a value equal to the username entered for **login** action and with a value of **$enab[priv-lvl]$** for **enable** action. ASA will use for both **login** and **enable** actions the RADIUS attribute **User-Name** with a value equal to the username entered. Essentially, when using RADIUS, ASA does regular username/password authentication for both the **login** and **enable** actions.

Configuration

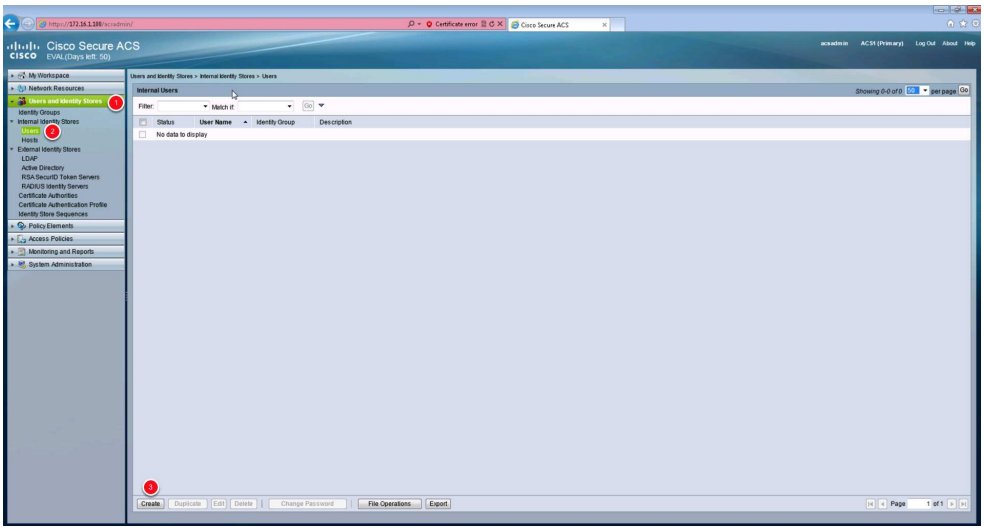
ASA3:

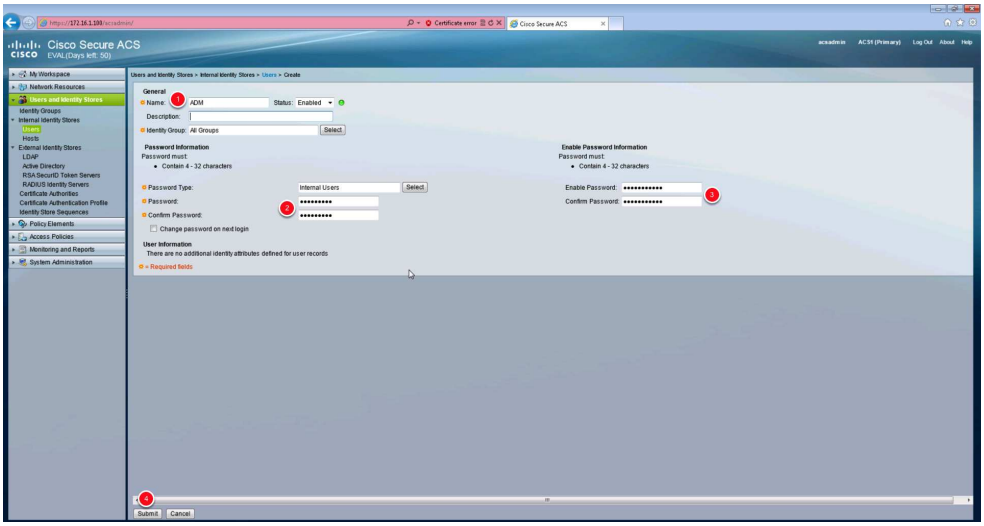
```
interface GigabitEthernet0/0
 security-level 10
!
telnet 136.1.37.0 255.255.255.0 VLAN37
```

https://t.me/learningnets

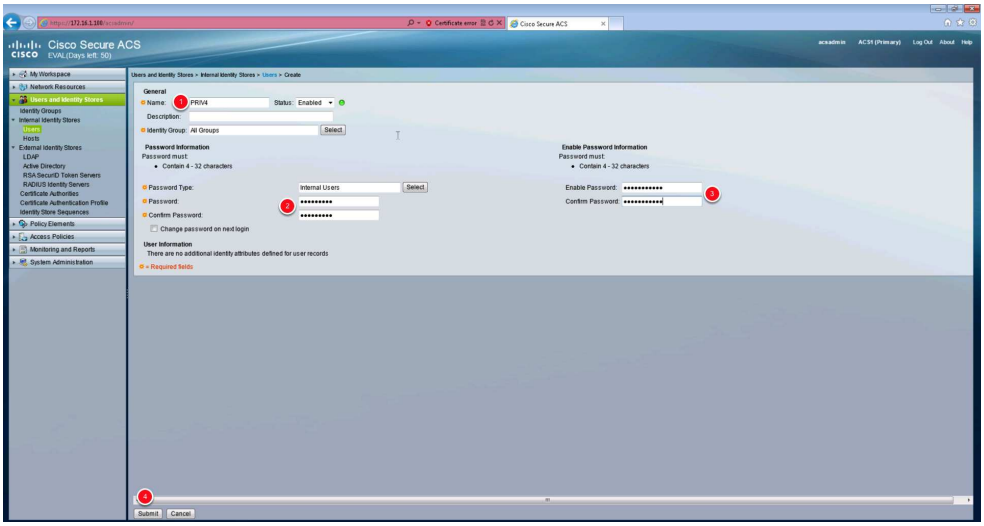
```
!
aaa-server RADIUS protocol radius
aaa-server RADIUS (VLAN17) host 172.16.1.100
 key 0 cisco
!
aaa authentication telnet console RADIUS
aaa authentication serial console RADIUS
aaa authentication enable console RADIUS
!
aaa authorization exec authentication-server
aaa authorization command LOCAL
!
privilege show level 4 command route
```

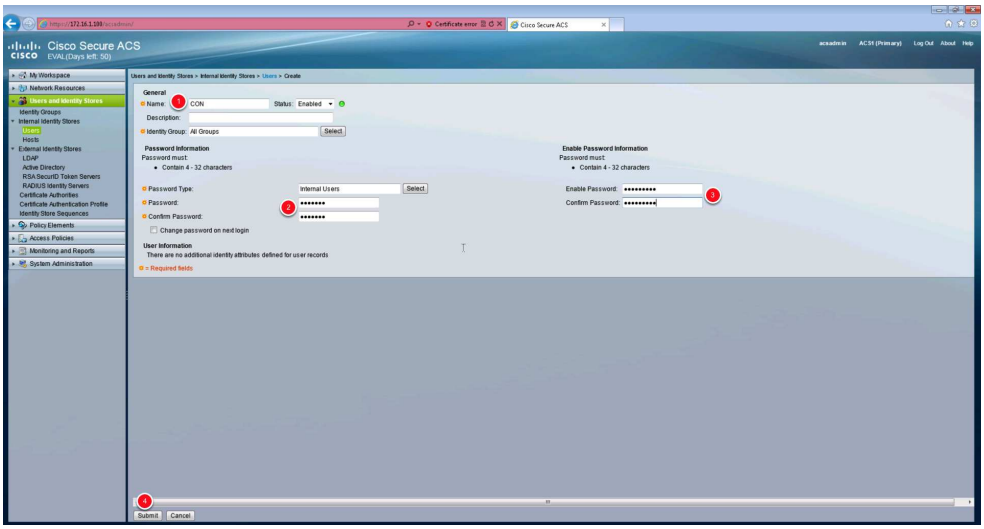
On ACS, the first step is to configure usernames with appropriate login and enable passwords as requested:



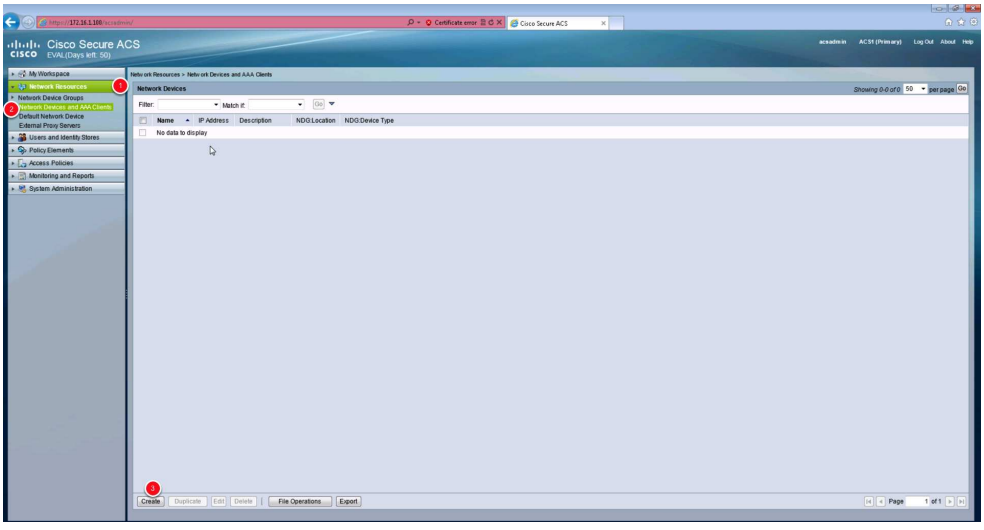


https://t.me/learningnets

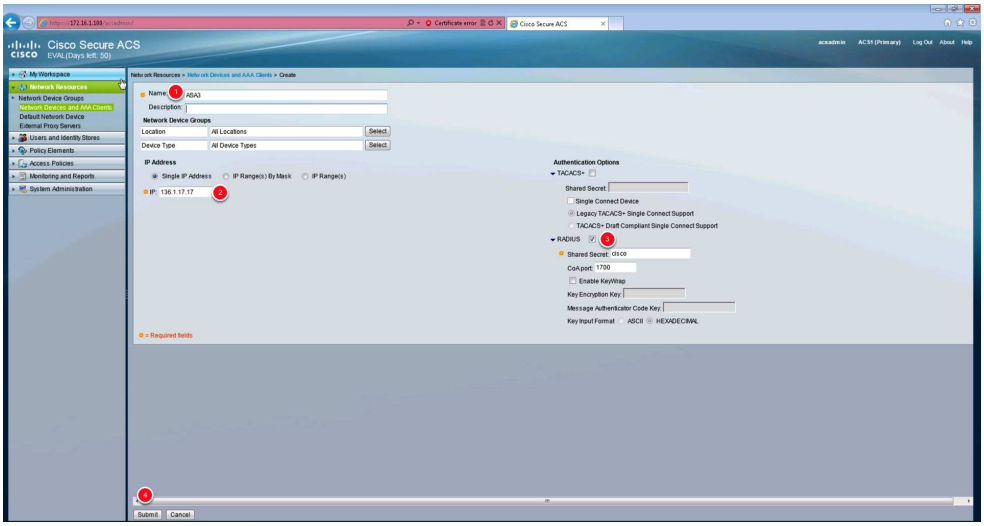




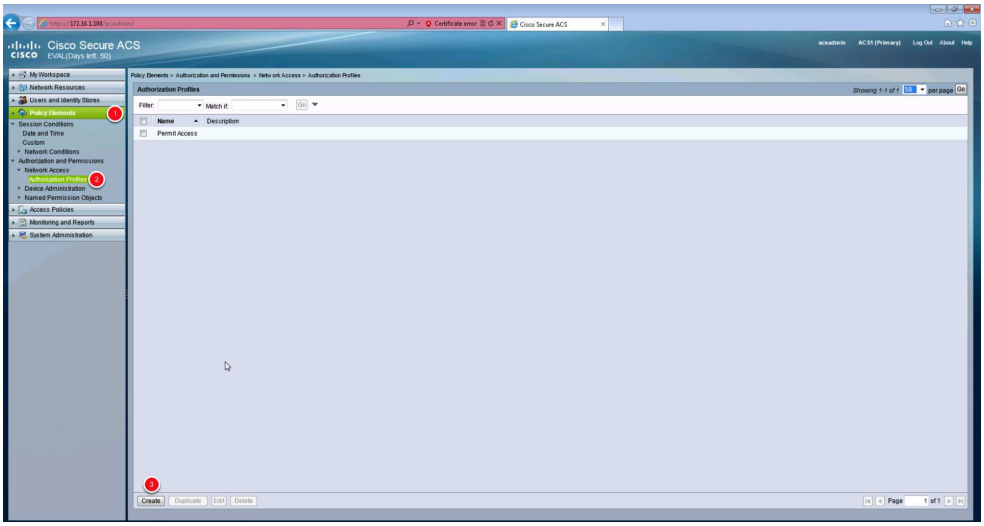
On ACS, the second step is to configure ASA3 as a NAD in ACS, using RADIUS protocol:



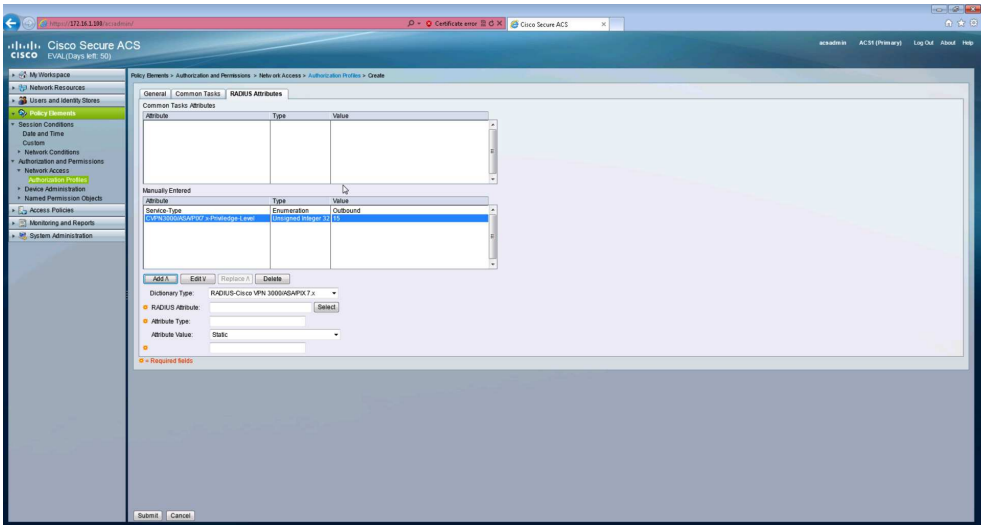
https://t.me/learningnets



On ACS, the third step is to configure three authorization profiles, because each user requires its own profile based on task requirements:

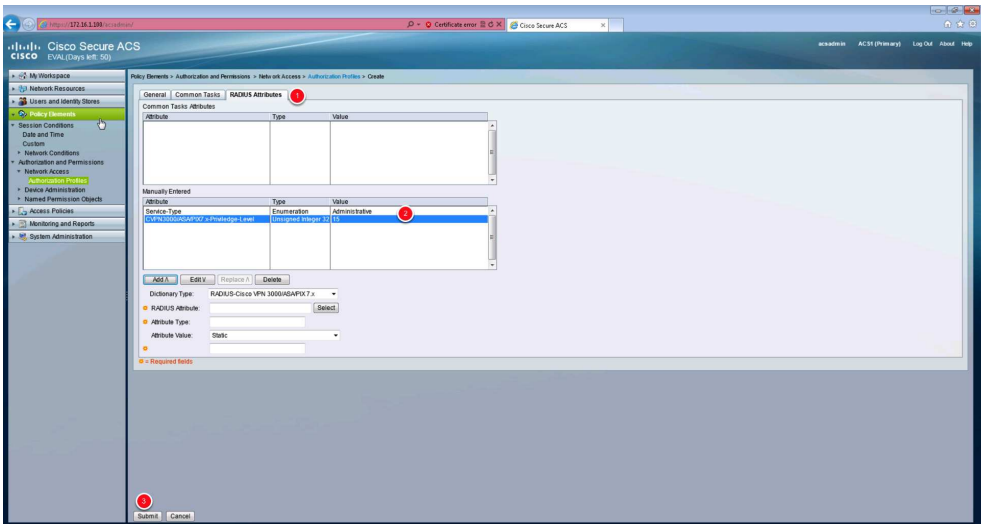


Authorization Profile for user **CON**:

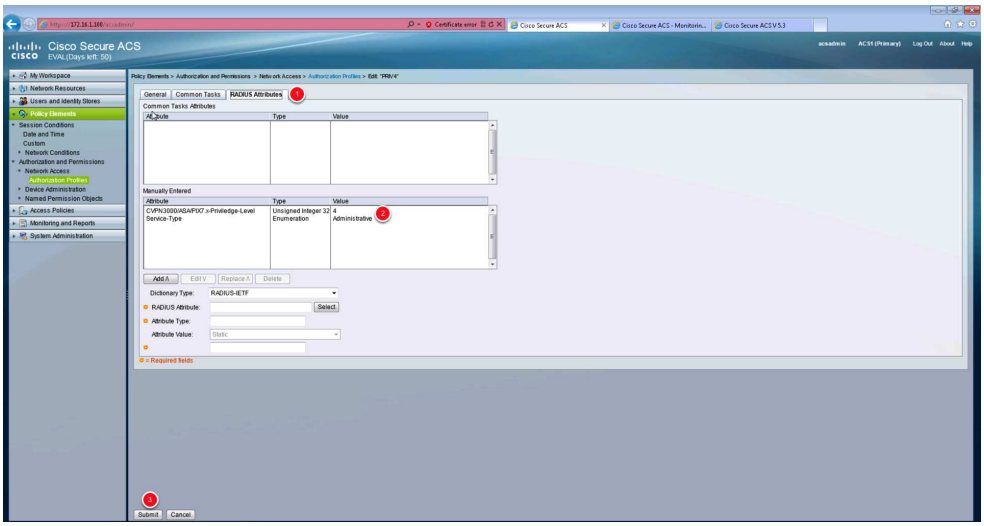


https://t.me/learningnets

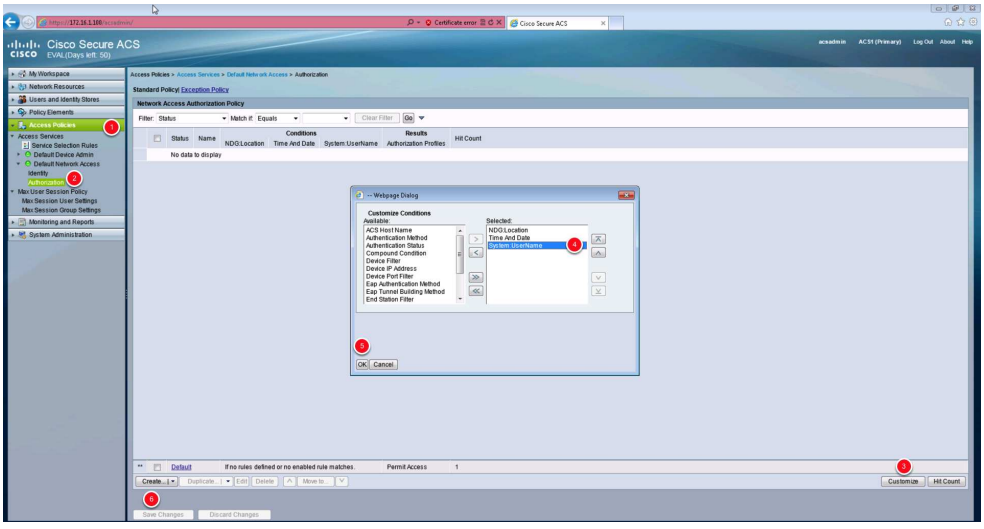
Authorization Profile for user **ADMIN**:



Authorization Profile for user **PRIV4**:



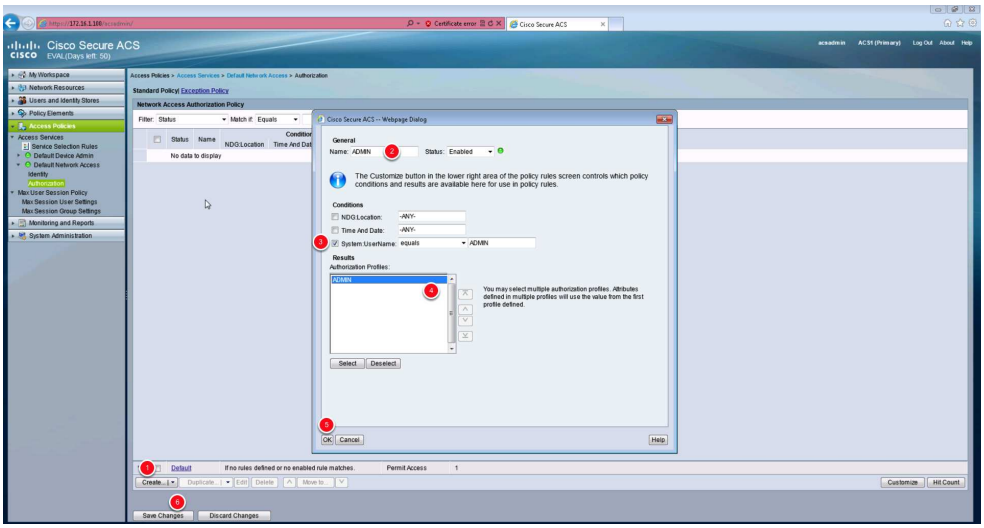
On ACS, add the username as a condition for the RADIUS access policies:

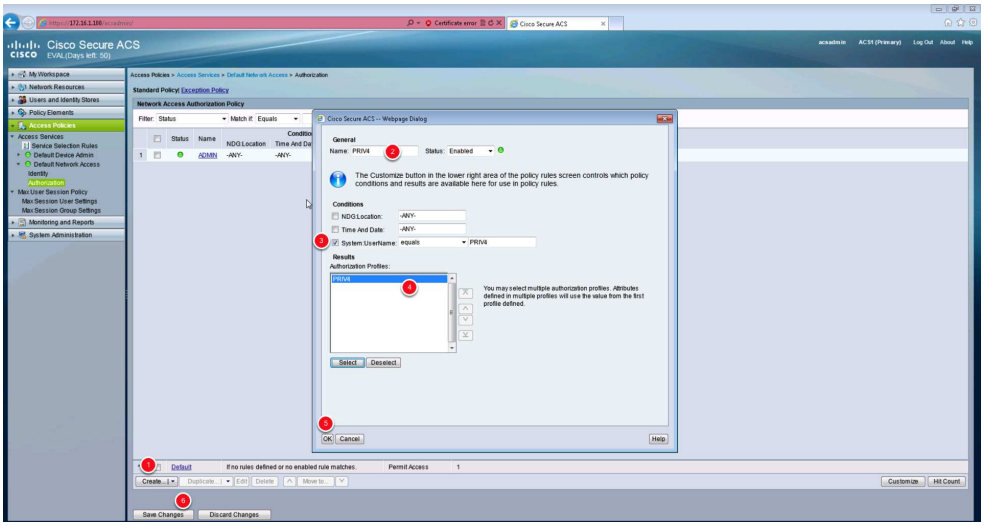


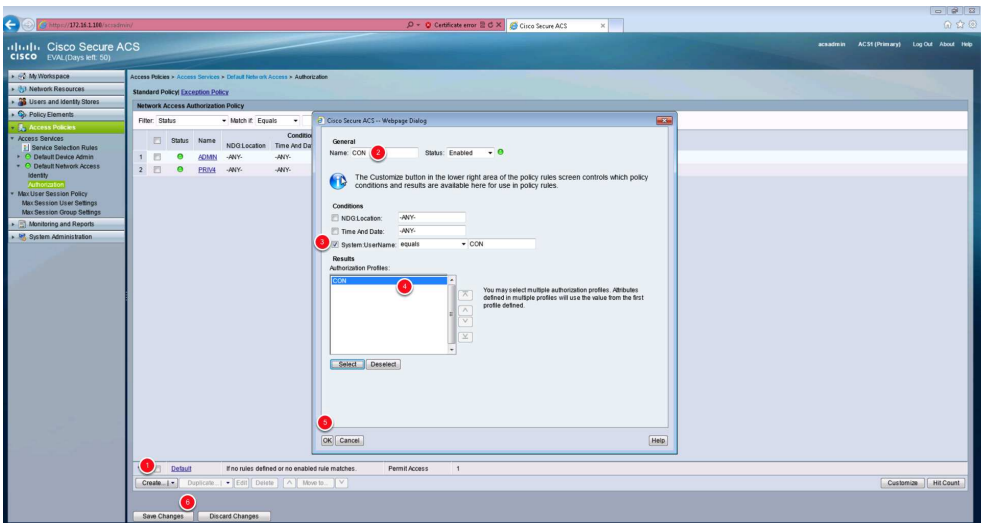
On ACS, configure three RADIUS access policies, the condition being the

https://t.me/learningnets

authenticating username:







Verification

https://t.me/learningnets

Authenticate with user **ADMIN** and verify the assigned privilege-level:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open


User Access Verification
Username: ADMIN
Password: *********
Type help or '?' for a list of available commands.ASA3> show curpriv
Username : ADMINCurrent privilege level : 1
Current Mode/s : P_UNPRASA3> enable
Password:*********ASA3# show curpriv
Username : ADMINCurrent privilege level : 15
Current Mode/s : P_PRIVASA3# configure terminal
ASA3(config)#
```

Verify that user **CON** cannot connect via telnet, but can authenticate on the console:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open

User Access Verification
Username: CON
Password: *******
[ CON ] You do NOT have Admin Rights to the console !
[Connection to 136.1.37.17 closed by foreign host]

ASA3# exit
Logoff
Username: CON
Password: *******
Type help or '?' for a list of available commands.ASA3> enable

Password: *******
[ CON ] You do NOT have enable Admin Rights to the console
```

Verify that user **PRIV4** is assigned privilege level 4 and can issue the `show route` command:

```
R3#telnet 136.1.37.17
```

https://t.me/learningnets

```
Trying 136.1.37.17 ... Open



User Access Verification
Username: PRIV4
Password: *********
Type help or '?' for a list of available commands.ASA3> show curpriv
Username : PRIV4Current privilege level : 1
Current Mode/s : P_UNPRASA3> enable
Password: *********ASA3# show curpriv
Username : PRIV4Current privilege level : 4
Current Mode/s : P_PRIVASA3# show route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
   D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
   N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
   E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
   i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
   * - candidate default, U - per-user static route, o - ODR
   P - periodic downloaded static route

Gateway of last resort is not set

D    136.1.27.0 255.255.255.0 [90/3072] via 136.1.17.9, 21:12:51, VLAN17
D    136.1.29.0 255.255.255.0 [90/3072] via 136.1.17.9, 21:12:51, VLAN17
C    136.1.17.0 255.255.255.0 is directly connected, VLAN17
D    136.1.19.0 255.255.255.0 [90/3072] via 136.1.17.9, 21:12:51, VLAN17
D    136.1.47.0 255.255.255.0 [90/3072] via 136.1.17.9, 21:12:51, VLAN17
C    136.1.37.0 255.255.255.0 is directly connected, VLAN37
D    172.16.10.0 255.255.255.0 [90/26368] via 136.1.17.9, 21:12:51, VLAN17
D    172.16.1.0 255.255.255.0 [90/26368] via 136.1.17.9, 21:12:51, VLAN17
D    172.16.3.0 255.255.255.0 [90/26368] via 136.1.17.9, 21:12:51, VLAN17
D    150.1.9.9 255.255.255.255 [90/130816] via 136.1.17.9, 21:12:51, VLAN17
D    150.1.1.1 255.255.255.255 [90/131072] via 136.1.17.9, 21:12:51, VLAN17
D    150.1.3.3 255.255.255.255 [90/130816] via 136.1.37.3, 21:12:51, VLAN37ASA3# configure terminal
    ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
```

Verify the RADIUS Authentication logs from the ACS server, and note that all users appear to have authenticated twice, because of the enable authentication:

https://t.me/learningnets

AAA Protocol > RADIUS Authentication

Authentication Status : Pass or Fail
Date : August 24, 2013 09:59 AM - August 24, 2013 10:29 AM (Last 30 Minutes | Last Hour | Last 12 Hours | Today | Yesterday | Last 7 Days | Last 30 Days)

Generated on August 24, 2013 10:29:36 AM UTC

Reload

✔=Pass ✖=Fail 🔍=Click for details 🖱=Mouse over item for additional information

ACS View Timestamp	ACS Timestamp	RADIUS Status	NAS Failure	Details	Username	MAC/IP Address	Access Service	Authentication Method	Network Device	NAS IP Address	NAS Port ID	CTS Security Group	ACS Instance	Failure	Reason
Aug 24,13 10:29:25.656 AM	Aug 24,13 10:29:25.626 AM	✔		🔍	CON		Default Network Access	PAP_ASCII	ASA3	136.1.17.17			ACS1		
Aug 24,13 10:29:21.573 AM	Aug 24,13 10:29:21.533 AM	✔		🔍	CON		Default Network Access	PAP_ASCII	ASA3	136.1.17.17			ACS1		
Aug 24,13 10:29:10.643 AM	Aug 24,13 10:29:10.623 AM	✔		🔍	PRIV4	ip:source-ip=136.1.37.3	Default Network Access	PAP_ASCII	ASA3	136.1.17.17			ACS1		
Aug 24,13 10:29:06.923 AM	Aug 24,13 10:29:06.893 AM	✔		🔍	PRIV4	ip:source-ip=136.1.37.3	Default Network Access	PAP_ASCII	ASA3	136.1.17.17			ACS1		
Aug 24,13 10:28:52.233 AM	Aug 24,13 10:28:52.226 AM	✔		🔍	ADMIN	ip:source-ip=136.1.37.3	Default Network Access	PAP_ASCII	ASA3	136.1.17.17			ACS1		
Aug 24,13 10:28:47.293 AM	Aug 24,13 10:28:47.266 AM	✔		🔍	ADMIN	ip:source-ip=136.1.37.3	Default Network Access	PAP_ASCII	ASA3	136.1.17.17			ACS1		

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA LDAP Authentication

Task

- Configure ASA3 to authenticate telnet sessions from VLAN 37 and console sessions via LDAP server on TEST-SRV-A.
- Use **ldapuser** with the password **1qaz@WSX**, which resides in the **Users** organizational unit, to query the LDAP server.
- The following users should already be configured on the LDAP server:

Username	Password	Organizational Unit	Department	Company
USER1	1qaz@WSX	FINANCIAL	FINANCIAL	FINANCIAL
USER2	1qaz@WSX	TECHNICAL	TECHNICAL	TECHNICAL
USER3	1qaz@WSX	CCIELAB	CCIELAB	CCIELAB

- Map the **department** LDAP attribute to the **Service-Type** RADIUS attribute, and the **company** LDAP attribute to the **Privilege-Level** RADIUS attribute, to achieve the following:
  - User **USER3** can fully manage the ASA.
  - User **USER2** is assigned a privilege-level of 9 and has access to the `show nameif` command.
  - User **USER1** can authenticate to the ASA only via console.
- If connection to the LDAP server fails, ensure that user **FALLBACK** can fully manage the ASA.

https://t.me/learningnets

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Telnet and console access must be enabled for LDAP-based authentication, as well as access to privileged exec mode thorough `enable` command, to be assigned the privilege-level from the LDAP server; both methods are configured with fallback on local username database in case connectivity with LDAP server fails. To allow user **USER2** to issue the `show nameif` command at privilege level 9 and because LDAP does not support command authorization, we must configure local command authorization and lower the privilege level of the command.

For user **USER1** to be allowed only console access, LDAP server must assign to the user a service-type value of 5. For ASA to take this value into consideration, exec authorization must be configured as well through the LDAP server.

Because LDAP server does not support the RADIUS attributes required to assign privilege-levels and service-type values, LDAP attribute map is configured, which basically creates a one-to-one mapping between the value of an LDAP attribute and the value of a RADIUS attribute. The LDAP attribute map is useless unless applied under the LDAP server configuration. By default, the ASA has no mappings between LDAP to RADIUS attributes.

LDAP server configuration is straightforward; the scope is configured to **subtree** because the users required for authentication are not found directly under the **inelab.local** root location of the tree, where the search (base-DN) starts. Although ASA uses by default the **sAMAccountName** for LDAP requests, which is the actual username login-id, it is still configured just to make it clear.

Configuration

ASA3:

```
interface GigabitEthernet0/0
 security-level 10
```

https://t.me/learningnets

```
!
telnet 136.1.37.0 255.255.255.0 VLAN37
!
ldap attribute-map LDAP-MAP
 map-name company Privilege-Level
 map-value company CCIELAB 15
 map-value company TECHNICAL 9
 map-value company FINANCIAL 15
 map-name department IETF-Radius-Service-Type
 map-value department CCIELAB 6
 map-value department TECHNICAL 6
 map-value department FINANCIAL 5
!
aaa-server LDAP protocol ldap
aaa-server LDAP (VLAN17) host 172.16.10.100
 ldap-base-dn DC=inelab, DC=local
 ldap-scope subtree
 ldap-naming-attribute sAMAccountName
 ldap-login-password 1qaz@WSX
 ldap-login-dn CN=ldapuser,CN=Users,DC=inelab,DC=local
 server-type microsoft
 ldap-attribute-map LDAP-MAP
!
aaa authentication telnet console LDAP LOCAL
aaa authentication serial console LDAP LOCAL
aaa authentication enable console LDAP LOCAL
!
aaa authorization exec authentication-server
aaa authorization command LOCAL
!
privilege show level 9 command nameif
username FALLBACK password FALLBACK privilege 15
```

Verification

Verify that the LDAP server is active:

```
ASA3# show aaa-server protocol ldap

Server Group:    LDAP
Server Protocol: ldapServer Address:  172.16.10.100
Server port:     0Server status:   ACTIVE, Last transaction at unknown

Number of pending requests              0
Average round trip time                 0ms
```

https://t.me/learningnets

```
Number of authentication requests       11
Number of authorization requests        0
Number of accounting requests           0
Number of retransmissions               0
Number of accepts                       9
Number of rejects                       0
Number of challenges                    0
Number of malformed responses           0
Number of bad authenticators            0
Number of timeouts                      0
Number of unrecognized responses        0
```

Authenticate with **USER3** and verify that you can configure the ASA:

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open

User Access Verification
Username: USER3
Password: ********
Type help or '?' for a list of available commands.ASA3> enable
Password: ********ASA3# show curpriv
Username : USER3Current privilege level : 15
Current Mode/s : P_PRIVASA3# configure terminal
ASA3(config)#
```

Authenticate with **USER2** and verify the assigned privilege level and that `show nameif` is allowed:

https://t.me/learningnets

```
R3#telnet 136.1.37.17

Trying 136.1.37.17 ... Open

User Access Verification
Username: USER2

Password: ********
Type help or '?' for a list of available commands.ASA3> enable

Password: ********ASA3# show curpriv

Username : USER2Current privilege level : 9

Current Mode/s : P_PRIVASA3# show nameif

Interface                Name                    Security
GigabitEthernet0/0       VLAN37                   10
GigabitEthernet0/1       VLAN17                    0ASA3# configure terminal

    ^
ERROR: % Invalid input detected at '^' marker.ERROR: Command authorization failed
```

Verify that **USER1** can only authenticate via the console:

https://t.me/learningnets

```
R3#telnet 136.1.37.17

Trying 136.1.37.17 ... Open

User Access Verification
Username: USER1

Password: ********
[ USER1 ] You do NOT have Admin Rights to the console !


[Connection to 136.1.37.17 closed by foreign host]
!
!ASA3# exit


Logoff
Username: USER1

Password: ********
Type help or '?' for a list of available commands. ASA3>
```

Bring down the switch port toward LDAP server and verify that you can still authenticate on the ASA:

https://t.me/learningnets

```
SW1(config)#interface fastethernet1/0/4
SW1(config-if)#shutdown

!R3#telnet 136.1.37.17

Trying 136.1.37.17 ... Open


User Access Verification
Username: FALLBACK

Password: ********
Type help or '?' for a list of available commands.ASA3> enable

Password: ********ASA3# show curpriv

Username : FALLBACKCurrent privilege level : 15

Current Mode/s : P_PRIVASA3# configure terminal
ASA3(config)#
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA TACACS Command Authorization and Accounting

Task

- Configure ASA3 for AAA services using the ACS server as follows:
  - Ensure that telnet, console, and privilege exec mode are accounted for using RADIUS.
  - Using TACACS, authenticate telnet and console sessions.
  - Account for configuration commands.
- Ensure that telnet management sessions from VLAN37 are allowed by the ASA.
- Configure users on ACS as follows:
  - User **USER2** should be allowed access to all commands except `clear` type of commands.
  - User **USER1** should be allowed access to all `show` type of commands, add or remove security-levels on interfaces, and disconnect from the management session.
  - Both users should be assigned the maximum privilege level required.
- Your access policies should be configured based on Identity Groups.
- To ensure password confidentiality and security, configure the ACS so that:
  - Users are required to change their password on first use.
  - Password cannot use the string **cisco**.
  - The password minimum character length is 5.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config`

https://t.me/learningnets

instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

RADIUS does not support command authorization and accounting, so TACACS+ protocol is used for this. Note that ASA accounts for all except `show` type of commands, and this is not configurable. Accounting for start and stop of management sessions such as console and telnet is supported by both RADIUS and TACACS+, but the task requires use of RADIUS.

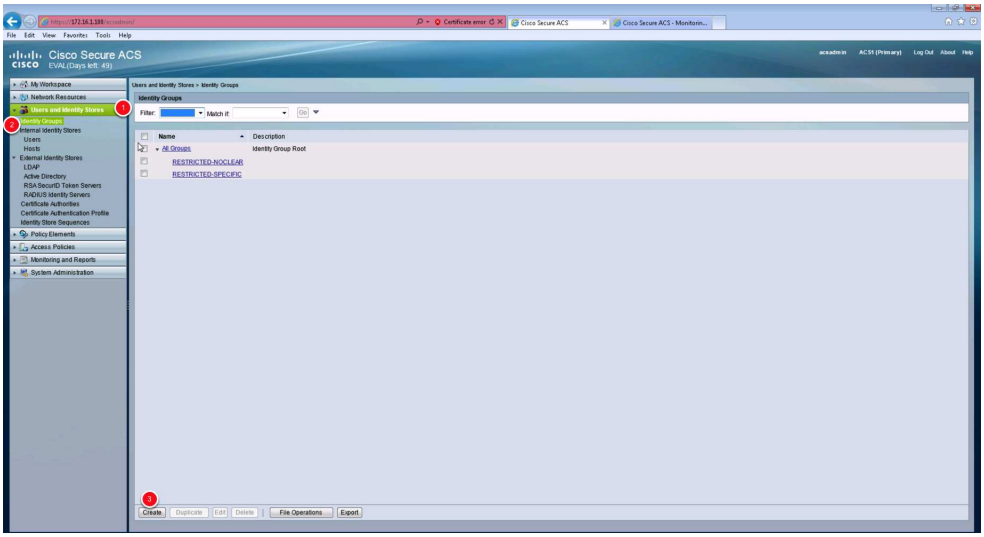
Configuration

ASA3:

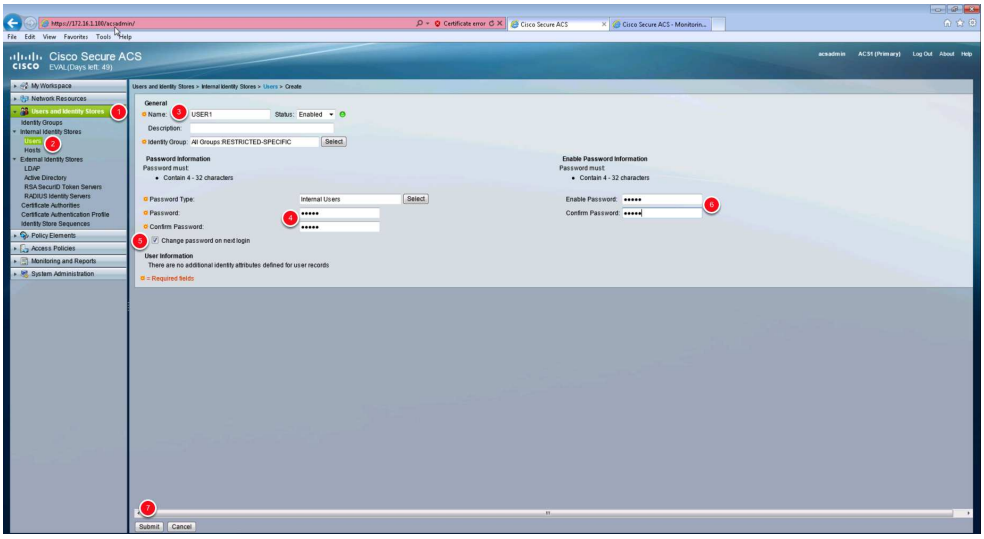
```
interface GigabitEthernet0/0
 security-level 10
!
telnet 136.1.37.0 255.255.255.0 VLAN37
!
aaa-server TACACS protocol tacacs+
aaa-server TACACS (VLAN17) host 172.16.1.100
 key 0 cisco
!
aaa-server RADIUS protocol radius
aaa-server RADIUS (VLAN17) host 172.16.1.100
 key 0 cisco
!
aaa accounting telnet console RADIUS
aaa accounting serial console RADIUS
aaa accounting enable console RADIUS
aaa accounting command TACACS
!
aaa authentication telnet console TACACS
aaa authentication serial console TACACS
aaa authentication enable console TACACS
aaa authorization command TACACS
```

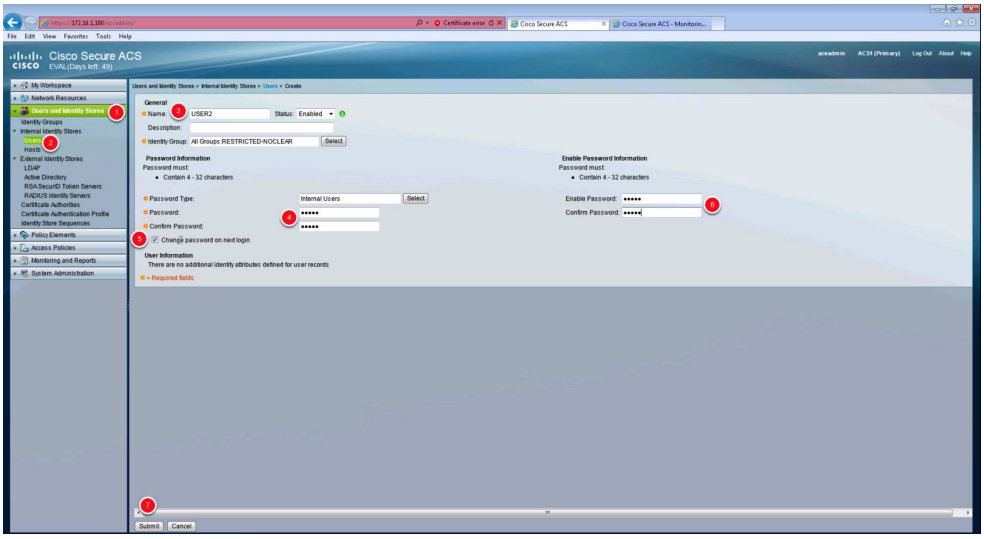
To configure access-policies based on Identity Groups, configure two groups on ACS, one for each user:

https://t.me/learningnets



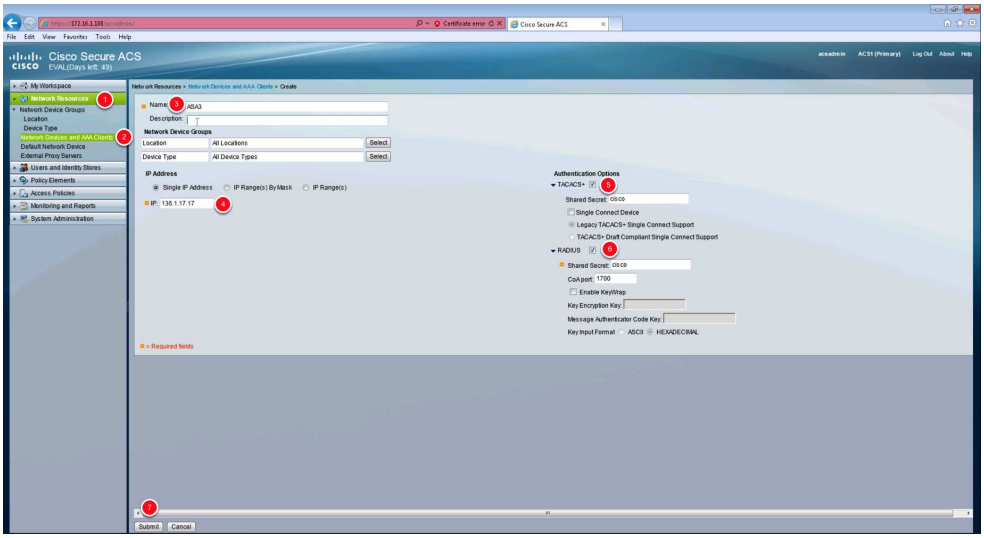
Create the two users, assign them to appropriate Identity Groups, and enable the change password functionality:



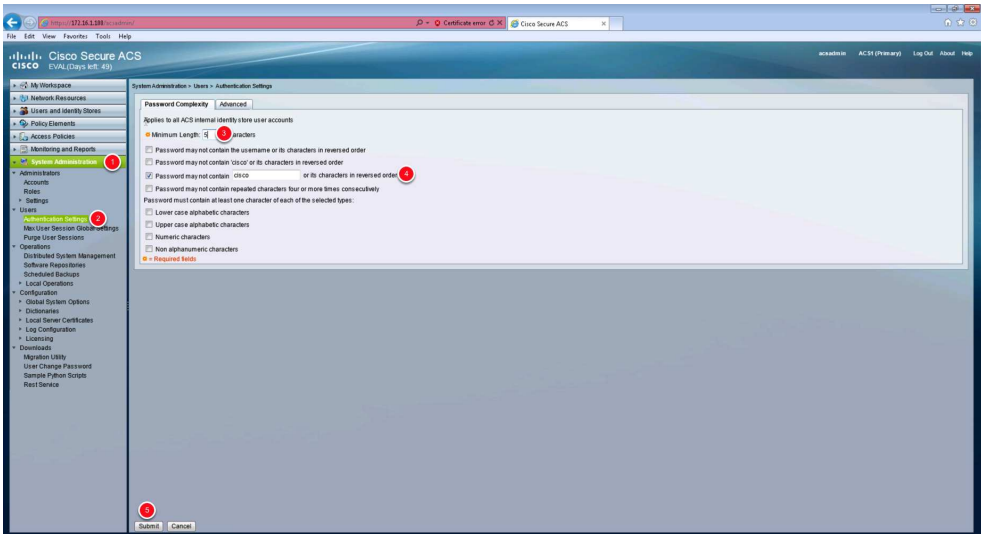


Add the ASA as a NAD in ACS, and configure it for both RADIUS and TACACS+:

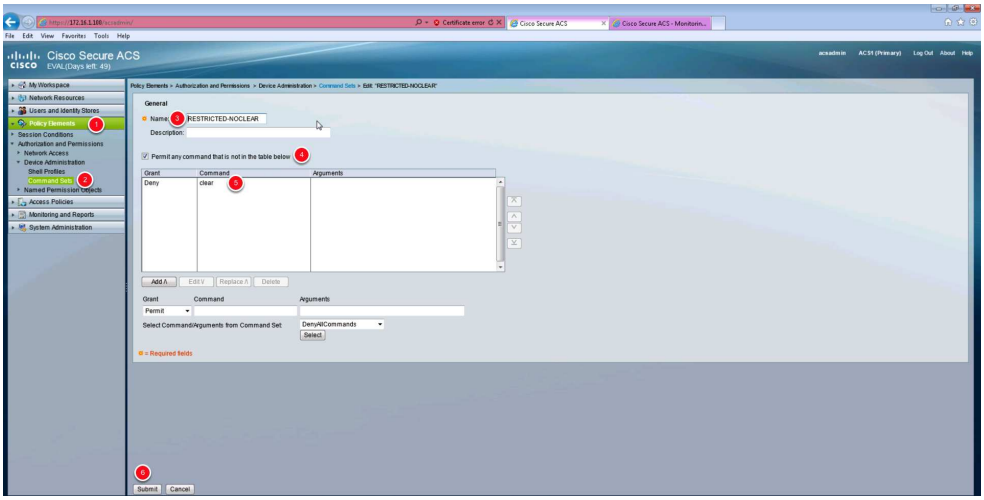
https://t.me/learningnets



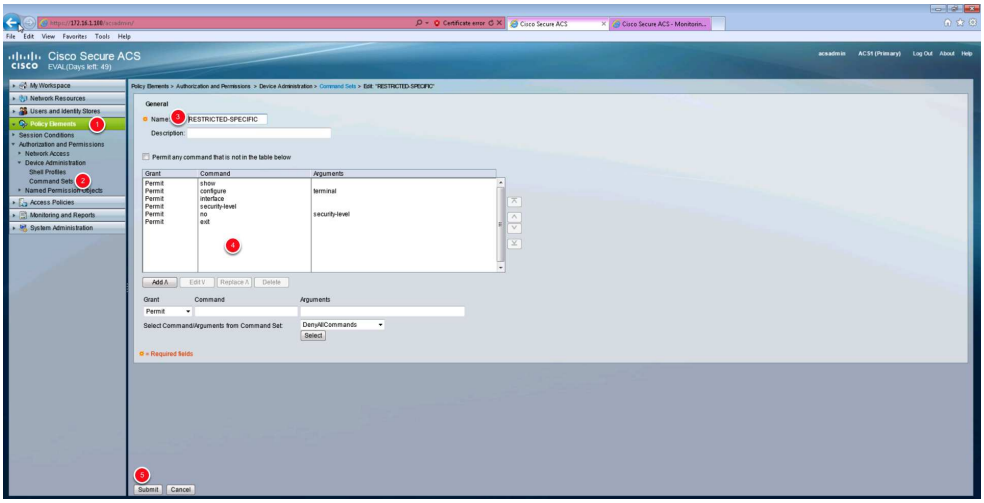
Modify the user password policy to match the task requirements:



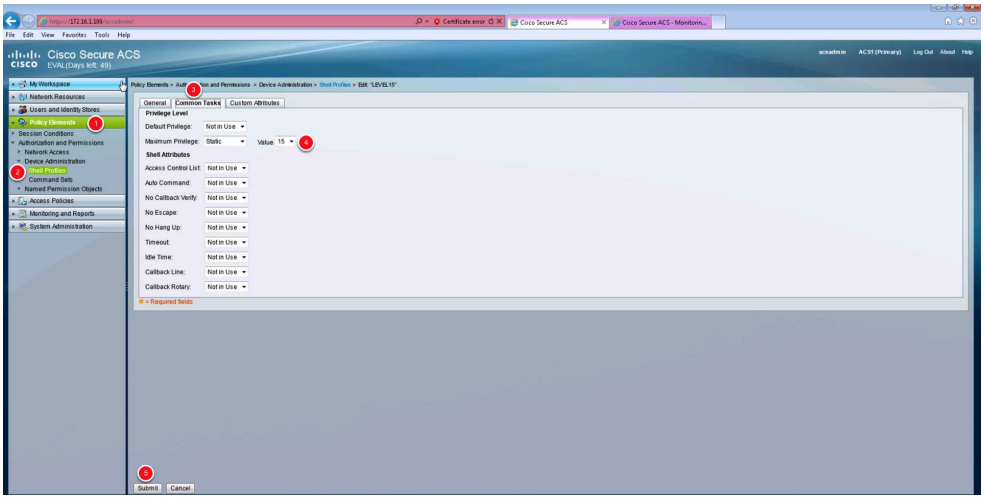
Configure the command sets for each user as specified by the task requirements:



https://t.me/learningnets

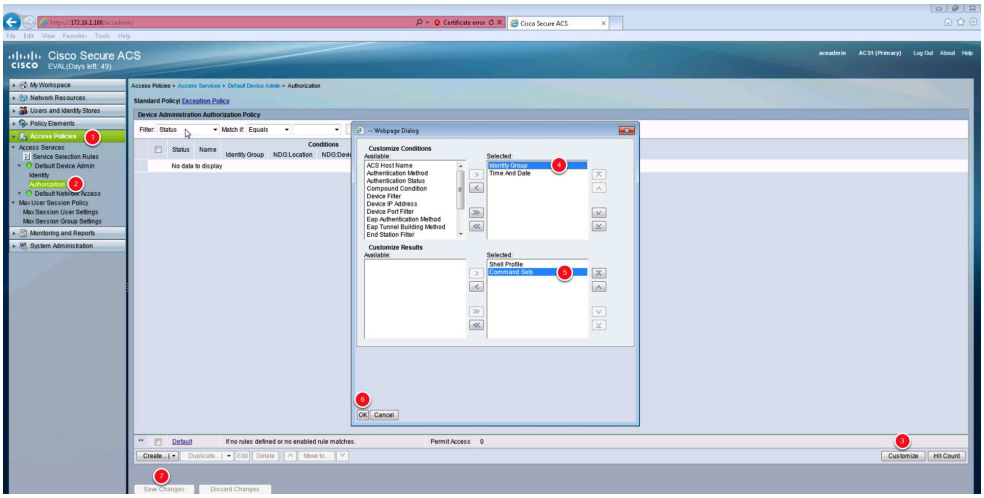


Configure a shell profile to allow privilege level 15 upon successful enable authentication. Note that because ASA always puts the user into privilege-level 1 upon initial authentication and before enable authentication, it is not required to set the **Default Privilege** value:

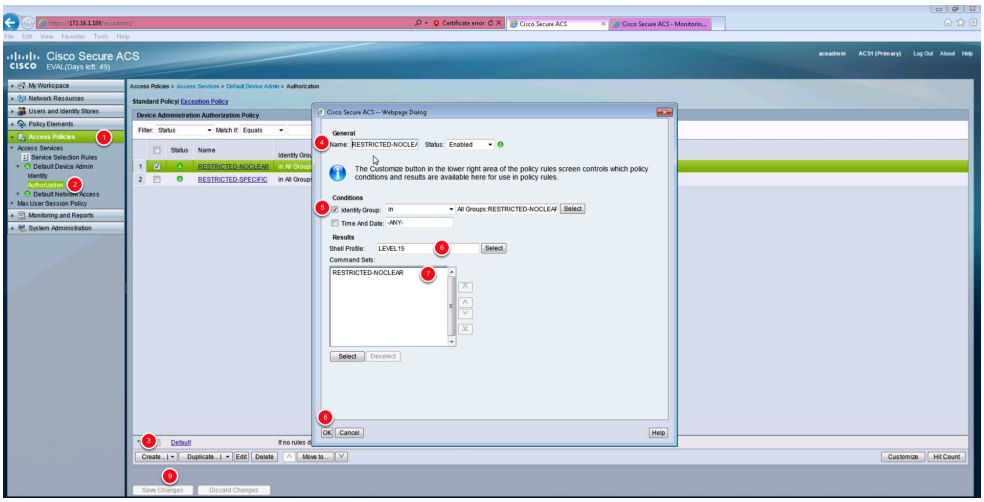


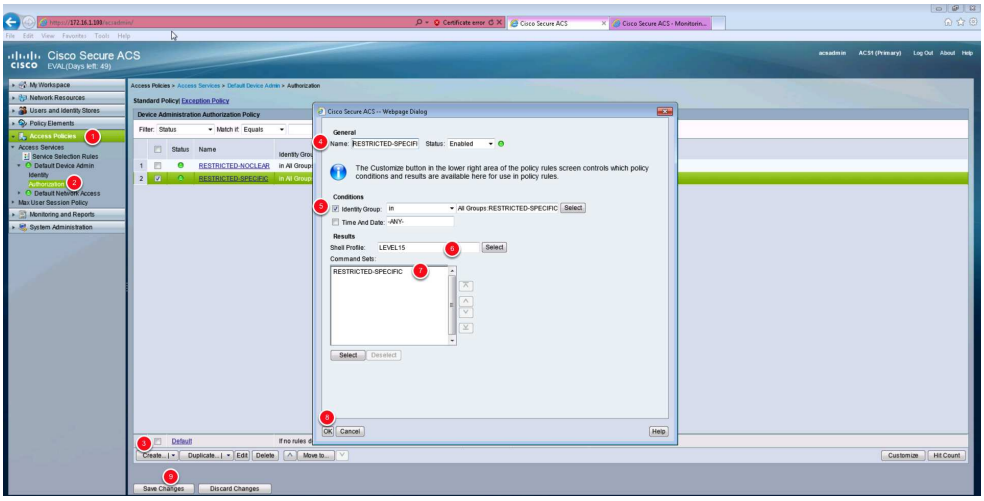
For **Default Device Admin** access policy, where by default all incoming TACACS sessions are redirected, add **Identity Group** to the available conditions and **Command Sets** to the available results:

https://t.me/learningnets



Configure the access policies for both users to match task requirements:





Verification

Telnet and authenticate using user **USER2**, and verify that you are required to change the password and that trying to use **cisco** as the password fails, as well as trying to use a length lower than 5 characters (note that no error messages are

https://t.me/learningnets

given; it just asks again for authentication):

```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open

User Access Verification
Username: USER2
Password: ******Enter new password: *****
Enter new password confirmation: *****
Username: USER2
Password: ******
Enter new password: ****
Enter new password confirmation: ****
Username: USER2
Password: ******
Enter new password: ******
Enter new password confirmation: ******
Type help or '?' for a list of available commands.ASA3> show curpriv
Username : USER2Current privilege level : 1
Current Mode/s : P_UNPRASA3> enable
Password: ******ASA3#
```

Verify that **USER2** is authorized to use all but **clear** type of commands:

```
ASA3# configure terminal
ASA3(config)# show running-config telnet
telnet 136.1.37.0 255.255.255.0 VLAN37
telnet timeout 5ASA3(config)# router eigrp 100
Too many IP routing processes for this routing protocol
ERROR: Unable to create router processASA3(config)# clear counters
Command authorization failed
```

Telnet and authenticate using user **USER1**, and verify that the user is authorized to use only commands specified by the task:

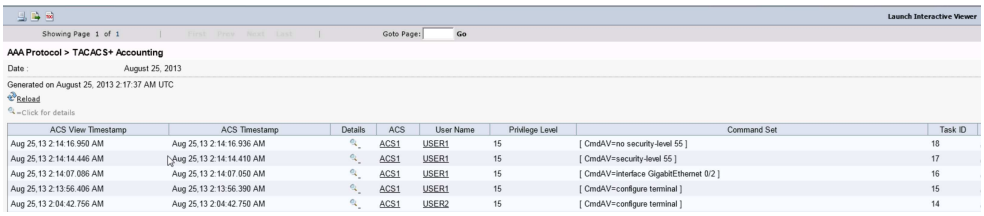
```
R3#telnet 136.1.37.17
Trying 136.1.37.17 ... Open

User Access Verification
Username: USER1
Password: ******
Enter new password: ******
```

https://t.me/learningnets

```
Enter new password confirmation: ******
Type help or '?' for a list of available commands.ASA3> enable
Password: ******ASA3# show curpriv
Username : USER1Current privilege level : 15
Current Mode/s : P_PRIVASA3# configure terminal
ASA3(config)# router ospf 10
Command authorization failed
ASA3(config)# interface GigabitEthernet0/2
ASA3(config-if)# nameif TEST
Command authorization failed
ASA3(config-if)# security-level 55
ASA3(config-if)# no security-level 55
ASA3(config-if)# exit
```

On ACS, launch the Monitoring and Report Viewer and review the **TACACS Accounting** logs to verify your configuration:



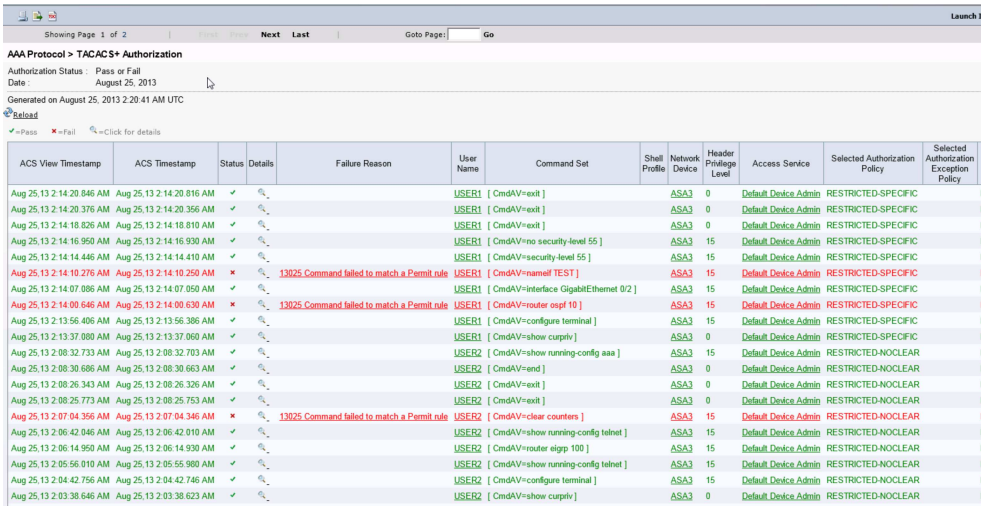
AAA Protocol > TACACS+ Accounting

Date : August 25, 2013

Generated on August 25, 2013 2:17:37 AM UTC

ACS View Timestamp	ACS Timestamp	Details	ACS	User Name	Privilege Level	Command Set	Task ID
Aug 25,13 2:14:16.950 AM	Aug 25,13 2:14:16.936 AM		ACS1	USER1	15	[CmdAV=no security-level 55]	18
Aug 25,13 2:14:14.446 AM	Aug 25,13 2:14:14.410 AM		ACS1	USER1	15	[CmdAV=security-level 55]	17
Aug 25,13 2:14:07.086 AM	Aug 25,13 2:14:07.050 AM		ACS1	USER1	15	[CmdAV=interface GigabitEthernet 0/2]	16
Aug 25,13 2:13:56.406 AM	Aug 25,13 2:13:56.390 AM		ACS1	USER1	15	[CmdAV=configure terminal]	15
Aug 25,13 2:04:42.756 AM	Aug 25,13 2:04:42.750 AM		ACS1	USER2	15	[CmdAV=configure terminal]	14

On ACS, launch the Monitoring and Report Viewer and review the **TACACS Authorization** logs to verify your configuration:



AAA Protocol > TACACS+ Authorization

Authorization Status : Pass or Fail

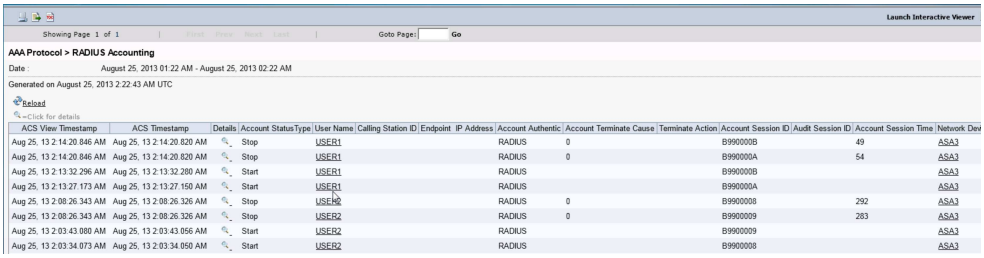
Date : August 25, 2013

Generated on August 25, 2013 2:20:41 AM UTC

ACS View Timestamp	ACS Timestamp	Status	Details	Failure Reason	User Name	Command Set	Shell Profile	Network Device	Header Privilege Level	Access Service	Selected Authorization Policy	Selected Authorization Exception Policy
Aug 25,13 2:14:20.846 AM	Aug 25,13 2:14:20.816 AM	✓			USER1	[CmdAV=exit]		ASA3	0	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:20.376 AM	Aug 25,13 2:14:20.356 AM	✓			USER1	[CmdAV=exit]		ASA3	0	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:18.826 AM	Aug 25,13 2:14:18.810 AM	✓			USER1	[CmdAV=exit]		ASA3	0	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:16.950 AM	Aug 25,13 2:14:16.930 AM	✓			USER1	[CmdAV=no security-level 55]		ASA3	15	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:14.446 AM	Aug 25,13 2:14:14.410 AM	✓			USER1	[CmdAV=security-level 55]		ASA3	15	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:10.276 AM	Aug 25,13 2:14:10.250 AM	✗		13025 Command failed to match a Permit rule	USER1	[CmdAV=nameif TEST]		ASA3	15	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:07.086 AM	Aug 25,13 2:14:07.050 AM	✓			USER1	[CmdAV=interface GigabitEthernet 0/2]		ASA3	15	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:14:00.646 AM	Aug 25,13 2:14:00.630 AM	✗		13025 Command failed to match a Permit rule	USER1	[CmdAV=router ospf 10]		ASA3	15	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:13:56.406 AM	Aug 25,13 2:13:56.386 AM	✓			USER1	[CmdAV=configure terminal]		ASA3	15	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:13:37.080 AM	Aug 25,13 2:13:37.060 AM	✓			USER1	[CmdAV=show curpriv]		ASA3	0	Default Device Admin	RESTRICTED-SPECIFIC	
Aug 25,13 2:08:32.733 AM	Aug 25,13 2:08:32.703 AM	✓			USER2	[CmdAV=show running-config aaa]		ASA3	15	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:08:30.686 AM	Aug 25,13 2:08:30.663 AM	✓			USER2	[CmdAV=end]		ASA3	0	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:08:26.343 AM	Aug 25,13 2:08:26.326 AM	✓			USER2	[CmdAV=exit]		ASA3	0	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:08:25.773 AM	Aug 25,13 2:08:25.753 AM	✓			USER2	[CmdAV=exit]		ASA3	0	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:07:04.356 AM	Aug 25,13 2:07:04.346 AM	✗		13025 Command failed to match a Permit rule	USER2	[CmdAV=clear counters]		ASA3	15	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:06:42.046 AM	Aug 25,13 2:06:42.010 AM	✓			USER2	[CmdAV=show running-config telnet]		ASA3	15	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:06:14.950 AM	Aug 25,13 2:06:14.930 AM	✓			USER2	[CmdAV=router eigrp 100]		ASA3	15	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:05:56.010 AM	Aug 25,13 2:05:55.980 AM	✓			USER2	[CmdAV=show running-config telnet]		ASA3	15	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:04:42.756 AM	Aug 25,13 2:04:42.746 AM	✓			USER2	[CmdAV=configure terminal]		ASA3	15	Default Device Admin	RESTRICTED-NOCLEAR	
Aug 25,13 2:03:38.646 AM	Aug 25,13 2:03:38.623 AM	✓			USER2	[CmdAV=show curpriv]		ASA3	0	Default Device Admin	RESTRICTED-NOCLEAR	

On ACS, launch the Monitoring and Report Viewer and review the **RADIUS Accounting** logs to verify your configuration:

https://t.me/learningnets



Launch Interactive Viewer

Showing Page 1 of 1 | First Prev Next Last | Goto Page: Go

AAA Protocol > RADIUS Accounting

Date : August 25, 2013 01:22 AM - August 25, 2013 02:22 AM

Generated on August 25, 2013 2:22:43 AM UTC

Reload

=Click for details

ACS View Timestamp	ACS Timestamp	Details	Account Status Type	User Name	Calling Station ID	Endpoint IP Address	Account Authentic	Account Terminate Cause	Terminate Action	Account Session ID	Audit Session ID	Account Session Time	Network Dev
Aug 25, 13 2:14:20.846 AM	Aug 25, 13 2:14:20.820 AM		Stop	USER1			RADIUS	0		B990000B		49	ASA3
Aug 25, 13 2:14:20.846 AM	Aug 25, 13 2:14:20.820 AM		Stop	USER1			RADIUS	0		B990000A		54	ASA3
Aug 25, 13 2:13:32.296 AM	Aug 25, 13 2:13:32.280 AM		Start	USER1			RADIUS			B990000B			ASA3
Aug 25, 13 2:13:27.173 AM	Aug 25, 13 2:13:27.150 AM		Start	USER1			RADIUS			B990000A			ASA3
Aug 25, 13 2:08:26.343 AM	Aug 25, 13 2:08:26.326 AM		Stop	USER2			RADIUS	0		B9900008		292	ASA3
Aug 25, 13 2:08:26.343 AM	Aug 25, 13 2:08:26.326 AM		Stop	USER2			RADIUS	0		B9900009		283	ASA3
Aug 25, 13 2:03:43.080 AM	Aug 25, 13 2:03:43.056 AM		Start	USER2			RADIUS			B9900009			ASA3
Aug 25, 13 2:03:34.073 AM	Aug 25, 13 2:03:34.050 AM		Start	USER2			RADIUS			B9900008			ASA3

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA Cut-Through Proxy Authentication

Task

Option 1

- Configure ASA3 to authenticate users before allowing telnet traffic through the firewall, for sessions initiated from behind VLAN 17.
- Use the local username database.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Cut-through proxy allows the ASA to authenticate and optionally authorize specific traffic before being allowed to pass through the firewall. Authorization gives you strict control over the traffic that each user is allowed to pass through the firewall. Authentication can be performed locally or through a remote RADIUS, TACACS, or LDAP server, whereas authorization can be performed only through RADIUS or TACACS. Regardless of the protocol being used, authorization cannot be performed without the user first authenticating.

ASA can perform the authentication in two ways:

- It can perform authentication transparently for protocols that support an

https://t.me/learningnets

authentication challenge inline, such as FTP, Telnet, HTTP, and HTTPS, but only on the default ports used by these protocols; so if telnet uses a non-standard port, the second method must be used. If a user tries to telnet to a destination for which the ASA requires authentication before allowing the traffic through, ASA prompts the user inline to first authenticate to it; if successful, the sessions toward the remote destination is allowed, and the user can be prompted for authentication by the remote server.

- If the above protocols do not use standard ports, or the ASA requires authentication for traffic that does not support an inline authentication challenge, such as ICMP, RDP, or SIP, users must first manually initiate a telnet or HTTP session to the ASA and authenticate for traffic to be allowed, and then initiate the ICMP, RDP, or SIP session toward the remote destination.

The second option supports two implementations methods by itself:

- Virtual HTTP or Virtual Telnet service must be enabled on the ASA, and users must telnet/HTTP/HTTPS to the Virtual IP address to authenticate.
- ASA can be configured to listen on a certain port for HTTP/HTTPS authentication requests, and users must manually browse to a well-known URL of **http://*IP-ADDRESS:PORT*/netaccess/connstatus.html**

Traffic to be authenticated and authorized can be identified by using the old statement of `match include` , where source/destination/service are configured inline, or the newer version of the `match` command, which requires an access-list. The only major functional difference between the two is that the old version does not allow use of the LOCAL user database; only remote AAA servers are supported. Because both methods achieve the same purpose, they are not allowed to be configured at the same time on the firewall.

Configuration

ASA3:

```
!
access-list TELNET permit tcp any any eq telnet
!
aaa authentication match TELNET VLAN17 LOCAL
username authproxy password authproxy
```

https://t.me/learningnets

Verification

Initiate a telnet session from SW1 destined to R3 and verify that ASA requires authentication before allowing the session through:

```
SW1#telnet 136.1.37.3
Trying 136.1.37.3 ... OpenUsername: authproxy
Password:
Password required, but none set
[Connection to 136.1.37.3 closed by foreign host]
!
!ASA3# show uauth
                     Current    Most Seen
Authenticated Users       1           2
Authen In Progress        0           2user 'authproxy' at 136.1.17.9, authenticated
   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00
!ASA3# clear uauth
```

Option 2

- Configure ASA3 to authenticate users before allowing them to pass ICMP traffic through the firewall, for sessions initiated from behind VLAN 17.
- Use the local username database.
- ICMP traffic sourced from a host with the MAC address of **AAAA.BBBB.CCCC** should be allowed through without prior authentication.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

https://t.me/learningnets

Overview

Because ICMP traffic does not support inline authentication, the telnet virtual service is used, so users must first telnet to the virtual IP address. To trigger the authentication process, traffic for the virtual telnet service must be matched by the authentication access-list. Telnet service is enabled by default on the ASA, so no extra configurations are required.

Configuration

```
ASA3:

!
virtual telnet 136.1.17.18
!
access-list ICMP permit icmp any any
access-list ICMP permit tcp any host 136.1.17.18 eq telnet
!
aaa authentication match ICMP VLAN17 LOCAL
username authproxy password authproxy
!
mac-list NOAUTH permit AAAA.BBBB.CCCC FFFF.FFFF.FFFF
aaa mac-exempt match NOAUTH
```

Verification

Try to ping from SW1 to R3; it should fail:

```
SW1#ping 136.1.37.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)
```

Telnet and authenticate to the ASA, and then verify that ICMP traffic is allowed:

```
SW1#telnet 136.1.17.18
Trying 136.1.17.18 ... Open
LOGIN Authentication Username: authproxy
```

https://t.me/learningnets

```
Password: Authentication Successful

[Connection to 136.1.17.18 closed by foreign host]
!SW1#ping 136.1.37.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
!ASA3(config)# show uauth
                    Current    Most Seen
Authenticated Users       1           2
Authen In Progress        0           2user 'authproxy' at 136.1.17.9, authenticated
   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00
!   ASA3(config)# clear uauth
```

Option 3

- Configure ASA3 to authenticate users before allowing them to pass HTTP traffic through the firewall, for sessions initiated from behind VLAN 17.
- Use the local username database.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA.

Configuration

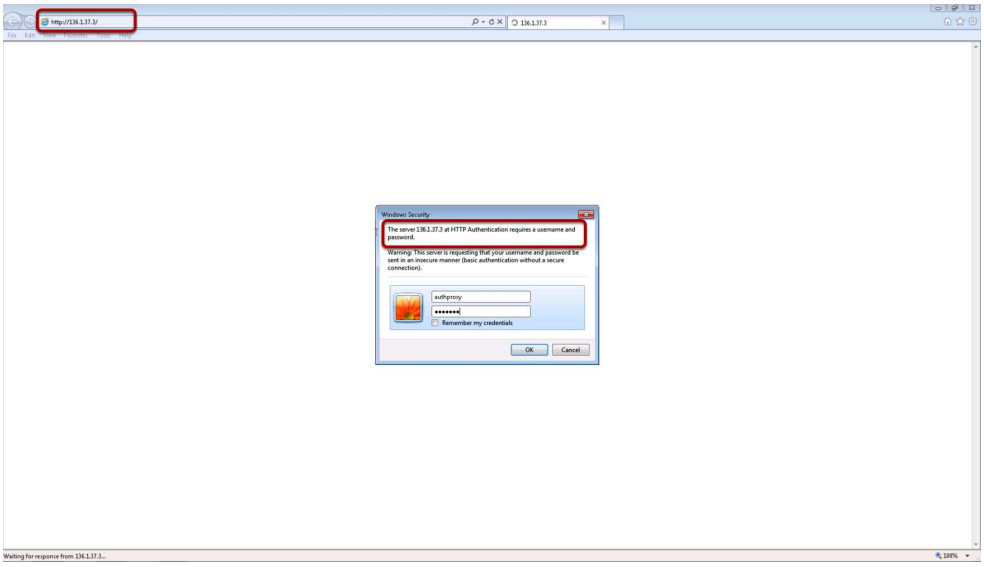
https://t.me/learningnets

ASA3:

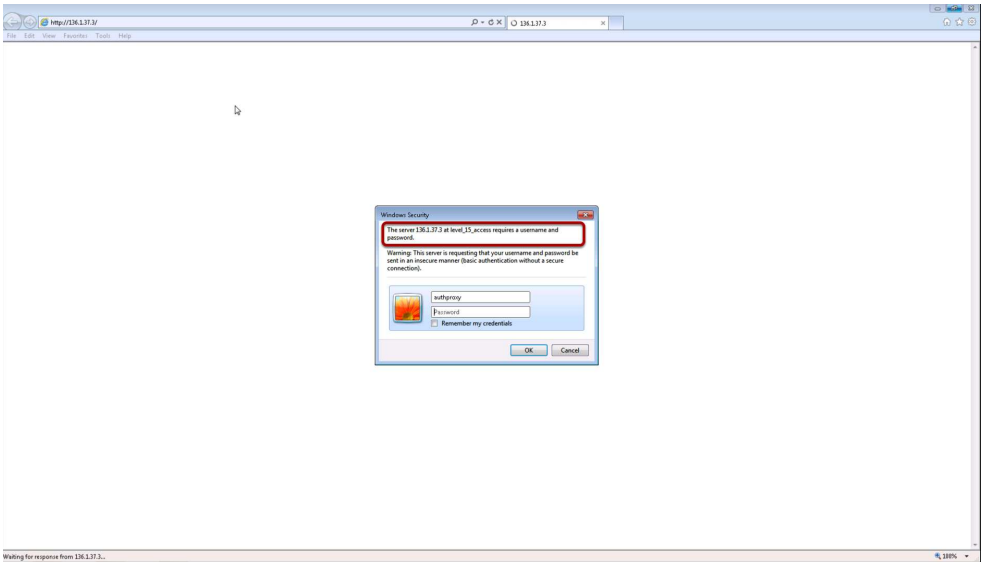
```
!
access-list HTTP permit tcp any any eq www
aaa authentication match HTTP VLAN17 LOCAL
username authproxy password authproxy
```

Verification

Initiate an HTTP session from TEST-PC-A destined to R3 and verify that ASA requires authentication before allowing the session through:



After successful authentication, traffic is redirected to R3, which asks for authentication:



https://t.me/learningnets

Verify the authentication on ASA3:

```
ASA3# show uauth
                     Current    Most Seen
Authenticated Users       1           2
Authen In Progress        0           2user 'authproxy' at 136.1.27.51, authenticated
   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00ASA3# clear uauth
```

Option 4

- Configure ASA3 to authenticate users before allowing them to pass ICMP traffic through the firewall, for sessions initiated from behind VLAN 17.
- Use the local username database.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Because ICMP traffic does not support inline authentication, the HTTP virtual service is used, so users must first HTTP to the virtual IP address. To trigger the authentication process, traffic for the virtual HTTP service must be matched by the authentication access-list. Also, unlike the virtual telnet service, the HTTP server must be enabled because it is disabled by default.

Configuration

```
ASA3:

!
```

https://t.me/learningnets

```
virtual http 136.1.17.18
!
access-list ICMP permit icmp any any
access-list ICMP permit tcp any host 136.1.17.18 eq www
aaa authentication match ICMP VLAN17 LOCAL
!
http server enable
username authproxy password authproxy
```

Verification

Try to ping from TEST-PC to R3; it should fail:

https://t.me/learningnets

Image not readable or empty
/uploads/workbooks/images/originals/6f4c7ea0a8f8dc868a5ba46fd4b07f5be339d644.png



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA Cut-Through Proxy RADIUS Authorization

Task

- Configure ASA3 to perform cut-through proxy authorization for traffic passing through the firewall from behind VLAN 17, using RADIUS.
- Users must authenticate by using telnet to the IP address of 136.1.17.100.
- Configure the ASA and ACS to match the following requirements:
  - User **user1** will be authorized for TCP traffic destined to R3's VLAN 37 IP address.
  - User **user2** will be authorized for ICMP traffic destined to R3's VLAN 37 IP address.
  - User **user3** will be authorized for both TCP and ICMP destined to R3's VLAN 37 IP address.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

Authorization gives you strict control over the traffic that each user is allowed to pass through the firewall, whereas performing only authentication allows you to apply the same traffic policy to all authenticated users. Authorization can be achieved through RADIUS or TACACS, but implementation is different because of the inherent differences between the two protocols.

Because RADIUS combines authentication and authorization in the same process (upon successful authentication, when the RADIUS server sends the ACCESS-ACCEPT message, authorization attributes also must be sent), cut-through proxy authorization requires only authentication to be configured on the ASA, and the authorization will be received from the RADIUS server in the format of an ACL. RADIUS can assign an access-list in three ways:

- **Downloadable ACL**, which requires the ACL to be configured on the RADIUS server and upon successful authentication is sent to the NAD.
- **Filter-ID ACL**, which requires the ACL to be configured on the NAD and upon successful authentication only the name of the ACL is sent to the NAD.
- **Proxy-ACL**, also known as Cisco AV Pair, which requires the ACL to be configured on the RADIUS server in a specific format and upon successful authentication is sent to the NAD.

Virtual telnet service must be configured for authentication to take place via telnet on 136.1.17.100. Note that if compared with the previous authentication task, the access-list used to trigger the authentication process contains only traffic destined to the virtual telnet service, because traffic allowed for the user will be received in the authorization process from the RADIUS server.

Access-list **user2** is configured in the ASA to use the **Filter-ID** method of delivering the authorization access-list. Also, without the **per-user-override** keyword, authorization rules received from the RADIUS server will not be applied to the interface ACL. Note that global ACL does not support this functionality.

Configuration

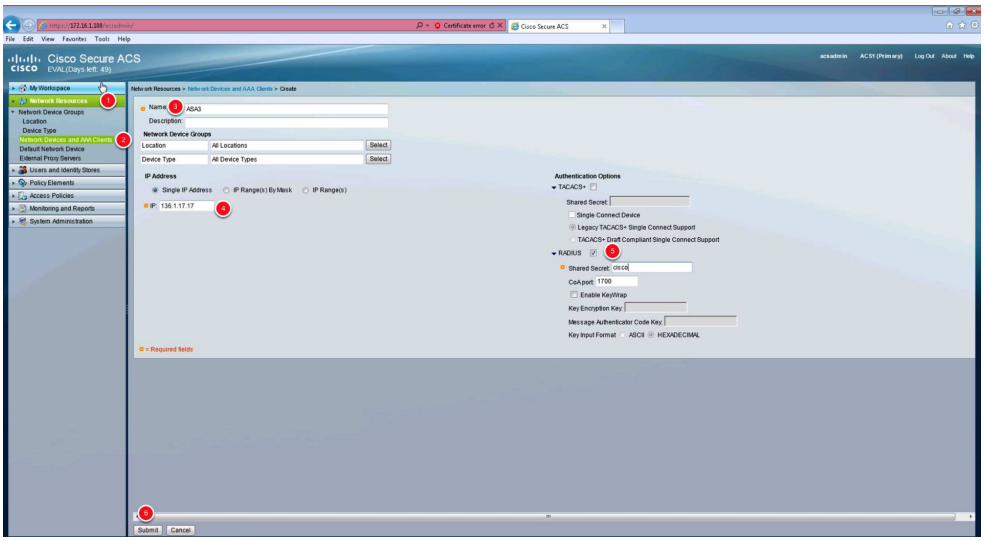
ASA3:

```
!
aaa-server RADIUS protocol radius
aaa-server RADIUS (VLAN17) host 172.16.1.100
 key cisco
!
virtual telnet 136.1.17.100
access-list AUTH permit tcp any host 136.1.17.100 eq telnet
!
aaa authentication match AUTH VLAN17 RADIUS
```

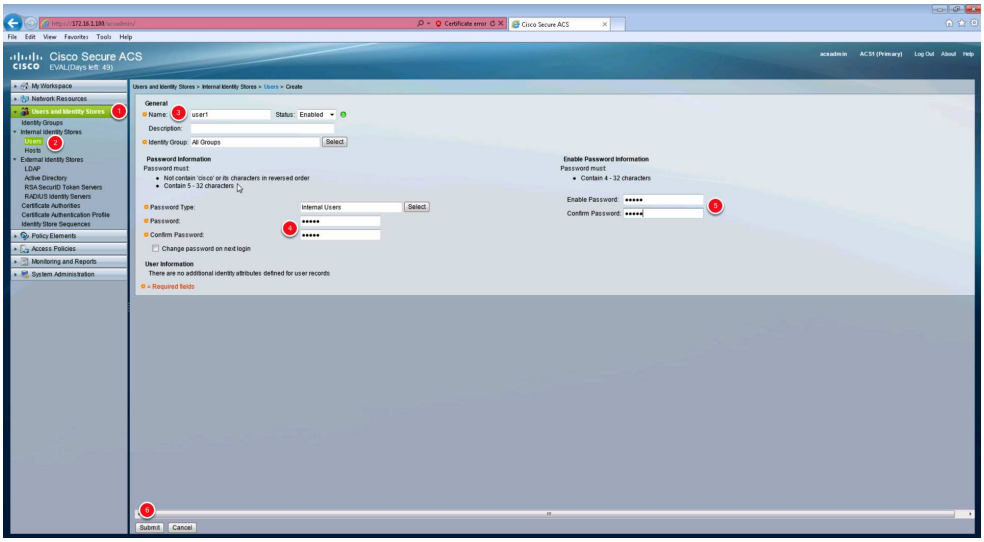
https://t.me/learningnets

```
!
access-list user2 permit icmp any host 136.1.37.3
!
access-list VLAN17_IN permit tcp any host 136.1.17.100 eq telnet
access-group VLAN17_IN in interface VLAN17 per-user-override
```

Add ASA3 as a device in ACS, and configure it for RADIUS only:

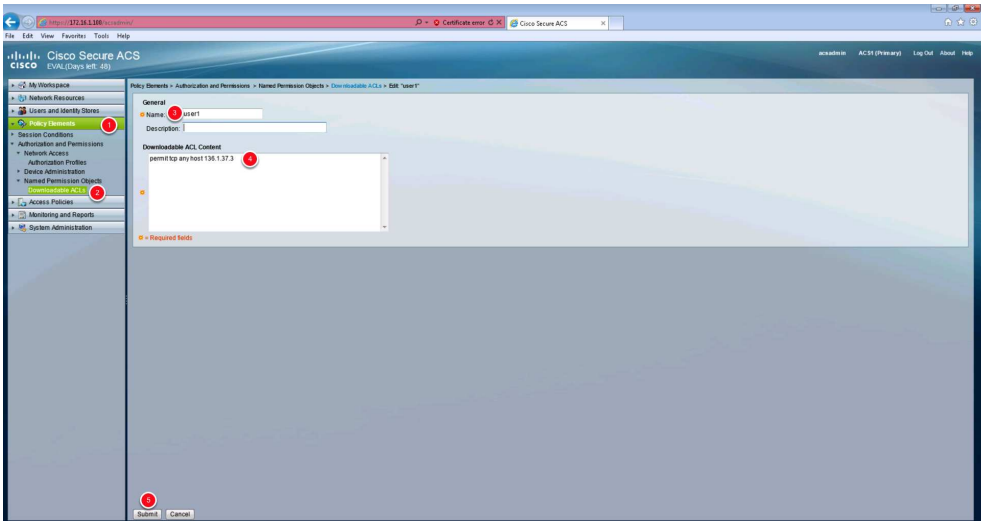


Configure all three users, with just the username and password. The following example is only for one user:

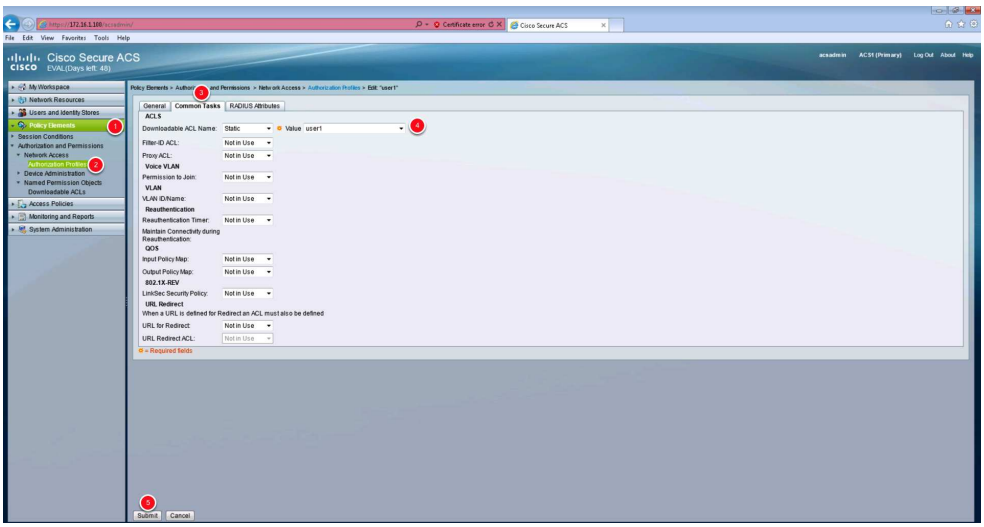


Configure the downloadable ACL that will be assigned to **user1** authorization profile:

https://t.me/learningnets

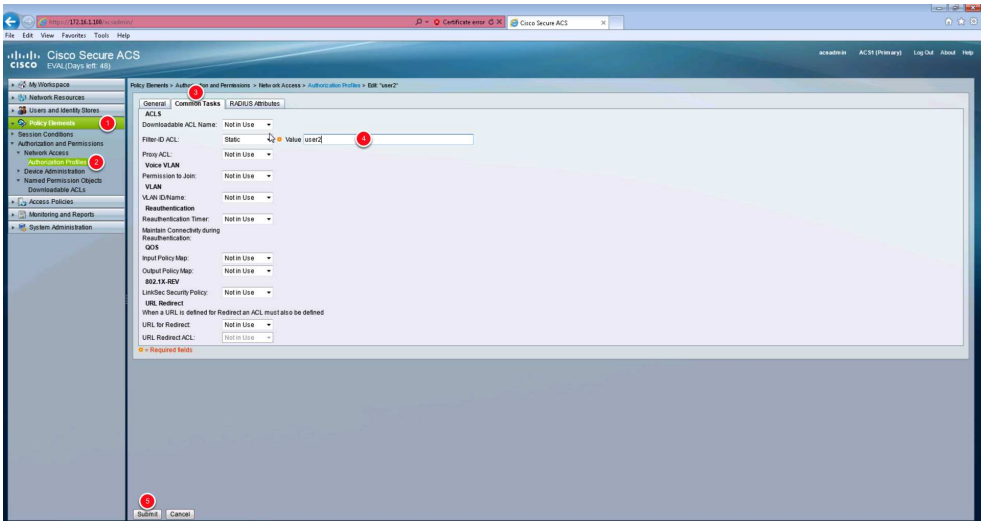


Configure authorization profiles for each user, with different settings. Note that even though ACS has a pre-configured **Proxy ACL** option in the **Common Tasks** section, while it uses the **cisco-av-pair** RADIUS attribute that is supported by the ASA, the format used by this option is for IOS and not for ASA, so ASA will not recognize the assigned value and reject it. For the ASA, the **cisco-av-pair** value must be in the format **ip:inacl#**, whereas IOS supports both **proxyacl#** and **ip:inacl#** formats. Note that there is also the option of using the **ip:outacl#** format, which means that the rule is applied in the outbound direction. Here is the authorization profile for user **user1**:

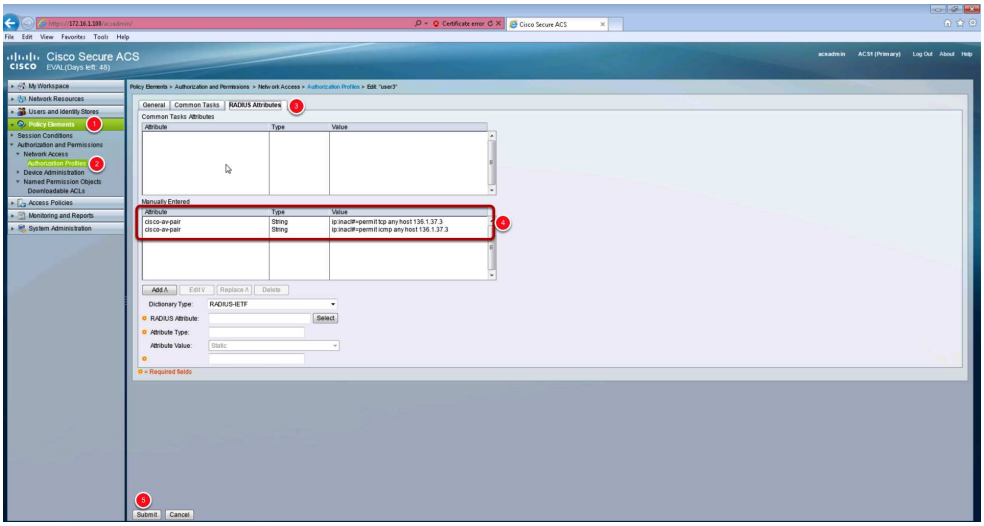


The authorization profile for user **user2**:

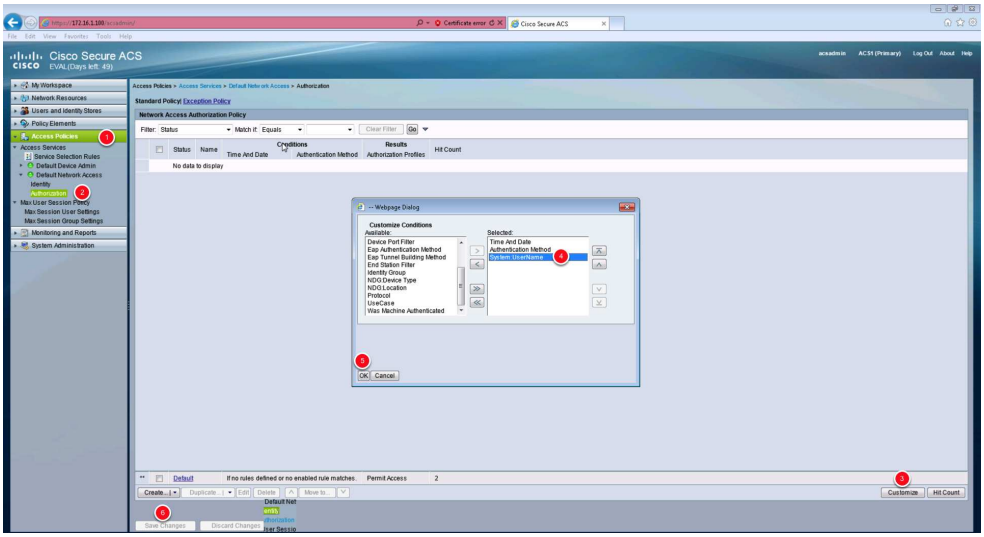
https://t.me/learningnets



The authorization profile for user **user3**:

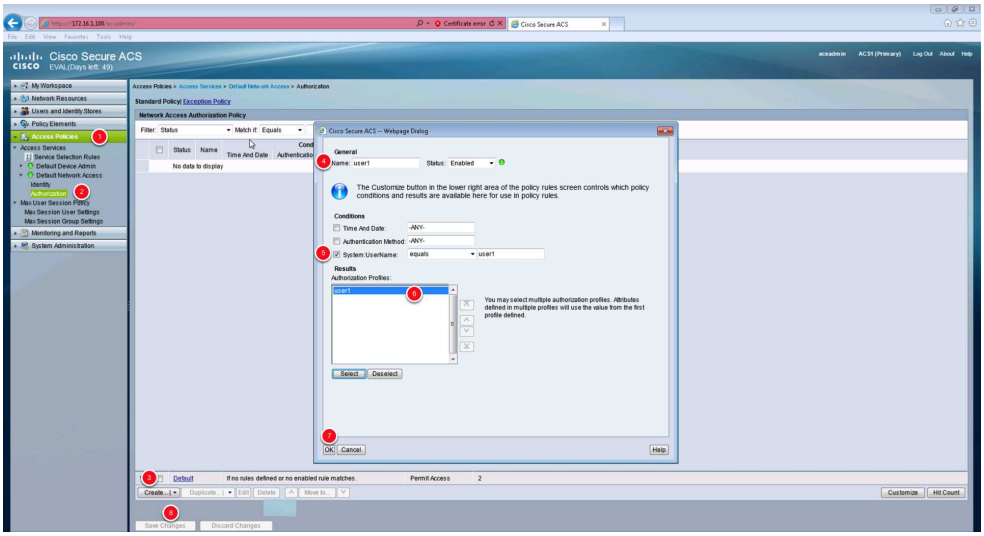


Add the **username** attribute as a condition in the authorization access policies:

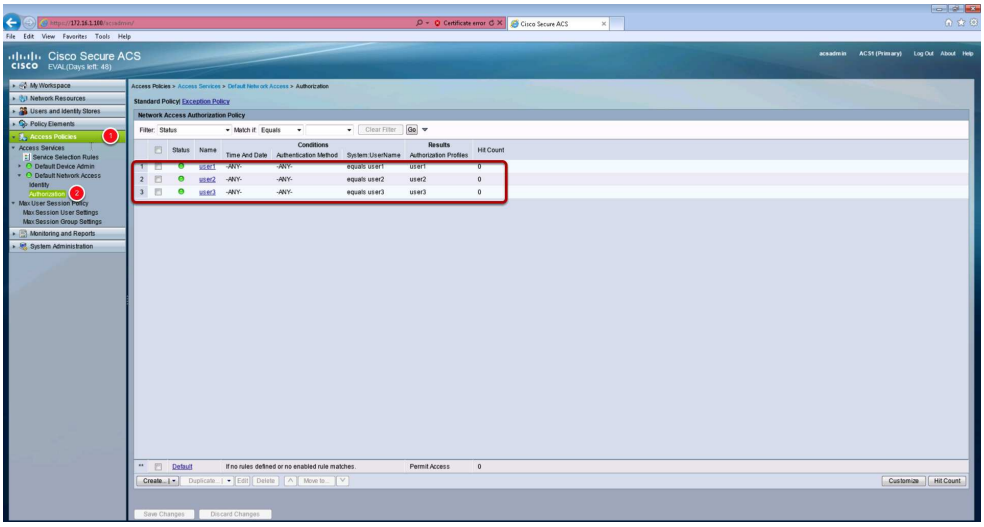


https://t.me/learningnets

Configure one access policy for each user, and apply the configured authorization profiles. Here is the access policy for user **user1**:



In the end, you should have three access policies, similar to this example:



Verification

Authenticate with **user1** from SW1, and verify what traffic is authorized afterward:

```
SW1#telnet 136.1.37.3
Trying 136.1.37.3 ... % Connection timed out; remote host not responding
!SW1#telnet 136.1.17.100
Trying 136.1.17.100 ... Open
LOGIN Authentication Username: user1
Password: Error:  acl authorization denied
[Connection to 136.1.17.100 closed by foreign host]SW1#telnet 136.1.37.3
Trying 136.1.37.3 ... Open Password required, but none set
[Connection to 136.1.37.3 closed by foreign host]
!
```

https://t.me/learningnets

```
!ASA3# show uauth

                    Current    Most Seen
Authenticated Users       1          2
Authen In Progress        0          2user 'user1' at 136.1.17.9, authenticated
access-list #ACSACL#-IP-user1-521a94a6 (*)
   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00
!ASA3# show access-list  #ACSACL#-IP-user1-521a94a6
access-list #ACSACL#-IP-user1-521a94a6; 1 elements; name hash: 0x75679416 (dynamic)
access-list #ACSACL#-IP-user1-521a94a6 line 1 extended permit tcp any host 136.1.37.3 (hitcnt=2) 0xfd95e552
ASA3# clear uauth
```

When authentication is successful against the virtual telnet service, the error **Error: acl authorization denied** is generated because traffic toward the virtual telnet service is not allowed in the downloaded ACL from the RADIUS server; however, the ACL is still applied and the result is the expected one. So this is basically a warning/informational message. However, for the CCIE lab, it would be wise to add this traffic to the downloadable ACL.

Now authenticate with **user2** from SW1, and observe what traffic is authorized afterward:

```
SW1#ping 136.1.37.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)SW1#telnet 136.1.17.100
Trying 136.1.17.100 ... Open
LOGIN AuthenticationUsername: user2

Password:
Error:  acl authorization denied

[Connection to 136.1.17.100 closed by foreign host]SW1#ping 136.1.37.3



Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2  seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
!
!ASA3# show uauth
                    Current    Most Seen
Authenticated Users       1          2
```

https://t.me/learningnets

```
Authen In Progress        0          2user 'user2' at 136.1.17.9, authenticated
 access-list user2 (*)
    absolute   timeout: 0:05:00
    inactivity timeout: 0:00:00
!   ASA3# show access-list user2
access-list user2; 1 elements; name hash: 0x77597ad3
access-list user2 line 1 extended permit icmp any host 136.1.37.3 (hitcnt=5) 0x256ebd16
ASA3# clear uauth
```

Authenticate with **user3** from SW1, and observe what traffic is authorized afterward:

https://t.me/learningnets

```
SW1#telnet 136.1.37.3

Trying 136.1.37.3 ... % Connection timed out; remote host not responding

SW1#ping 136.1.37.3


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)SW1#telnet 136.1.17.100

Trying 136.1.17.100 ... Open
LOGIN AuthenticationUsername: user3


Password:  Error:  acl authorization denied


[Connection to 136.1.17.100 closed by foreign host]SW1#ping 136.1.37.3


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 msSW1#telnet 136.1.37.3

Trying 136.1.37.3 ... OpenPassword required, but none set

[Connection to 136.1.37.3 closed by foreign host]
!
!ASA3# show uauth

                    Current    Most Seen
Authenticated Users       1          2
Authen In Progress        0          2user 'user3' at 136.1.17.9, authenticated
access-list AAA-user-user3-835BC30D (*)

   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00
!ASA3# show access-list AAA-user-user3-835BC30D
```

https://t.me/learningnets

```
access-list AAA-user-user3-835BC30D; 2 elements; name hash: 0xa8b7859a (dynamic)
access-list AAA-user-user3-835BC30D line 1 extended permit tcp any host 136.1.37.3 (hitcnt=2) 0xdc65ba52
access-list AAA-user-user3-835BC30D line 2 extended permit icmp any host 136.1.37.3 (hitcnt=5) 0xc4d62315
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ASA Cut-Through Proxy TACACS Authorization

Task

- Configure ASA3 to perform cut-through proxy authorization for traffic passing through the firewall from behind VLAN17, using TACACS.
- Authentication will be performed transparently inline through telnet.
- Configure ASA3 and ACS so that user **cut-through** is allowed to pass the following traffic through the firewall:
  - ICMP Echo traffic destined to R3's VLAN 37 IP address
  - Telnet traffic destined to R3's VLAN 37 IP address

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Authorization gives you strict control over the traffic that each user is allowed to pass through the firewall, whereas performing only authentication allows you to apply the same traffic policy to all authenticated users. Authorization can be achieved through RADIUS or TACACS, but implementation is different because of the inherent differences between the two protocols.

Because TACACS separates the authentication process from the authorization process, on the ASA we must match the traffic that triggers the authentication in the access-list used for cut-through proxy authentication, and traffic to be sent for

https://t.me/learningnets

authorization to TACACS servers in the access-list used for cut-through proxy authorization.

Cut-through proxy authorization is implemented with TACACS as follows: Traffic matching the authorization access-list is sent live/on-demand to the TACACS server for authorization, just as TACACS command authorization works. As we'll see, traffic to be authorized is actually defined in a **Command Set** policy element, but the syntax differs from how command authorization is configured; for cut-through proxy authorization, we define the protocol/service as the command and the destination as the command argument. Basically, when the ASA sends traffic to be authorized for cut-through proxy, it asks the TACACS server the following:

- Is this user authorized for shell access?
- Is this user authorized to use this command?

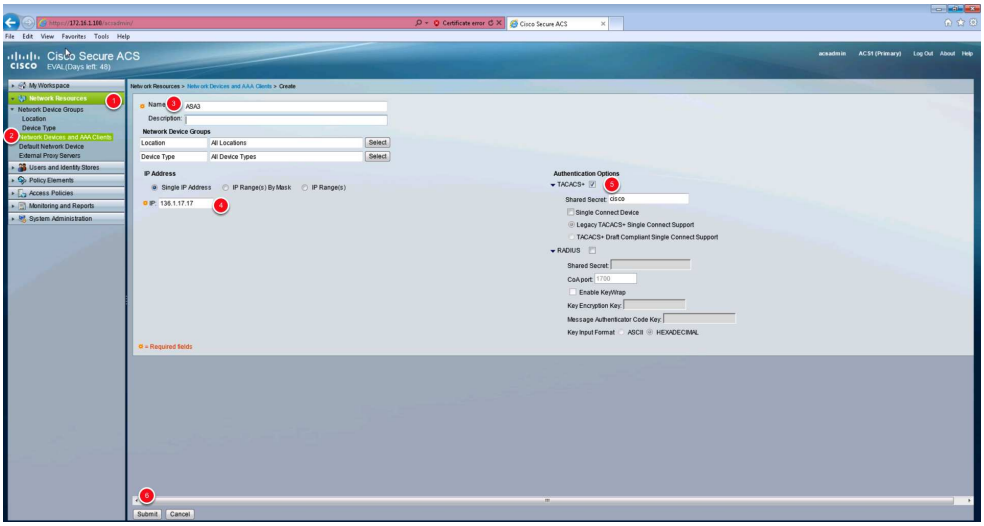
Configuration

ASA3:

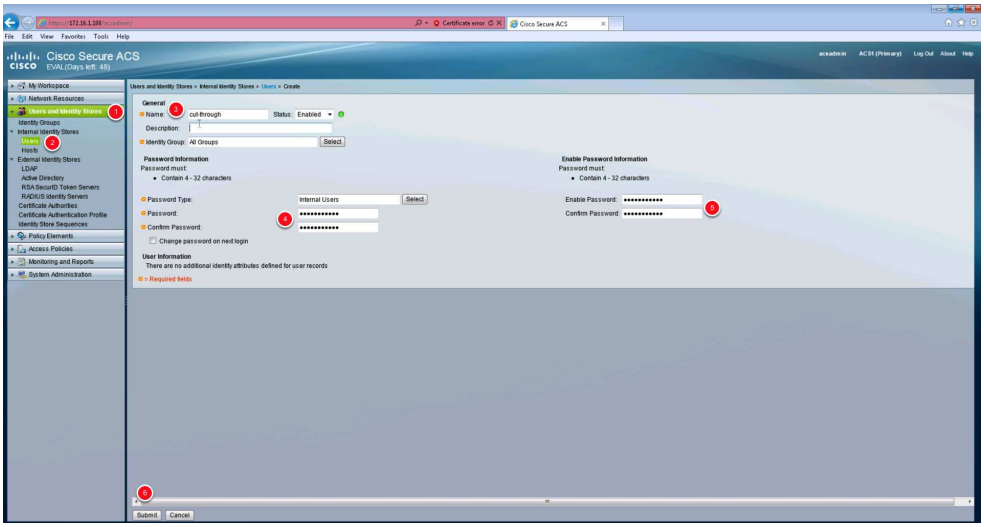
```
!
aaa-server TACACS protocol tacacs+
aaa-server TACACS (VLAN17) host 172.16.1.100
 key cisco
!
access-list AUTHENTICATION extended permit tcp any any eq telnet
access-list AUTHORIZATION extended permit tcp any host 136.1.37.3 eq telnet
access-list AUTHORIZATION extended permit icmp any host 136.1.37.3
!
aaa authentication match AUTHENTICATION VLAN17 TACACS
aaa authorization match AUTHORIZATION VLAN17 TACACS
```

Add ASA3 as a device into ACS, and configure it for TACACS:

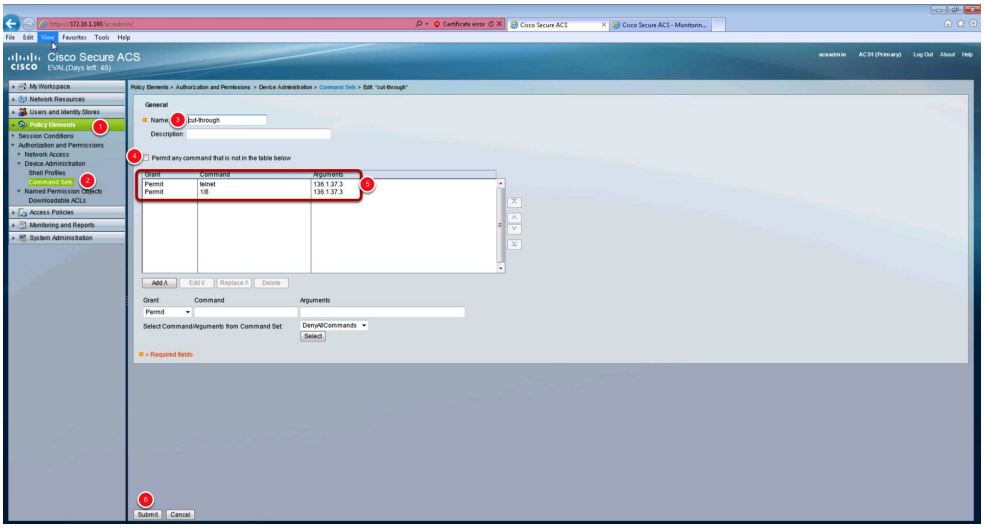
https://t.me/learningnets



Configure a user to be used for cut-through proxy authentication:



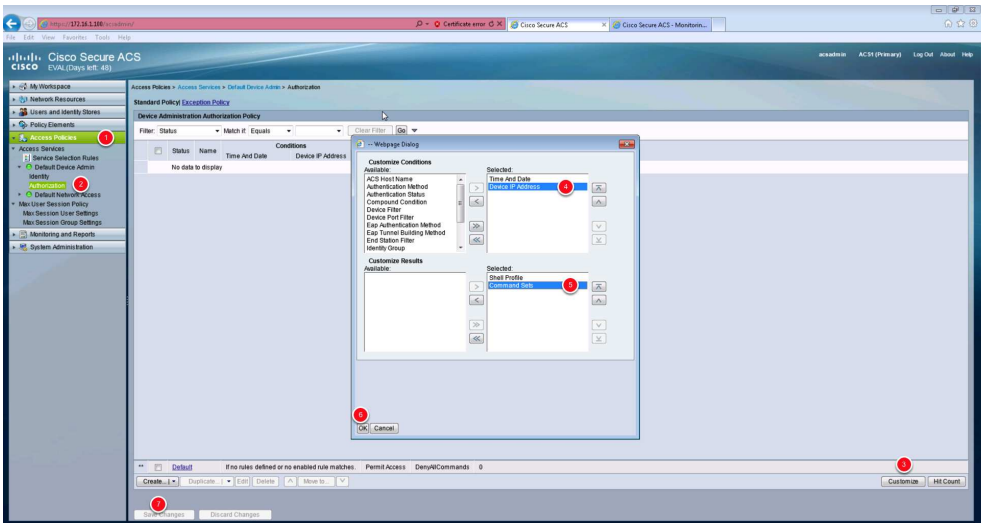
Configure the Command Set allowing telnet and ICMP echo destined to R3's VLAN 37 IP address:



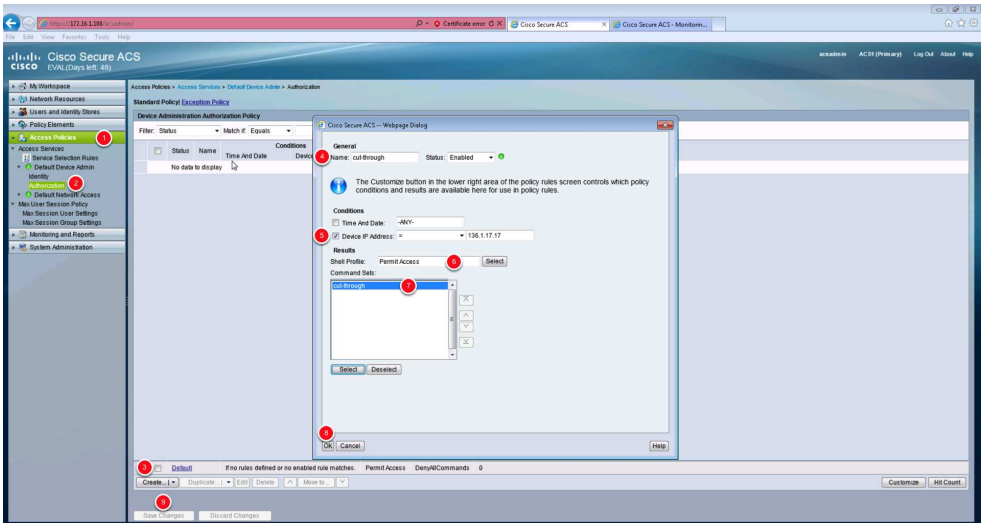
Add a condition to be used for the access policy and the command sets as available

https://t.me/learningnets

authorization result:



Configure an access policy allowing shell access and performing command authorization:



Verification

Verify that, for example, ICMP traffic toward R3 is not allowed prior to authentication:

```
SW1#ping 136.1.37.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)
```

Authenticate from SW1, and verify that telnet/ICMP traffic is allowed:

https://t.me/learningnets

```
SW1#telnet 136.1.37.3

Trying 136.1.37.3 ... Open
Username: cut-through

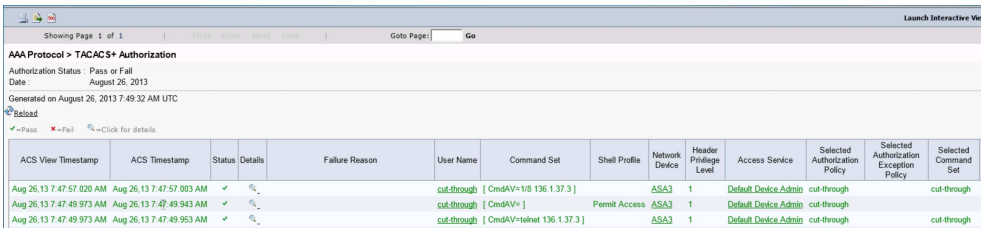
Password:
Password required, but none set

[Connection to 136.1.37.3 closed by foreign host]
!SW1#ping 136.1.37.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.37.3, timeout is 2 seconds:.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
!
!ASA3# show uauth
                    Current    Most Seen
Authenticated Users       1           1
Authen In Progress        0           1user 'cut-through' at 136.1.17.9, authorized to
port 136.1.37.3/telnet         136.1.37.3/icmp/8

   absolute   timeout: 0:05:00
   inactivity timeout: 0:00:00
```

Verify the TACACS Authorization reports/logs on ACS:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS Local Authentication

Task

- Configure SW1 for local authentication as follows:
  - Console sessions should be password authenticated and automatically placed into privilege exec mode upon successful authentication.
  - Telnet sessions should be user-based authenticated; **user1** should be automatically granted privilege exec mode and **user2** should need the **enable** password to be granted privilege exec mode.
- Configure R1 for local authentication as follows, by enabling AAA:
  - Do not configure any enable passwords or users.
  - Console and telnet should be authenticated using a user/password prompt.
  - VTY lines should automatically be granted privilege level 15 upon successful authentication.
  - Console should require the **enable** command before granting privilege exec access.
- Configure R3 for local authentication as follows, by enabling AAA:
  - No configuration should be done on the console, but it should remain in the default **no login** mode.
  - VTY lines access should be authenticated using a case-sensitive username-based authentication; if no users configured, access should be granted.
  - Access to privilege exec mode should require the enable password; if not available, access should be granted.
  - Ensure that **cisco1** is automatically placed into privilege level 15 upon authentication and its password is configured with the most secure method.
  - Ensure that **cisco2** needs the enable password to access privilege exec mode; upon authentication it should be granted privilege level 10.

A Note On Task Initial Configuration Files: For this task, you must

https://t.me/learningnets

load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

By default, without any AAA-related configurations applied, the IOS behaves as follows for management sessions:

- Does not require users to authenticate at the console, so it works in **no login mode**, just like ASA.
- Requires users to authenticate for telnet/SSH/HTTP, so it works in **login mode**, asking for password authentication, just like ASA.
- By default, telnet requires password authentication, but because there is no default password configured on the VTY lines, telnet connections will fail with the error message **Password required, but none set**; if password is configured with the `password <string>` command under VTY lines, the user will be placed into privilege level 1. To obtain privilege level 15, the `enable` command is required, but because there is no default enable password it will fail with the error message `% No password set` ; if the enable password is configured with the command `enable [password|secret] <string>` , the user will be placed into privilege level 15.
- If SSH is activated, which it is not by default, it will still not work because SSH requires username/password-based authentication. So it needs AAA to be configured, or VTY lines to authenticate via username/password, which is activated through the VTY command `login local` .
- If HTTP is activated, which it is not by default, it will ask for username/password authentication; however, it uses enable password authentication. It could still work using blank as the username and the enable password if this is configured, which it is not by default.

For AAA-related commands to become active, as well as global RADIUS and TACACS commands, the `aaa new-model` command must be configured. When this command is configured, VTY and HTTP access will require username/password-based authentication. To configure a login authentication method, the command `aaa authentication login [default|<name>]` is used; if the **default** keyword is used, the

https://t.me/learningnets

configured method applies for both VTY lines and console access, whereas if a named method is used, it must be manually applied using the console/VTY command `login authentication <name> .`

With the local user database, IOS supports only authentication and exec/command authorization for management sessions; command accounting is only supported using a remote TACACS server (RADIUS does not support command accounting). Local command accounting can be activated using a feature known as **Configuration Change Notification and Logging**, but this is not part of AAA sub-system; all entered commands will be saved locally on flash to provide a history of who (user) did what configuration changes (commands) and when (time).

To add users to the local database, use one of the following commands. The first one allows users to have a blank password, the second one is the regular way of configuring users, and the third option allows the user to authenticate one time only (after this, the user is removed from the configuration). The privilege-level is optional; if not configured, a default level of 1 is assigned to the user, which is different from ASA where a user has a default level of 2.

```
username <string> nopassword privilege <0-15>
username <string> privilege <0-15> [password|secret] <string>
username <string> privilege <0-15> one-time
```

When users are authenticated to console/telnet/ssh, they will be placed into privilege level 1, identified by the **>** character by default. If, however, the `aaa authorization exec [default|<name>]` command is configured, the user will be placed into the privilege level assigned to the username; for example, if the user has a configured privilege level of 15, when authenticated it will automatically be placed into privilege exec mode. From the user exec mode, which is privilege level 1, to be granted privileged exec mode, the following options are available:

- If the `aaa authentication enable default` command has not been configured, use the `enable` command and authenticate using the global enable password to be assigned privilege level 15.
- If the `aaa authentication enable default` command has been configured, use the `enable` command and authenticate using the specified method to be assigned privilege level 15.

https://t.me/learningnets

All IOS commands have a default privilege level assigned by default, just like ASA, and users will be authorized to use commands based on their assigned privilege-level. To perform command authorization, even if being granted privilege level 15, TACACS must be used.

Configuration

SW1:

```
!
username user1 privilege 15 secret user1
username user2 password user2
!
enable secret cisco
!
line vty 0 15
 login local
!
line console 0
 password cisco
 login
 privilege level 15
```

R1:

```
aaa new-model
aaa authentication login default local line
aaa authentication enable default enable line
!
line vty 0 15
 password vtypass
 privilege level 15
!
line console 0
 password consolepass
```

R3:

```
aaa new-model
aaa authentication login VTY local-case none
aaa authentication enable default enable none
aaa authorization exec VTY local
!
username cisco1 privilege 15 secret cisco1
```

https://t.me/learningnets

```
username cisco2 privilege 10 password cisco2
!
line vty 0 15
 login authentication VTY
 authorization exec VTY
```

Verification

Verify that SW1 console is authenticated and authorized as requested:

```
Press RETURN to get started.


User Access Verification
 Password:
SW1#show privilege
 Current privilege level is 15
```

Verify that **user1** and **user2** are correctly authorized when logging through telnet on SW1:

https://t.me/learningnets

```
SW1#telnet 150.1.9.9

Trying 150.1.9.9 ... Open

User Access Verification
Username: user1

Password: SW1#show privilege
Current privilege level is 15
SW1#exit


[Connection to 150.1.9.9 closed by foreign host]
!
!SW1#telnet 150.1.9.9

Trying 150.1.9.9 ... Open

User Access Verification
Username: user2

Password: SW1>show privilege
Current privilege level is 1
SW1>enable

Password: SW1#show privilege
Current privilege level is 15
```

Verify settings for VTY lines on R1; you can use any value for the username while the password is actually the line password, because there are no local users configured:

```
R1#telnet 150.1.1.1
Trying 150.1.1.1 ... Open

User Access Verification
Username: testuser
Password: R1#show privilege
Current privilege level is 15
```

https://t.me/learningnets

Verify settings for console on R1. Both username and password are actually the console password; the same is true for the enable password:

```
User Access Verification
Username: consolepass
Password:
R1>show privilege
Current privilege level is 1
R1>enable
Password: R1#show privilege
Current privilege level is 15
```

Verify R3's authentication and exec authorization for both users:

```
R3#telnet 150.1.3.3
Trying 150.1.3.3 ... Open

User Access Verification
Username: cisco1
Password:
R3#show privilege
Current privilege level is 15
!
!R3#telnet 150.1.3.3
Trying 150.1.3.3 ... Open

User Access Verification
Username: cisco2
Password:
R3#show privilege
Current privilege level is 10
R3#enable
R3#show privilege
Current privilege level is 15
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS Local Command Authorization

Task

- Configure AAA on R3 for local command authorization to allow **user1** access at privilege level 10 to:
  - Use RIP debugging commands and disable any currently active debugging using a single command
  - Configure any interface IP settings and administratively enable or disable interfaces
  - See only commands it has access to configure in the running-configuration
- Configuration should only affect VTY lines.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IOS command authorization can be performed locally or remotely through a TACACS server. With TACACS, the server is queried for each command the user types; with local, the privilege level of the command being typed is compared against the privilege level assigned to the user's session.

Local command authorization uses the concept of privilege levels. There are sixteen levels supported, 0 to 15. Every subsequent level includes the commands assigned to every previous level, so privilege 5 includes commands assigned to levels 0-5,

https://t.me/learningnets

and privilege 15 includes levels 0-15. By default, all IOS commands are assigned to one of the following three privilege levels:

- **Level 0**, with just a few basic commands, such as enable, login, and exit
- **Level 1**, the default exec user level; has some show commands available, but no configuration commands
- **Level 15**, the maximum privilege level, also known as privileged mode or enable mode; includes all the commands available in IOS

Local command authorization is always on by default and works according to the privilege levels assigned to the users and commands associated with the levels. To create custom command sets, the following options are available:

- Assign some level 15 commands to level 1, effectively making them available to all users who may log in to the router (if the user has the default privilege level settings).
- Re-assign some commands from level 1 to a higher level, thus disallowing all unprivileged users the use of those commands.
- Assign some level 15 commands to a new custom level, such as level 6, for example. After that, you can assign the custom privilege level to a specific user, allowing the use of some privileged commands to this particular user only, or to all users authorized for that privilege level.

To demonstrate how IOS performs command authorization, let's look at the command structure `command sub-command [arguments] [argument-values] [options]`. Here, **command** is the first portion of the command string—for example, **ip** in the **ip address** command entered under the interface configuration mode. The sub-command field makes the command more specific and usually is not available for all commands. The **arguments** list covers all mandatory named parameters that might have values assigned. For example, in the **ip address** command, the **address** field is an argument and it may take a value such as **1.2.3.4**. The **options** may cover various command attributes that are not mandatory. From the local command authorization standpoint, you can only match the mandatory fields such as command, sub-command, and arguments. The system will automatically allow any argument values and options if the command that the user enters matches the configured pattern.

To re-assign a command to a different privilege level, use the command `privilege <mode> level <level> <command>`. This command tells the router shell to assign the command matching the string **command** to the level specified by the **level** argument. The match occurs against all mandatory parts of the command that a user enters in a particular exec mode. For example, if you assigned the command **snmp-server** but not the command **snmp-server host** to the privilege level 7, a

https://t.me/learningnets

user will not be able to configure the SNMP traps destination because **host** is a mandatory (non-optional) part of the command. The following features ease the local command authorization configuration:

- When you enter commands as a shortcut, such as `privilege exec level 7 conf t` , the shell automatically expands it to the full syntax, thus to the string `privilege exec level 7 configure terminal .`
- When you assign a compound command to a particular level, such as `privilege interface level 7 ip address` , the shell automatically adds extra lines assigning all initial components of the compound command to the same level, so it adds the `privilege interface level 7 ip` command.
- You can use the special keyword **all** to allow any sub-commands behind a specific keyword. For example, the command `privilege interface all level 7 ip` allows all **ip** sub-commands in the interface configuration mode, such as `ip address` , `ip proxy-arp` , and `ip redirects .`
- When issuing the `show running-config` command, a user will only see the commands authorized to use. For example, if a user is authorized to use the `interface` command and all `ip` sub-commands, the `show running-config` command will only display those specific commands.

Configuration

R3:

```
aaa new-model
aaa authentication login VTY local
aaa authorization exec VTY local
!
username user1 privilege 10 password user1
!
line vty 0 15
 login authentication VTY
 authorization exec VTY
!
privilege exec level 10 undebug all
privilege exec level 10 debug ip rip
!
privilege exec level 10 configure terminal
privilege configure level 10 interface
privilege interface level 10 shutdown
privilege interface level 10 no shutdown
privilege interface all level 10 ip
```

https://t.me/learningnets

```
!
privilege exec level 10 show running-config
```

Verification

Authenticate with **user1** and verify the assigned privilege level:

```
R3#telnet 150.1.3.3
Trying 150.1.3.3 ... Open

User Access Verification
Username: user1
Password:
R3#show privilege
Current privilege level is 10
```

Verify the commands that the user is allowed to use:

```
R3#debug ip eigrp
        ^% Invalid input detected at '^' marker.
R3#debug ip rip
RIP protocol debugging is on
!R3#undebug all
All possible debugging has been turned off
!R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#interface FastEthernet0/1
R3(config-if)#shutdown
R3(config-if)#no shutdown
R3(config-if)#ip address 10.10.10.3 255.255.255.0
R3(config-if)#end
!R3#show running-config

Building configuration...

Current configuration : 318 bytes
!
! Last configuration change at 15:01:18 UTC Tue Aug 27 2013 by  user1
!
boot-start-marker
boot-end-marker
!
!
interface Loopback0
```

https://t.me/learningnets

```
 ip address 150.1.3.3 255.255.255.255
!
interface FastEthernet0/0
 ip address 136.1.37.3 255.255.255.0
!
interface FastEthernet0/1
 ip address 10.10.10.3 255.255.255.0
!
!
end
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS Role-Based CLI Access

Task

- Configure the following user roles on R3:
  - **INTERFACE1** should have access to all IP configuration commands of interface FastEthernet0/0 only.
  - **INTERFACE2** should have access to all IP configuration commands of interface FastEthernet0/1 only.
  - **DEBUG** should be able to use any debug and undebug commands, and it should be able to inspect the running configuration.
  - **SUPER** should be a sum of all the other roles.
- Use the password **CISCO** whenever required.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Role-based CLI enhances the local command authorization model. Instead of using hierarchical privilege levels, you may define user roles in Cisco IOS, called **views**, and assign command sets accessible to each view. You may further group views in hierarchies using the concept of a **super-view** that inherits commands from other views instead of a list of the allowed commands. Each **view** is password-protected, and users are required to enter the respective password when switching to the view. Alternatively, a **view** may be associated with a user of the local database, or an

https://t.me/learningnets

attribute from external AAA server specifying the role name. Note that enabling AAA with the **aaa new-model** command is mandatory for the role-based access-control to work.

One special **root** view always exists. By default, even when in privileged exec mode, you are not part of the root view. You must explicitly switch to the root view to be able to create other views. This is the only difference from the regular enable privilege mode. You can switch to the root view using the command **enable view**, which actually enables the functionality of role-based CLI. Within the root view, you have privilege level 15 access. To access the **root** view, you must authenticate using the enable authentication method, which by default asks for the enable password.

A new view can be defined using the command `parser view <VIEW_NAME>`. Under the view configuration mode, you define the password for the view using the command `secret <password>` and define commands accessible to the view using a syntax similar to command authorization:

```
commands parser-mode {include|exclude|include-exclusive} [all] [interface interface-name] [command]
```

The **include-exclusive** means that the commands included in this view are not accessible to other views. By using the **interface** keyword, you can limit the interfaces accessible to the role, considering that the user has access to interface configuration mode. You can switch to any view using the command `enable view <VIEW_NAME>` from the router CLI. If you need to define a superview, use the command `parser view <NAME> superview` and add views using the nested command `view <SUBVIEW_NAME>`. Note that all views, including superviews, must have passwords defined to function.

Configuration

```
R3:

aaa new-model
enable password cisco
exit
!
enable view

configure terminal
!
parser view DEBUG
 secret CISCO
```

https://t.me/learningnets

```
 commands exec include show running-config
 commands exec include all debug
 commands exec include all undebug
!
parser view INTERFACE1
 secret CISCO
 commands interface include all ip
 commands configure include interface
 commands exec include configure terminal
 commands configure include interface FastEthernet0/0
!
parser view INTERFACE2
 secret CISCO
 commands interface include all ip
 commands configure include interface
 commands exec include configure terminal
 commands configure include interface FastEthernet0/1
!
parser view SUPER superview
 secret CISCO
 view DEBUG
 view INTERFACE1
 view INTERFACE2
```

Verification

Verify, for example, the **INTERFACE2** view:

```
R3#enable view INTERFACE2
Password:

R3#%PARSER-6-VIEW_SWITCH: user unknown successfully set to view 'INTERFACE2'.
R3#?
Exec commands:
  configure   Enter configuration mode
  credential  load the credential info from file system
  enable      Turn on privileged commands
  exit        Exit from the EXEC
  show        Show running system information
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#?
Configure commands:
  do         To run exec commands in config mode
  exit       Exit from configure mode
  interface  Select an interface to configure
```

https://t.me/learningnets

```
R3(config)#interface FastEthernet0/0

                           ^% Invalid input detected at '^' marker.
R3(config)#interface FastEthernet0/1
R3(config-if)#?
Interface configuration commands:
  exit  Exit from interface configuration mode
  ip    Interface Internet Protocol config commands
R3(config-if)#ip address 10.10.10.10 255.255.255.0
R3(config-if)#no shutdown
          ^% Invalid input detected at '^' marker.
R3(config-if)#exit
R3(config)#exit
```

Verify the superview:

```
R3#enable view SUPER
Password:

R3#%PARSER-6-VIEW_SWITCH: user unknown successfully set to view 'SUPER'.
R3#?
Exec commands:
  configure   Enter configuration mode
  credential  load the credential info from file system
  debug       Debugging functions (see also 'undebug')
  enable      Turn on privileged commands
  exit        Exit from the EXEC
  show        Show running system information
  undebug     Disable debugging functions (see also 'debug')
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#interface FAstEthernet0/0
R3(config-if)#exit
R3(config)#interface FastEthernet0/1
R3(config-if)#exit
R3(config)#exit
R3#show running-config

Building configuration...

Current configuration : 212 bytes
!
! Last configuration change at 16:22:36 UTC Tue Aug 27 2013
!
!
```

https://t.me/learningnets

```
!
interface FastEthernet0/0
 ip address 136.1.37.3 255.255.255.0
!
interface FastEthernet0/1
 ip address 10.10.10.10 255.255.255.0
!
!
end
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS RADIUS Authentication

Task

- Configure AAA on R1 using RADIUS with fallback on local user database as follows:
  - Authenticate console, VTY, and privileged exec mode access.
  - RADIUS communication should use Loopback0 interface.
  - Account for exec session initiation and termination.
  - User **user1** should be assigned privilege level 5; additional access to privilege level 7 or privilege level 15 should be possible.
  - User **user2** should be assigned privilege level 15.
- Configure AAA on R3 using RADIUS as follows:
  - Authenticate VTY and privileged exec mode access.
  - RADIUS communication should use Loopback0 interface.
  - Do not use global commands for defining RADIUS communication with ACS.
  - User **user1** should be assigned privilege level 3; additional access to privilege level 7 or privilege level 15 should be possible.
  - User **user2** should not be allowed to authenticate.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Unlike TACACS+ protocol, RADIUS does not implement separate authentication

https://t.me/learningnets

and authorization phases. When a client sends authentication requests to the AAA server, the server returns a set of RADIUS attributes that are used to authorize the particular service requested. Various TACACS+ attributes have been mapped to RADIUS using VSA (Vendor Specific Attributes), also known as Cisco AV Pair. Cisco AV Pair syntax is usually in the format of `<service>:<tacacas-av-pair>` , like **shell:priv-lvl** or **ip:inacl**. For example, if you want exec authorization via RADIUS, you may use the AAA command `aaa authorization exec default radius` and associate the AV-pair **shell:priv-lvl=15** with the respective authorization profile. For the purpose of exec shell authorization, you may also set the IETF RADIUS attribute **Service-Type** to the value **Administrative**, as supported by the ASA. This will automatically authorize the respective user to login with privilege level of 15.

For the enable privilege authentication, IOS will send to the RADIUS server usernames in the format of `$enab<N>$` to authenticate the enable password. For example, create user `$enab15$` to authenticate the maximum enable privilege level of 15. RADIUS does not support per-user enable passwords as TACACS+ does, nor does it support per-command authorization. Additionally, the RADIUS server does not supply its own AAA banner messages, so the ones configured locally would take effect even with remote authentication.

Similar to the TACACS+ protocol settings, you may configure named AAA lists and define groups of RADIUS servers using the command `aaa group server` . The default RADIUS server group is the one configured using the commands `radius-server host` and `ip radius source-interface` from the global configuration mode. Also, RADIUS does not support command accounting, but does support exec session accounting.

Configuration

R1:

```
aaa new-model
!
aaa authentication login default group radius local
aaa authentication enable default group radius enable
!
aaa authorization exec default group radius local
aaa authorization console
!
aaa accounting  exec default start-stop group radius
!
enable secret level 7 enable7
enable secret level 15 enable15
!
username user1 privilege 5 password cisco5
username user2 privilege 15 password cisco15
```

https://t.me/learningnets

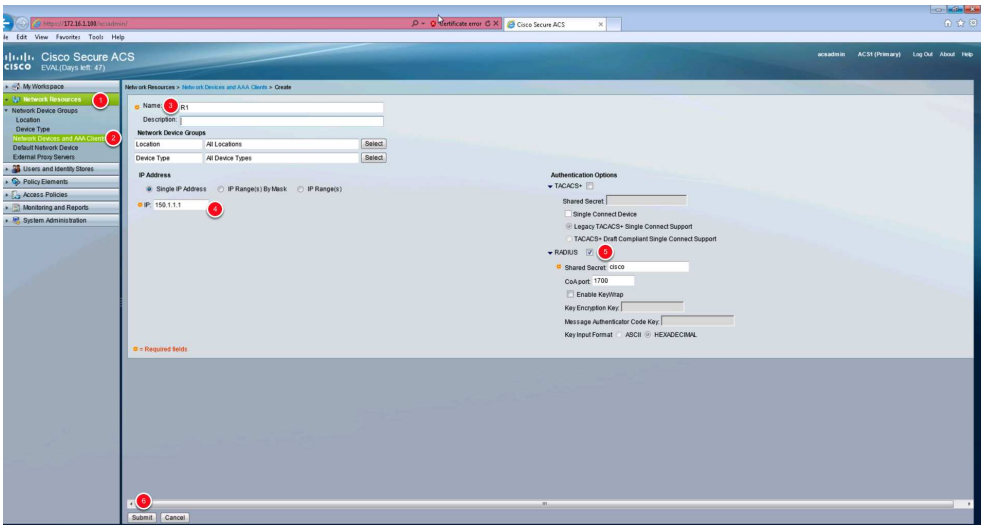
```
!
radius-server host 172.16.1.100 key cisco
ip radius source-interface Loopback0


R3:

aaa new-model
!
aaa group server radius ACS
 server-private 172.16.1.100 key 0 cisco
 ip radius source-interface Loopback0

aaa authentication login CONSOLE none
aaa authentication login default group ACS local
aaa authentication enable default group ACS enable
!
aaa authorization exec default group ACS local
!
enable secret level 7 enable7
enable secret level 15 enable15
!
username user1 privilege 3 password cisco3
!
line con 0
 login authentication CONSOLE
 privilege level 15
```

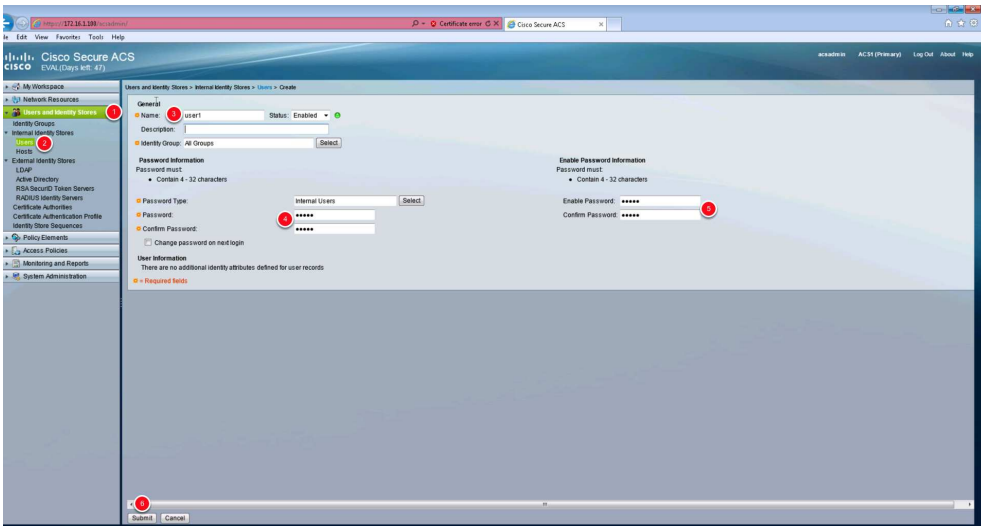
On ACS, first add both routers as NAD, use Loopback0 addresses, and configure it only for RADIUS; the following example is only for R1:



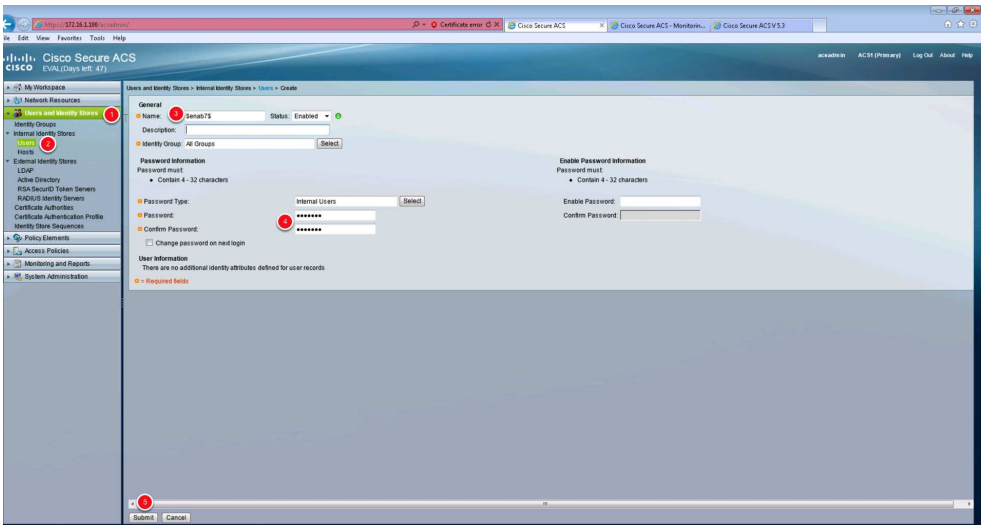
Add both **user1** and **user2** to the Internal User store; these will be used for login

https://t.me/learningnets

access authentication. The following example is only for **user1**:

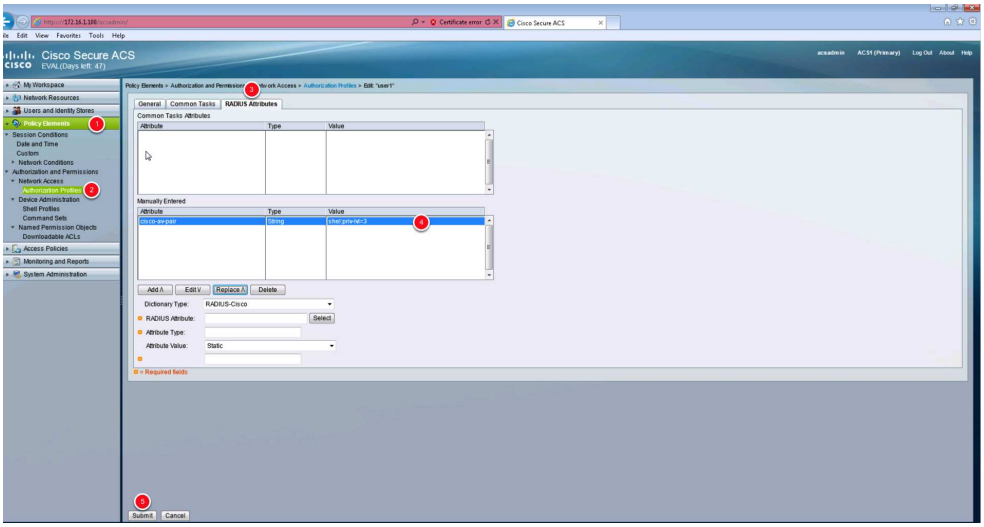


Add both **$enab7$** and **$enab15$** to the Internal user store; these will be used for enable access authentication:

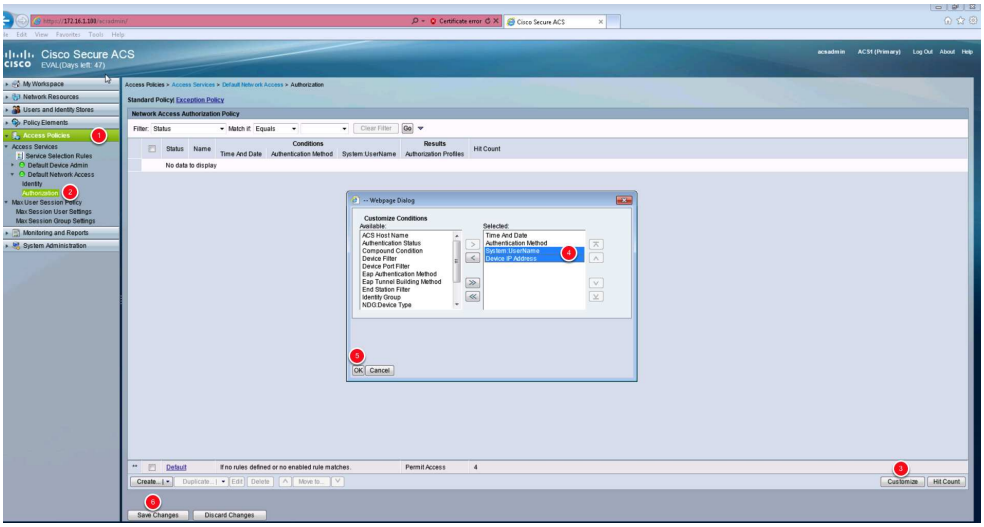


Configure three RADIUS authorization profiles to assign privilege levels 3, 5 and 15; the following example is only for privilege-level 3:

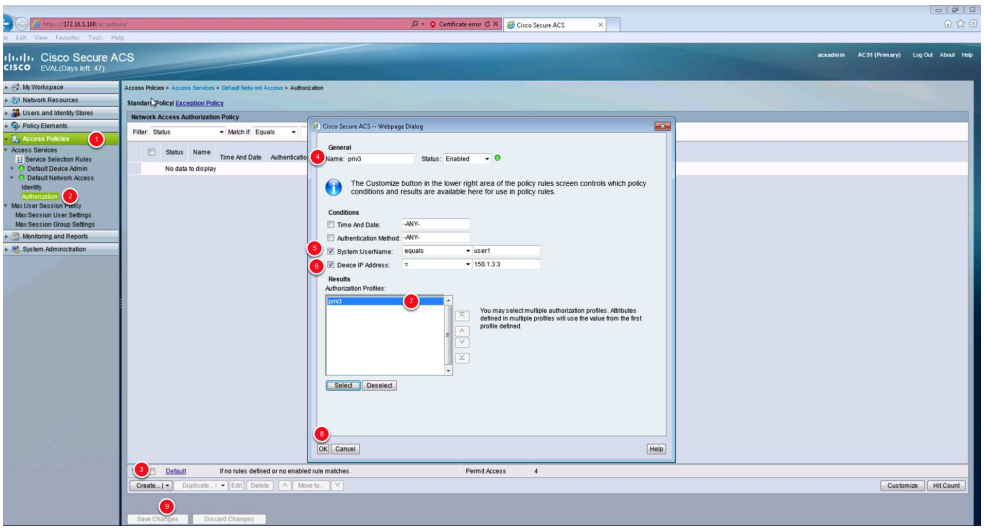
https://t.me/learningnets



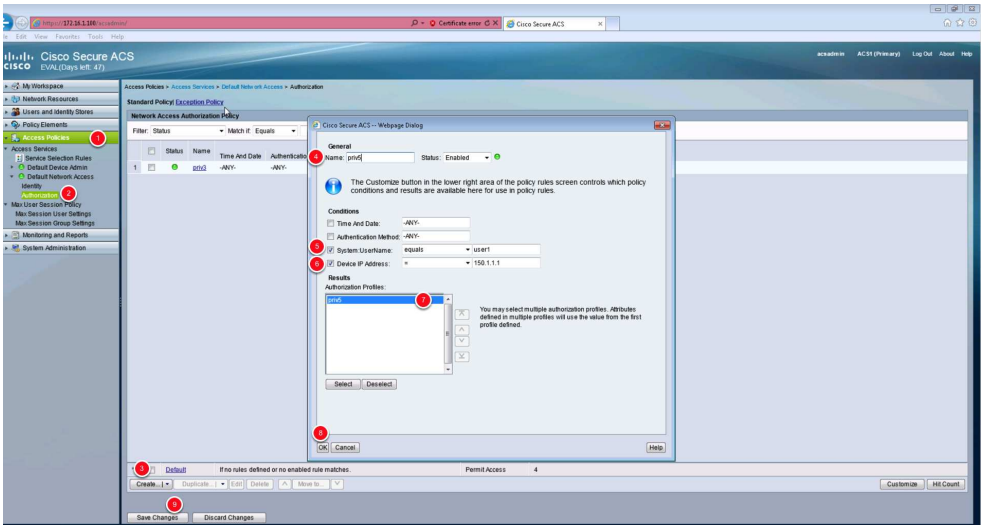
Add the NAD IP address and authenticating user as conditions in the RADIUS access policies:

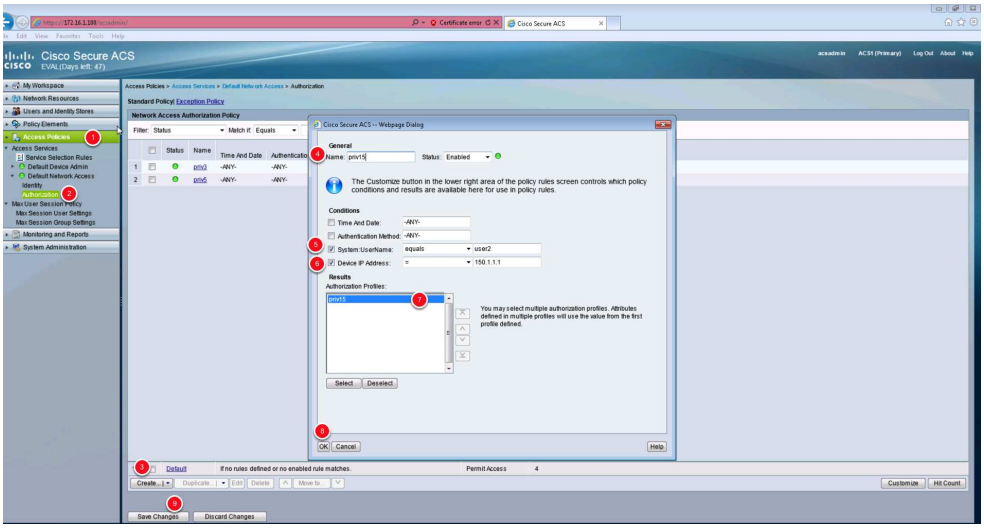


Configure three RADIUS authorization access policies to assign privilege levels as requested:

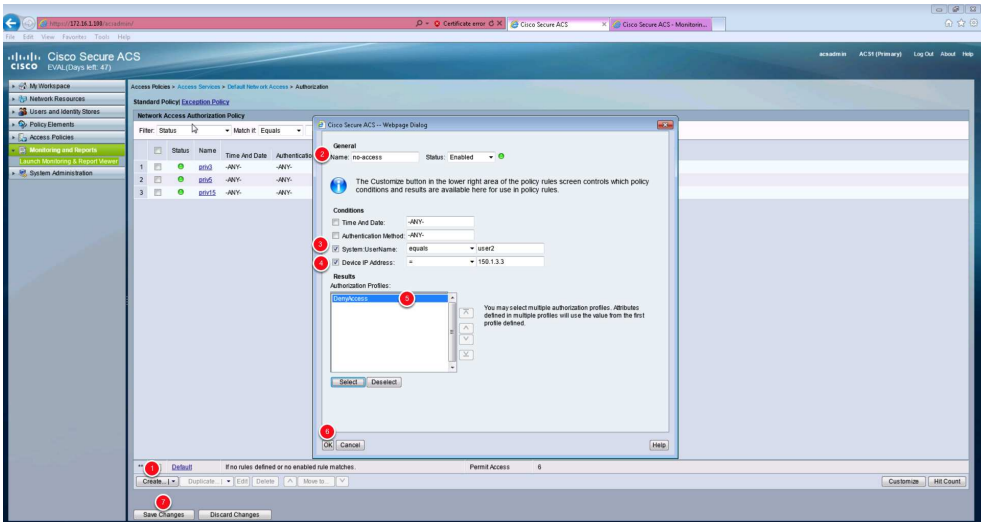


https://t.me/learningnets





Configure one RADIUS authorization policy so that **user2** cannot log in from R3; note that modifying the default RADIUS authorization policy from **PermitAccess** to **DenyAccess** is not an option, because the default policy is matched when **user1** is trying to access privileged exec mode with the **enable** command:



https://t.me/learningnets

Verification

Log in on R1 from the console using **user1**, verify the assigned privilege level, and confirm that it can access privilege level 7 and 15:

```
R1#login

User Access Verification
Username: user1
Password:
R1#show privilege
Current privilege level is 5
R1#enable 7
Password: R1#show privilege
Current privilege level is 7
R1#enable 15
Password: R1#show privilege
Current privilege level is 15
R1#configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
```

Log in on R1 from the console using **user2**, and verify the assigned privilege level:

```
R1#login

User Access Verification
Username: user2
Password:
R1#show privilege
Current privilege level is 15
R1#configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
```

Because R3's console does not ask for authentication, telnet and authenticate with **user1**, verify the assigned privilege level, and confirm that it can access privilege level 7 and 15:

```
R3#telnet 150.1.3.3
Trying 150.1.3.3 ... Open

User Access Verification
```

https://t.me/learningnets

```
Username: user1

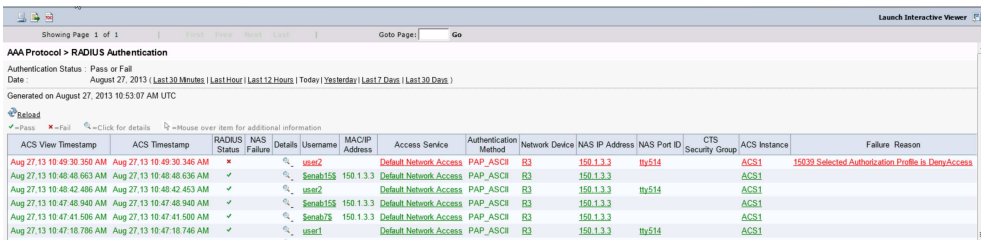
Password:
R3#show privilege
Current privilege level is 3
R3#enable 7
Password: R3#show privilege
Current privilege level is 7
R3#enable 15
Password: R3#show privilege
Current privilege level is 15
R3#configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
```

Verify that **user2** cannot authenticate via RADIUS on R3, and check the RADIUS Authentication logs from ACS:

```
R3#telnet 150.1.3.3
Trying 150.1.3.3 ... Open


User Access Verification
Username: user2
Password:
% Authentication failed
```



AAA Protocol > RADIUS Authentication

Authentication Status : Pass or Fail
Date : August 27, 2013 (Last 30 Minutes | Last Hour | Last 12 Hours | Today | Yesterday | Last 7 Days | Last 30 Days)

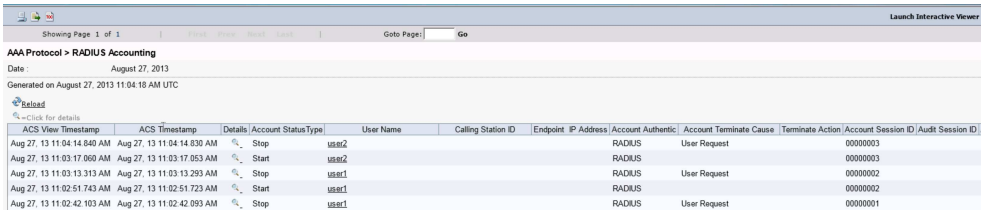
Generated on August 27, 2013 10:53:07 AM UTC

Reload

✔=Pass ✖=Fail 🔍=Click for details 🖱=Mouse over item for additional information

ACS View Timestamp	ACS Timestamp	RADIUS Status	NAS Failure	Details	Username	MAC/IP Address	Access Service	Authentication Method	Network Device	NAS IP Address	NAS Port ID	CTS Security Group	ACS Instance	Failure Reason
Aug 27,13 10:49:30.350 AM	Aug 27,13 10:49:30.346 AM	✖		🔍	user2		Default Network Access	PAP_ASCII	R3	150.1.3.3	tty514		ACS1	15039 Selected Authorization Profile is DenyAccess
Aug 27,13 10:48:48.663 AM	Aug 27,13 10:48:48.636 AM	✔		🔍	$enab15$	150.1.3.3	Default Network Access	PAP_ASCII	R3	150.1.3.3			ACS1	
Aug 27,13 10:48:42.486 AM	Aug 27,13 10:48:42.453 AM	✔		🔍	user2		Default Network Access	PAP_ASCII	R3	150.1.3.3	tty514		ACS1	
Aug 27,13 10:47:48.940 AM	Aug 27,13 10:47:48.940 AM	✔		🔍	$enab15$	150.1.3.3	Default Network Access	PAP_ASCII	R3	150.1.3.3			ACS1	
Aug 27,13 10:47:41.506 AM	Aug 27,13 10:47:41.500 AM	✔		🔍	$enab7$	150.1.3.3	Default Network Access	PAP_ASCII	R3	150.1.3.3			ACS1	
Aug 27,13 10:47:18.786 AM	Aug 27,13 10:47:18.746 AM	✔		🔍	user1		Default Network Access	PAP_ASCII	R3	150.1.3.3	tty514		ACS1	

Verify the RADIUS accounting logs, for exec accounting from R1:



AAA Protocol > RADIUS Accounting

Date : August 27, 2013

Generated on August 27, 2013 11:04:18 AM UTC

Reload

🔍=Click for details

ACS View Timestamp	ACS Timestamp	Details	Account Status Type	User Name	Calling Station ID	Endpoint IP Address	Account Authentic	Account Terminate Cause	Terminate Action	Account Session ID	Audit Session ID
Aug 27, 13 11:04:14.840 AM	Aug 27, 13 11:04:14.830 AM	🔍	Stop	user2			RADIUS	User Request		00000003	
Aug 27, 13 11:03:17.060 AM	Aug 27, 13 11:03:17.053 AM	🔍	Start	user2			RADIUS			00000003	
Aug 27, 13 11:03:13.313 AM	Aug 27, 13 11:03:13.293 AM	🔍	Stop	user1			RADIUS	User Request		00000002	
Aug 27, 13 11:02:51.743 AM	Aug 27, 13 11:02:51.723 AM	🔍	Start	user1			RADIUS			00000002	
Aug 27, 13 11:02:42.103 AM	Aug 27, 13 11:02:42.093 AM	🔍	Stop	user1			RADIUS	User Request		00000001	

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS LDAP Authentication

Task

- Configure R3 to integrate with the LDAP server at 172.16.10.100, using user **ldapuser** with a password of **1qaz@WSX**, which resides in the **Users** organizational unit.
- Ensure that authentication is performed based on the **sAMAccountName** LDAP attribute.
- The following users should already be configured on the LDAP server, available for testing:

Username	Password	Organizational Unit	Department	Company
USER1	1qaz@WSX	FINANCIAL	FINANCIAL	FINANCIAL
USER2	1qaz@WSX	TECHNICAL	TECHNICAL	TECHNICAL
USER3	1qaz@WSX	CCIELAB	CCIELAB	CCIELAB

A Note On Task Initial Configuration Files: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by

https://t.me/learningnets

the ASA.

Overview

The LDAP configuration is straightforward, very similar to the ASA LDAP configuration; the main commands differ, but configured attributes are the same. By default, IOS uses the full CN of the user defined in the LDAP server for authentication, which is different than how ASA behaves. For example, if you use LDAP for VPN user authentication and want to authenticate with the username **test1** , IOS queries the LDAP server for the user that has a CN value of **test1**. To use the **sAMAccountName**, an LDAP attribute map must be configured and applied under the LDAP server configuration.

Because of a Cisco bug, specifically **`CSCug65194`** , for IOS versions used in the CCIE Security lab, LDAP is not supported for administrative access authentication; it can only be used for VPN user authentication. On a **`debug ldap event`** and **`debug ldap error`** , within the output the following message is relevant and identifies the problem: **`LDAP doesn't suport interactive login .`**

Configuration

R3:

```
aaa new-model
!
aaa group server ldap SRV
 server TEST-SERVER
!
ldap attribute-map LDAP
 map type sAMAccountName username
!
ldap server TEST-SERVER
 ipv4 172.16.10.100
 attribute map LDAP
 base-dn DC=inelab,DC=local
 bind authenticate root-dn CN=ldapuser,CN=Users,DC=inelab,DC=local password 1qaz@WSX
```

Verification

Test successful user authentication:

https://t.me/learningnets

```
R3#test aaa group SRV USER1 1qaz@WSX new-code
User successfully authenticated


USER ATTRIBUTES

username             "USER1"
supplicant-group     "FINANCIAL"
username             "USER1"
!R3#test aaa group SRV USER2 1qaz@WSX new-code
User successfully authenticated

USER ATTRIBUTES

username             "USER2"
supplicant-group     "TECHNICAL"
username             "USER2"
!R3#test aaa group SRV USER3 1qaz@WSX new-code
User successfully authenticated


USER ATTRIBUTES

username             "USER3"
supplicant-group     "CCIELAB"
username             "USER3"
```

Verify LDAP server configuration:

```
R3#show ldap server TEST-SERVER
Server Information for TEST-SERVER
==================================
Server name             :TEST-SERVERServer IP                :172.16.10.100
Server listening Port   :389
Connection status       :UPBind Root-dn             :CN=ldapuser,CN=Users,DC=inelab,DC=local
Root Bind status        :Root-dn Bind Done
Server mode             :Non-Secure
Cipher Suite            :0x00
Authentication Seq      :Search first. Then Bind/Compare password next
Authentication Procedure:Bind with user passwordBase-Dn                  :DC=inelab,DC=local
Attribute map           :LDAP

Request timeout             :30
---------------------------------
* LDAP STATISTICS *
```

https://t.me/learningnets

```
Total messages  [Sent:5, Received:7]
Response delay(ms) [Average:504, Maximum:534]
Total search    [Request:2, ResultEntry:2, ResultDone:2]
Total bind      [Request:3, Response:3]
Total extended  [Request:0, Response:0]
Total compare   [Request:0, Response:0]
---------------------------------
```

Verify the default LDAP attribute mapping, as well as the custom configured one:

```
R3#show ldap attributes
 LDAP Attribute                    Format      AAA Attribute
 ===============                   ======      =============
 airespaceBwDataBurstContract      Ulong       bsn-data-bandwidth-burst-contr
 userPassword                      String      password
 airespaceBwRealBurstContract      Ulong       bsn-realtime-bandwidth-burst-c
 employeeType                      String      employee-type
 airespaceServiceType              Ulong       service-type
 airespaceACLName                  String      bsn-acl-name
 priv-lvl                          Ulong       priv-lvl
 memberOf                          String DN   supplicant-group
 cn                                String      username

 airespaceDSCP                     Ulong       bsn-dscp
 policyTag                         String      tag-name
 airespaceQOSLevel                 Ulong       bsn-qos-level
 airespace8021PType                Ulong       bsn-8021p-type
 airespaceBwRealAveContract        Ulong       bsn-realtime-bandwidth-average
 airespaceVlanInterfaceName        String      bsn-vlan-interface-name
 airespaceVapId                    Ulong       bsn-wlan-id
 airespaceBwDataAveContract        Ulong       bsn-data-bandwidth-average-con
 sAMAccountName                    String      sam-account-name
 meetingContactInfo                String      contact-info
 telephoneNumber                   String      telephone-number
 Map: LDAP
 sAMAccountName                    String      username
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS TACACS Command Authorization and Accounting

Task

- Configure R1 for AAA as follows:
  - All communication with the ACS server should be based on Loopback0 interface through a multiplexed connection.
  - User **admin** should be placed into privilege level 11 upon authentication, but allowed to access privilege level 15; if ACS server becomes unreachable, the local user database should be used with the same policies.
  - Authorize only privilege level 15 commands, for both console and VTY lines; ensure that unauthenticated users are authorized for all commands.
  - Allow user **admin** to configure IP address from the range of 150.1.0.0/16 to the Loopback1 interface.
  - Account for all commands.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Remote command authorization is supported only through TACACS+. Every time a user presses Enter after typing a command, IOS sends the fully expanded command line to the TACACS+ server in the context of the current user's authorization session. The server compares the string with the policy configuration

https://t.me/learningnets

(actually, a list of regular expressions) for the particular user's authorization policy, and responds whether the command is permitted or denied.

To enable per-command authorization at a particular exec privilege level, use the command **`aaa authorization commands <Level> [<NAME>|default] group [tacacs+|<GROUP_NAME>`** . It makes no sense to use the local database for per-command authorization because local configurations are always in effect. Thus, the only two meaningful options are tacacs+ and custom TACACS+ server group. However, you may specify local authorization as the fallback method in addition to the primary TACACS+ authorization, so that if the connection with the TACACS server fails, authorization is done locally based on the user's privilege level.

Command authorization is enabled per privilege-level, and by default applies only to the exec mode commands, not the configuration commands. To enable configuration mode commands authorization, enter the command **`aaa authorization config-commands`** . This command instructs the router to send the command strings entered in the configuration mode to the AAA server for authorization as well. Notice that the router sends both the exec and configuration mode commands in the same format, without any discriminator to distinguish those.

The console line will not be affected by authorization settings unless you enter the command **`aaa authorization console`** ; the default setting acts as a fail-safe mechanism to prevent you from locking yourself out of the console. In addition, if you are using the default authorization list, you may override the setting on a per-line basis using the line-level command **`authorization commands <Level> <LIST>`** after configuring the method using the global command **`aaa authorization commands <Level> <LIST>`** . In this case, the level numbers must match for the new authorization list to take effect.

Configuration

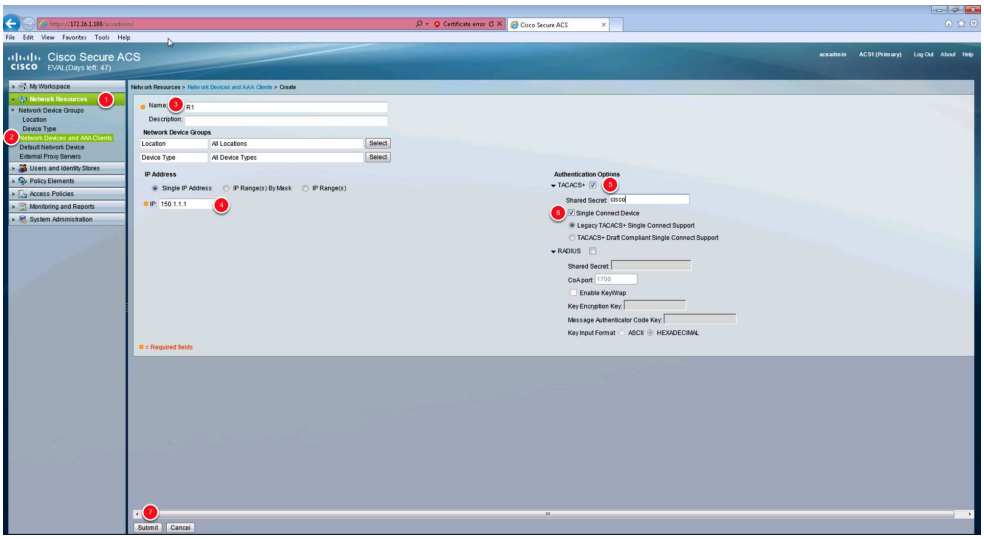
R1:

```
aaa new-model
!
tacacs server TACACS
 address ipv4 172.16.1.100
 key cisco
 single-connection
!
aaa group server tacacs+ ACS
 server name TACACS
 ip tacacs source-interface Loopback0
!
aaa authentication login default group ACS local
```

https://t.me/learningnets

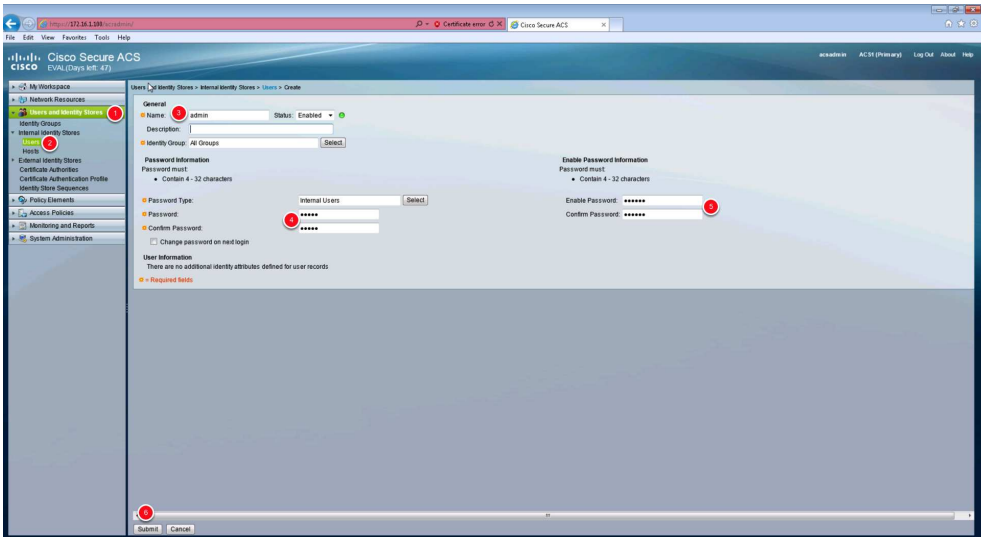
```
aaa authentication enable default group ACS enable
aaa authorization commands 15 default group ACS local none
!
aaa authorization config-commands
aaa authorization console
aaa authorization exec default group ACS local none
!
aaa accounting commands 0 default start-stop group ACS
aaa accounting commands 1 default start-stop group ACS
aaa accounting commands 15 default start-stop group ACS
!
username admin privilege 11 password admin
enable secret cisco
```

On ACS, add R1 as a NAD into ACS, configure it for TACACS communication, and enable the connection multiplexing functionality:

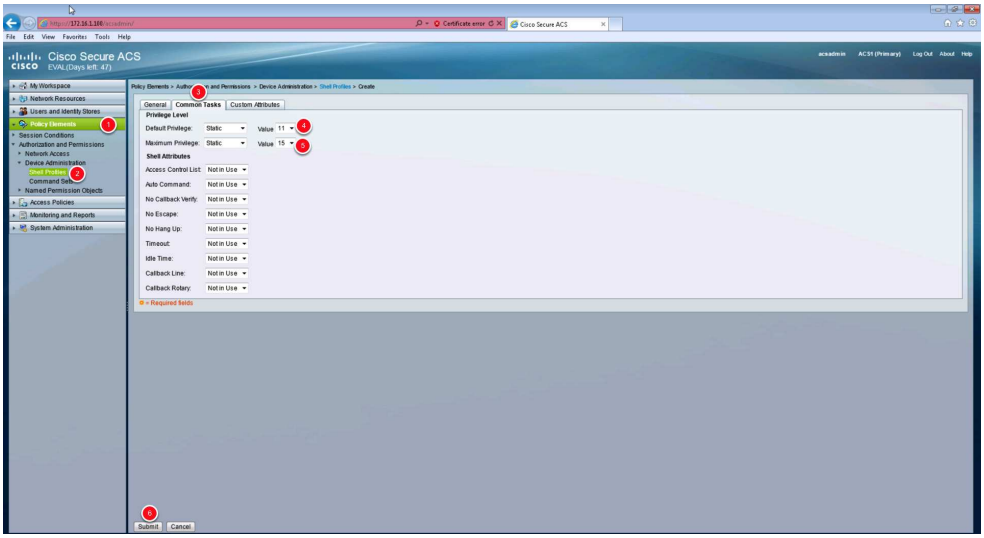


Add user **admin** into the Internal User Store:

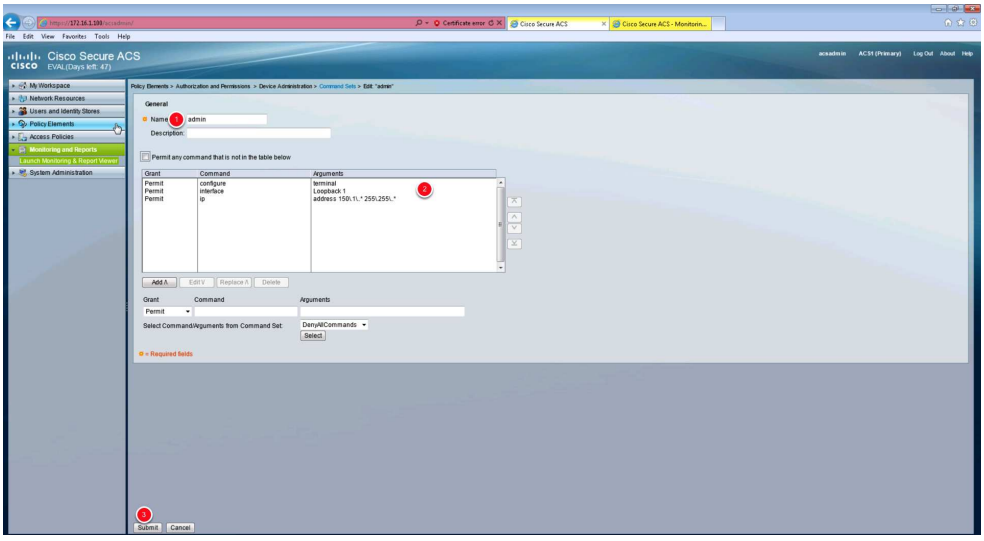
https://t.me/learningnets



Configure a shell profile, so that the user is assigned a privilege level of 11 upon login, and is allowed to access privilege level 15 upon enable authentication:



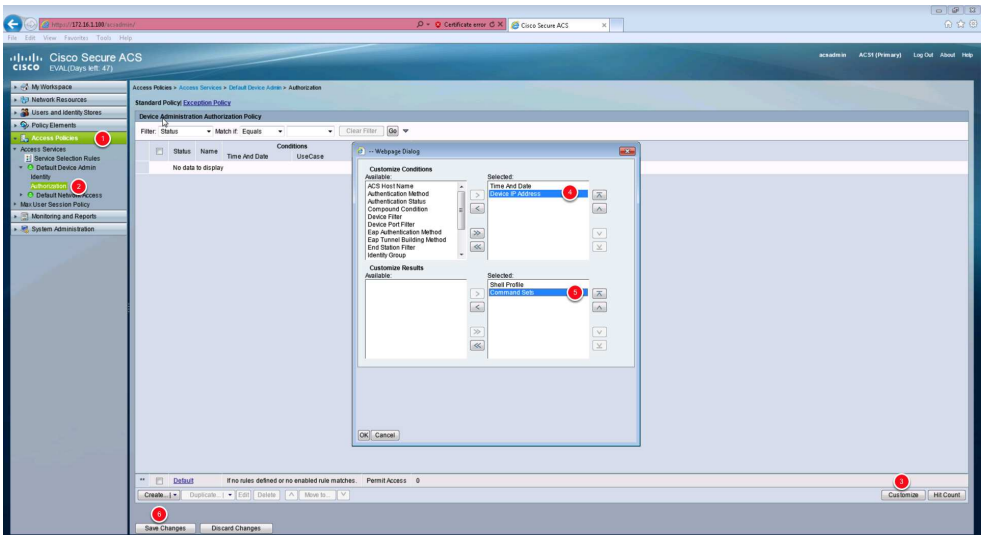
Configure the command set:



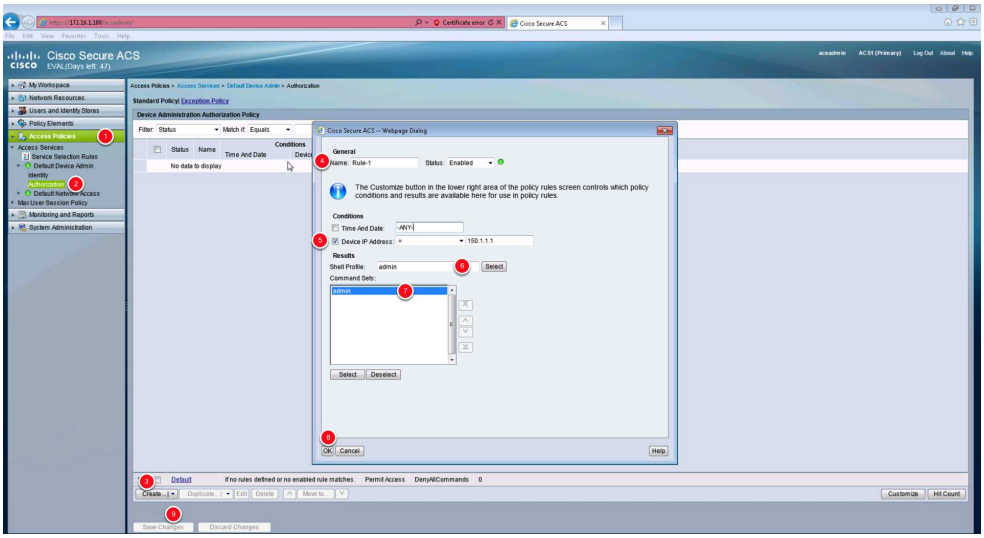
Add the device IP address as a condition for the authorization policy, and command

https://t.me/learningnets

set as a result:



Configure the TACACS authorization access policy:



To regain full router control, you can allow all commands after verification.

Verification

Telnet from R3 and verify assigned privilege level, as well as allowed commands (the same test can be performed from the console):

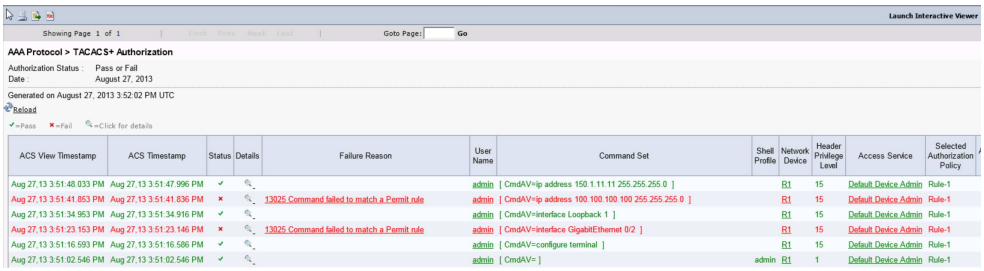
```
R3#telnet 150.1.1.1
Trying 150.1.1.1 ... Open username: admin
password:
R1#show privilege
Current privilege level is 11
R1#enable
```

https://t.me/learningnets

```
password: R1#show privilege
Current privilege level is 15


!R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#interface GigabitEThernet0/2
Command authorization failed.
R1(config)#interface Loopback1
R1(config-if)#ip address 100.100.100.100 255.255.255.0
Command authorization failed.
R1(config-if)#ip address 150.1.11.11 255.255.255.0
```

Verify the TACACS Authorization Logs:



Launch Interactive Viewer

Showing Page 1 of 1 | First Prev Next Last | Goto Page: Go

AAA Protocol > TACACS+ Authorization

Authorization Status : Pass or Fail
Date : August 27, 2013

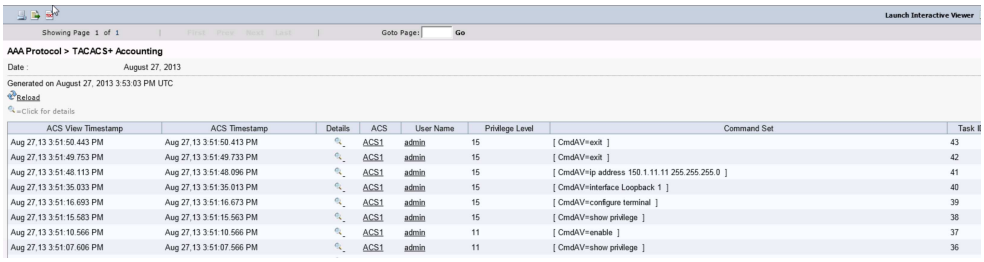
Generated on August 27, 2013 3:52:02 PM UTC

Reload

✔=Pass ✖=Fail 🔍=Click for details

ACS View Timestamp	ACS Timestamp	Status	Details	Failure Reason	User Name	Command Set	Shell Profile	Network Device	Header Privilege Level	Access Service	Selected Authorization Policy
Aug 27,13 3:51:48.033 PM	Aug 27,13 3:51:47.996 PM	✔	🔍		admin	[CmdAV=ip address 150.1.11.11 255.255.255.0]		R1	15	Default Device Admin	Rule-1
Aug 27,13 3:51:41.853 PM	Aug 27,13 3:51:41.836 PM	✖	🔍	13025 Command failed to match a Permit rule	admin	[CmdAV=ip address 100.100.100.100 255.255.255.0]		R1	15	Default Device Admin	Rule-1
Aug 27,13 3:51:34.953 PM	Aug 27,13 3:51:34.916 PM	✔	🔍		admin	[CmdAV=interface Loopback 1]		R1	15	Default Device Admin	Rule-1
Aug 27,13 3:51:23.153 PM	Aug 27,13 3:51:23.146 PM	✖	🔍	13025 Command failed to match a Permit rule	admin	[CmdAV=interface GigabitEthernet 0/2]		R1	15	Default Device Admin	Rule-1
Aug 27,13 3:51:16.593 PM	Aug 27,13 3:51:16.586 PM	✔	🔍		admin	[CmdAV=configure terminal]		R1	15	Default Device Admin	Rule-1
Aug 27,13 3:51:02.546 PM	Aug 27,13 3:51:02.546 PM	✔	🔍		admin	[CmdAV=]	admin	R1	1	Default Device Admin	Rule-1

Verify the TACACS Accounting logs:



Launch Interactive Viewer

Showing Page 1 of 1 | First Prev Next Last | Goto Page: Go

AAA Protocol > TACACS+ Accounting

Date : August 27, 2013

Generated on August 27, 2013 3:53:03 PM UTC

Reload

🔍=Click for details

ACS View Timestamp	ACS Timestamp	Details	ACS	User Name	Privilege Level	Command Set	Task ID
Aug 27,13 3:51:50.443 PM	Aug 27,13 3:51:50.413 PM	🔍	ACS1	admin	15	[CmdAV=exit]	43
Aug 27,13 3:51:49.753 PM	Aug 27,13 3:51:49.733 PM	🔍	ACS1	admin	15	[CmdAV=exit]	42
Aug 27,13 3:51:48.113 PM	Aug 27,13 3:51:48.096 PM	🔍	ACS1	admin	15	[CmdAV=ip address 150.1.11.11 255.255.255.0]	41
Aug 27,13 3:51:35.033 PM	Aug 27,13 3:51:35.013 PM	🔍	ACS1	admin	15	[CmdAV=interface Loopback 1]	40
Aug 27,13 3:51:16.693 PM	Aug 27,13 3:51:16.673 PM	🔍	ACS1	admin	15	[CmdAV=configure terminal]	39
Aug 27,13 3:51:15.583 PM	Aug 27,13 3:51:15.563 PM	🔍	ACS1	admin	15	[CmdAV=show privilege]	38
Aug 27,13 3:51:10.566 PM	Aug 27,13 3:51:10.566 PM	🔍	ACS1	admin	11	[CmdAV=enable]	37
Aug 27,13 3:51:07.606 PM	Aug 27,13 3:51:07.566 PM	🔍	ACS1	admin	11	[CmdAV=show privilege]	36

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS Authentication Proxy Local AAA

Task

- Configure authentication proxy for telnet traffic on R1, using the local user database:
  - Configure user **AUTH-USER** for this purpose.
  - Authorize the user so that it is not allowed any ICMP traffic.
  - The authentication process should be triggered only for traffic destined to R3's VLAN 37 IP address and received inbound on R1's VLAN 26 interface.
  - Force the user to authenticate again after two hours.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IOS Authentication Proxy is actually similar to the ASA Cut-Through Proxy functionality. The idea behind Authentication Proxy is to authenticate users before passing traffic through the router, download per-user access profile (ACL rules), and merge it with the interface-applied access-list so that each user receives differentiated access.

To authenticate the user, IOS can transparently intercept a telnet/FTP/HTTP session, these being protocols that support inline authentication, just like the ASA. For other protocols, ASA has the virtual telnet or virtual HTTP service available, but IOS does not. If, for example, a user needs to authenticate with an IOS box before

https://t.me/learningnets

issuing ICMP traffic, it first must initiate a telnet/FTP/HTTP session through the IOS box so that the authentication prompt is presented to the user. Note that even though the `ip admission` command has a `virtual-ip` option, this is not used for virtual HTTP service, but to redirect users for authentication purposes.

Normally, a remote authentication server (RADIUS or TACACS) is required to authenticate users and push down the Cisco AV pair values with the access-list contents, but the local IOS AAA service can be used as well through the use of attribute lists. Using the local database for authentication proxy, when the user is authenticated it will be automatically authorized for any IP traffic. By using an AAA attribute list, we can modify the default authorization policy by configuring additional access-list entries, using the same attributes (cisco-av-pair) as received from a remote RADIUS or TACACS server. Note that there is no need to configure a AAA login authentication method, because once AAA is enabled, the default is to use the local database anyway.

Authentication proxy is configured as a rule applied ingress to the interface. Additionally, there should be an access-list applied ingress to the same interface for preventing users from initiating any sessions before being authenticated. The access-list may need to allow control-plane traffic if required, such as routing protocols or even RADIUS/TACACS packets.

To configure the authentication proxy function, the old `ip auth-proxy` command or the new `ip admission` command can be used. The new command allows two features:

- Configure integration between a Cisco router and Scansafe (a Cloud Web proxy), which is not in the scope of the current CCIE Security Exam Blueprint.
- For HTTP authentication proxy, you have two additional buttons in the web portal, allowing the user to consent use of the service.

Configuration

R1:

```
aaa new-model
aaa authorization auth-proxy default local
!
aaa attribute list AUTH-LIST
 attribute type proxyacl "deny icmp any any" service auth-proxy
!
username AUTH-USER password 0 AUTH-USER
username AUTH-USER aaa attribute list AUTH-LIST
!
```

https://t.me/learningnets

```
ip access-list extended DEFAULT-ACL
 deny ip any any
!
ip access-list extended AUTH-ACL
 permit tcp any host 136.1.37.3
!
ip auth-proxy name LOCAL telnet absolute-timer 120 list AUTH-ACL
!
interface GigabitEthernet0/1.26
 ip access-group DEFAULT-ACL in
 ip auth-proxy LOCAL
```

Verification

Verify authentication proxy configuration:

```
R1#show ip auth-proxy configuration

Authentication Proxy Banner not configured
Consent Banner is not configured
Customizable Authentication Proxy webpage not configured
Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Watch-list is disabled

Authentication Proxy Max HTTP process is 7
Authentication Proxy Auditing is disabled
Max Login attempts per user is 5

Authentication Proxy Rule Configuration
 Auth-proxy name LOCAL telnet list AUTH-ACL inactivity-time 60 minutes
absolute-timer 120 minutes
```

Verify that telnet to SW1 does not trigger the authentication process:

https://t.me/learningnets

Image not readable or empty
/uploads/workbooks/images/originals/35e461b181e273a72fca9c8c4145ee29af05e2e8.png

Verify the authenticated user and its authorization (ACL) process, also verify that ICMP traffic to R3 is not allo

Image not readable or empty
/uploads/workbooks/images/originals/2d11ca03aa4b72cdec63ecf12f21c0b7f5d44169.png

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS Authentication Proxy Using RADIUS

Task

- Configure Authentication Proxy settings on R1 as follows:
  - Use the RADIUS server at 172.16.1.100 with Loopback0 interface.
  - The authentication proxy should be active for traffic received inbound on R1's VLAN 26 interface, and should intercept all FTP and telnet traffic for unauthenticated users.
  - Authenticated users should be allowed to send ICMP packets and initiate TCP sessions.
  - Do not modify the default console settings.
- Configure the ACS server with the following users and authorizations:
  - User **user1** is authorized only for ICMP traffic.
  - User **user2** is authorized only for TCP traffic.
  - User **user3** is authorized for both ICMP and TCP traffic.
- If users fail authentication two times, they should not be allowed to authenticate again for 10 minutes.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Just like the cut-through functionality on ASA, RADIUS can assign an access-list for

https://t.me/learningnets

authentication proxy in IOS in three different ways:

- **Downloadable ACL**, which requires the ACL to be configured on the RADIUS server; upon successful authentication, it will be sent to the NAD (using RADIUS Cisco VSA).
- **Filter-ID ACL**, which requires the ACL to be configured on the NAD; upon successful authentication, only the name of the ACL will be sent to the NAD (using IETF RADIUS attribute 11).
- **Proxy-ACL**, also known as Cisco AV Pair, which requires the ACL to be configured on the RADIUS server in a specific format; upon successful authentication, it will be sent to the NAD (using RADIUS Cisco VSA).

For downloadable ACL to work, the following additional configurations must be done:

- The command `aaa authorization network default` is required so that when IOS receives the dACL name from the RADIUS server, it accepts it.
- The command `radius-server vsa send authentication` is required so that IOS can request the ACL entries for the received dACL from the RADIUS server.

When the Proxy-ACL method is used, the above two are not required, because in the **ACCESS-ACCEPT** message, RADIUS also sends the ACL entries, so there is no ACL name involved.

For Filter-ID ACL to work, the `radius-server attribute 11 default direction in` command is required to make the IOS apply the ACL in the inbound direction of the interface where authentication proxy is configured, because by default the direction is outbound. Another option would be to signal the direction (inbound or outbound) from the RADIUS server. For example, if the router had the ACL name of **FILTER-ACL** configured, for RADIUS to signal the inbound direction, the ACL would have to be specified in RADIUS as **FILTER-ACL.in**.

On the IOS version used, 15.2(3)T3, by default the RADIUS **Service-Type** attribute is not sent in the **ACCESS-REQUEST** message toward the RADIUS server. This may differ based on the IOS code used. When authentication proxy is triggered through telnet, the attribute value is **Login**, whereas if it's triggered through HTTP, the attribute value is **Outbound**. To configure the IOS to send this attribute, the command `radius-server attribute 6 on-for-login-auth` is required.

Configuration

R1:

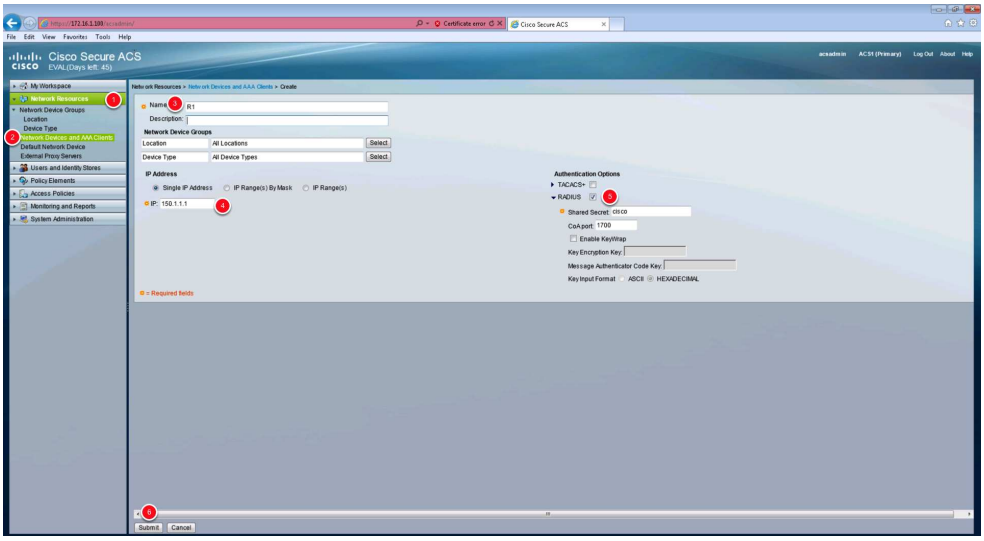
```
aaa new-model
```

https://t.me/learningnets

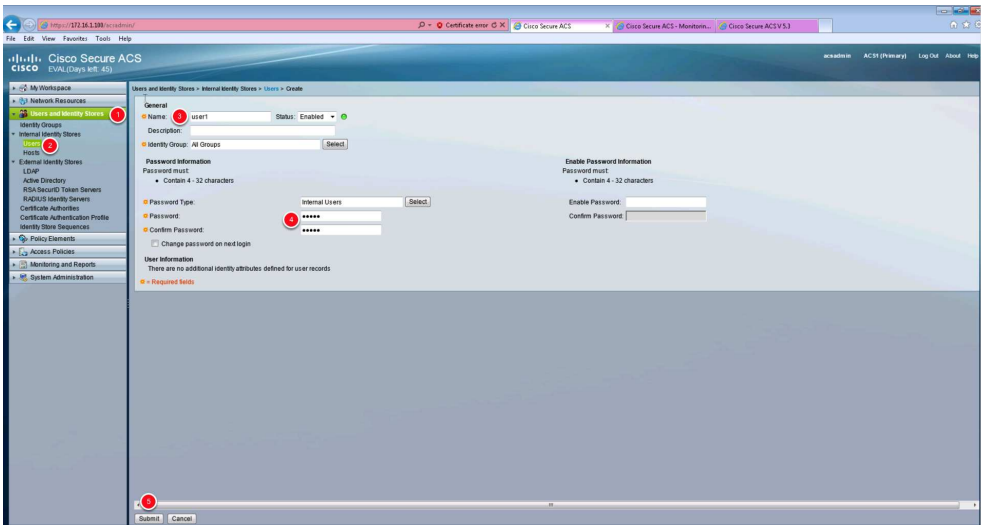
```
aaa authentication login CONSOLE none
aaa authentication login default group ACS
aaa authorization auth-proxy default group ACS
aaa authorization network default group ACS
!
line console 0
 login authentication CONSOLE
!
radius server RADIUS
 address ipv4 172.16.1.100
 key cisco
!
aaa group server radius ACS
 server name RADIUS
 ip radius source-interface Loopback0
!
radius-server vsa send authentication
radius-server attribute 11 default direction in
!
ip auth-proxy name RADIUS ftp
ip auth-proxy name RADIUS telnet
ip auth-proxy max-login-attempts 2
ip auth-proxy init-state-time 10
!
ip access-list extended DEFAULT-ACL
 deny ip any any
!
ip access-list extended user2
 permit tcp any any
!
interface GigabitEthernet0/1.26
 ip access-group DEFAULT-ACL in
 ip auth-proxy RADIUS
```

On ACS, configure R1 as a NAD for RADIUS communication:

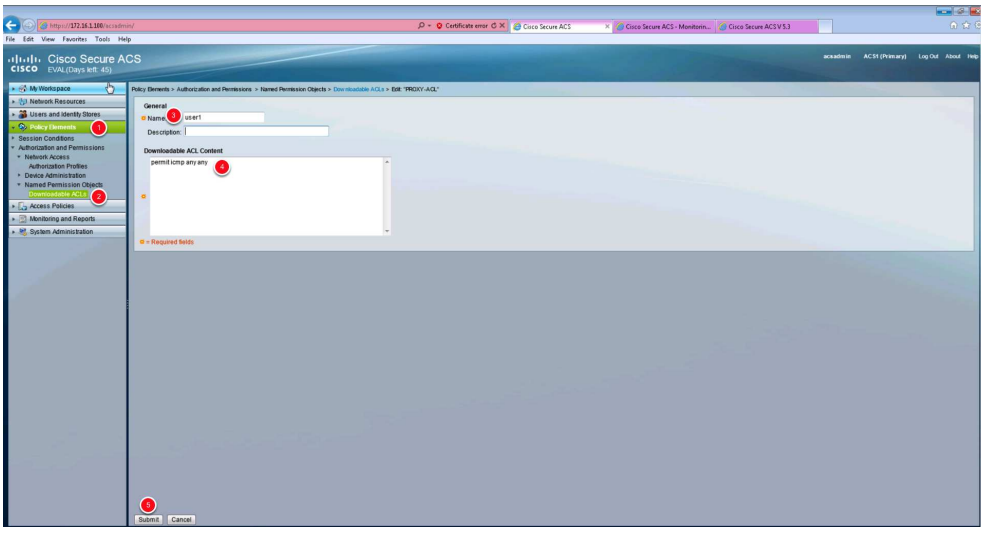
https://t.me/learningnets



Configure all three users in the ACS Internal Identity Store, for example **user1**:

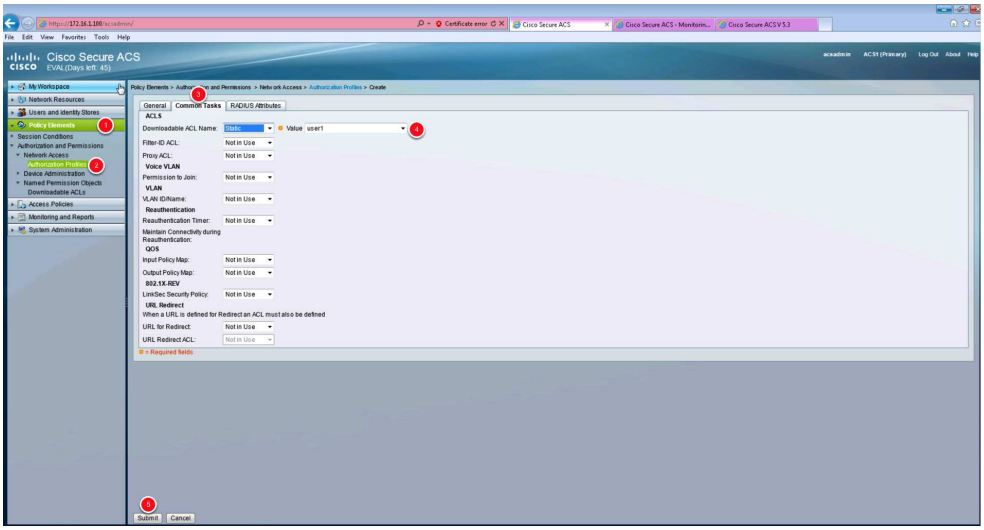


Configure the dACL for **user1** with appropriate authorization:

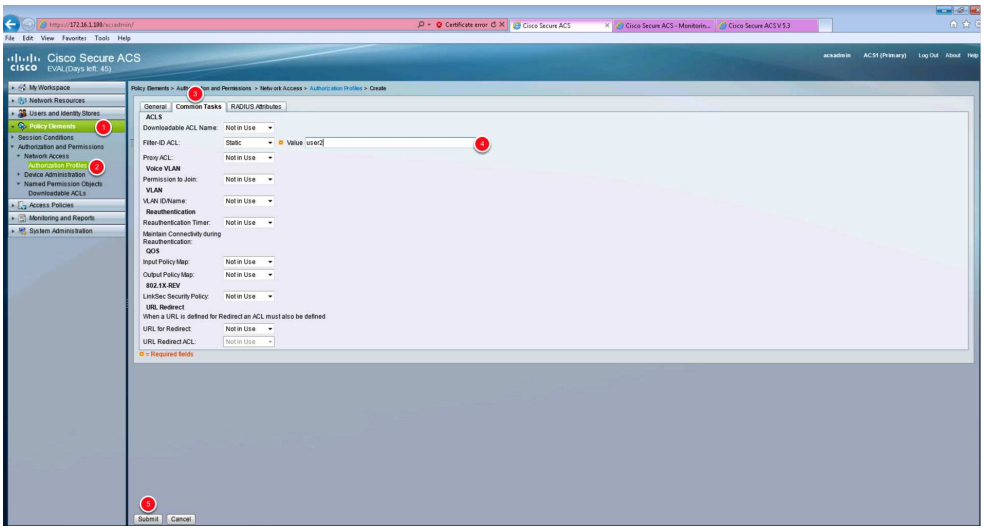


https://t.me/learningnets

Configure three authorization profiles, one for each user, each using one of the available methods described, example for **user1**:

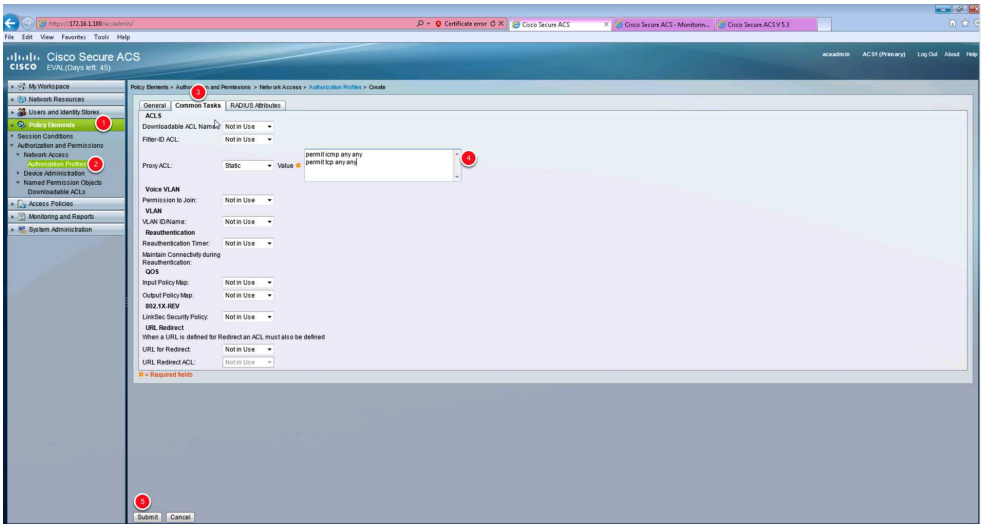


Authorization profile for **user2**:

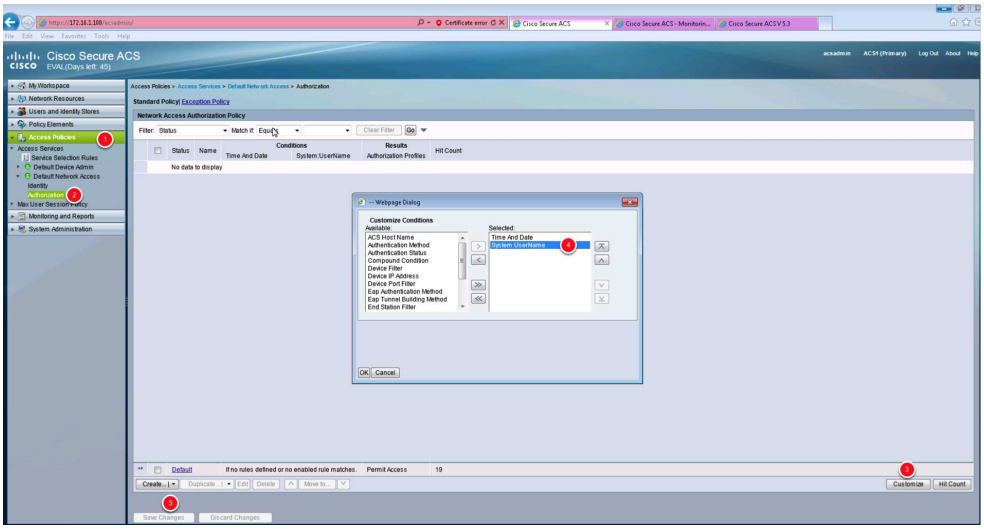


Authorization profile for **user3**:

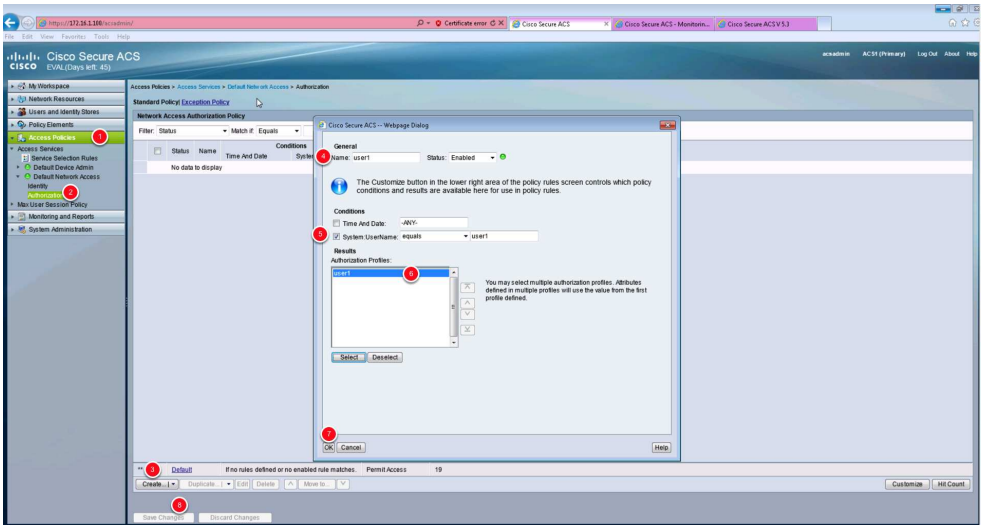
https://t.me/learningnets



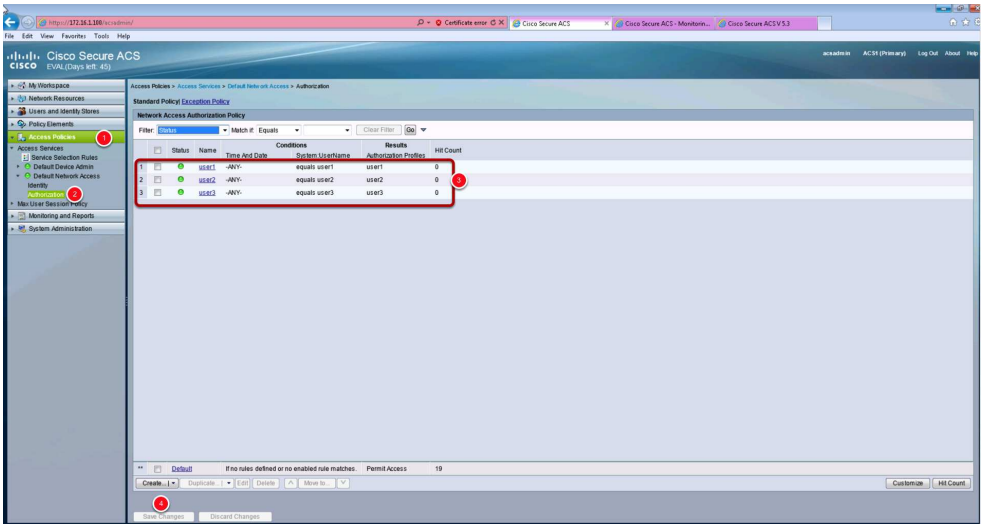
Add the username as condition for RADIUS authorization access policies:



Configure three authorization access policies, one for each user:



https://t.me/learningnets



Verification

Note that authentication proxy works only with the **default** AAA authentication login list, and not with a named list. We've used all three options of RADIUS authorization, on a per-user level.

Authenticate using **user1** and verify the authorization:

https://t.me/learningnets

Image not readable or empty
/uploads/workbooks/images/originals/ddf6f7a2ed5b8ca923a84903005085780c7e6d23.png

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

IOS Authentication Proxy Using TACACS

Task

- Configure Authentication Proxy settings on R1 to achieve the following:
  - Use the TACACS server at 172.16.1.100 with Loopback0 interface.
  - The authentication proxy should be active for traffic received inbound on R1's VLAN 26 interface, and should intercept HTTP traffic for unauthenticated users.
  - Authenticated users should be allowed to send ICMP packets and initiate TCP sessions.
  - Do not modify the default console settings.
- Users should see the following banner before authenticating: **This is iNE Authetication Portal**.
- Ensure that users are forced to consent when authenticating.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Configuration

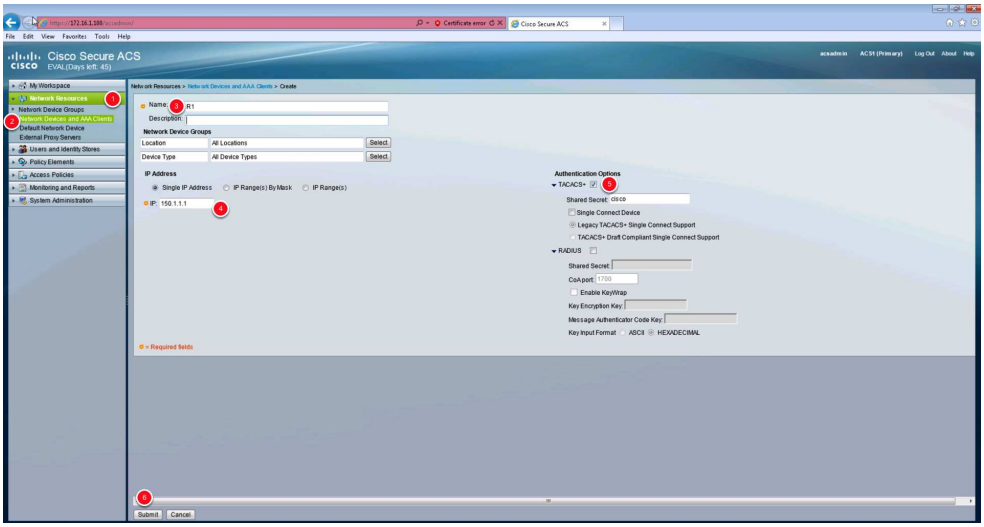
R1:

```
aaa new-model
aaa authentication login CONSOLE none
aaa authentication login default group ACS
```

https://t.me/learningnets

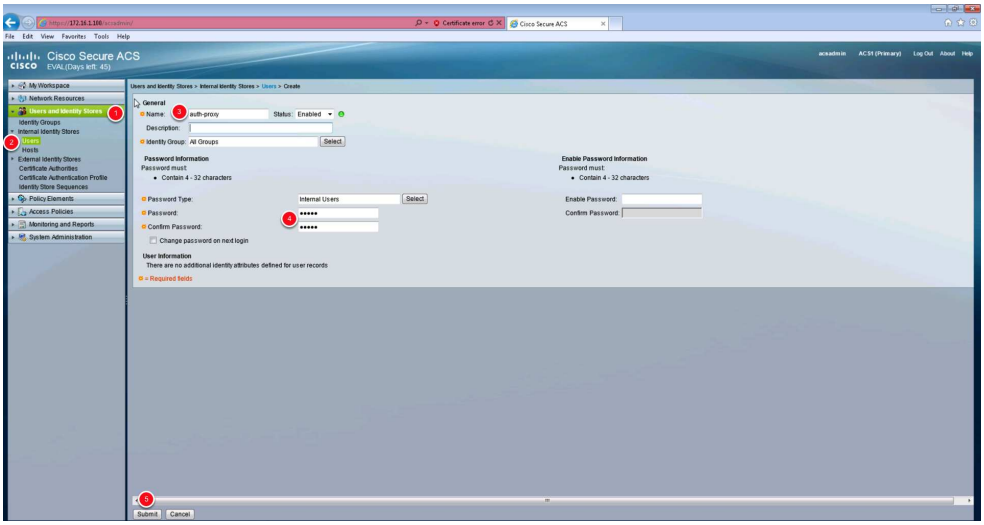
```
aaa authorization auth-proxy default group ACS
!
tacacs server TACACS
 address ipv4 172.16.1.100
 key cisco
!
aaa group server tacacs+ ACS
 server name TACACS
 ip tacacs source-interface Loopback0
!
ip admission name TACACS proxy http
ip admission name TACACS consent
ip admission auth-proxy-banner http q This is iNE Authetication Portal q
ip http server
!
ip access-list extended DEFAULT-ACL
 deny ip any any
!
interface GigabitEthernet0/1.26
 ip access-group DEFAULT-ACL in
 ip admission TACACS
```

Configure R1 as a NAD device in ACS; use TACACS:

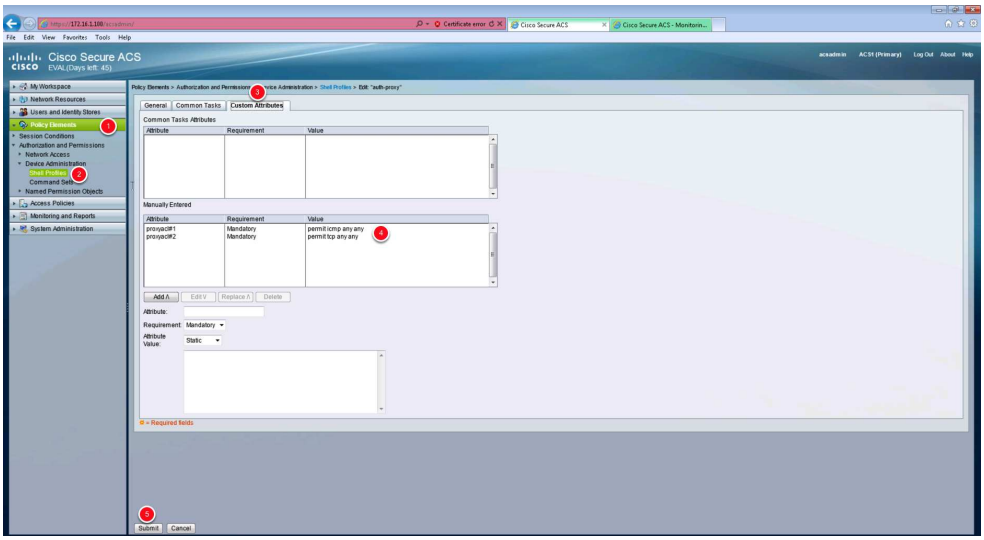


Configure a user for authentication:

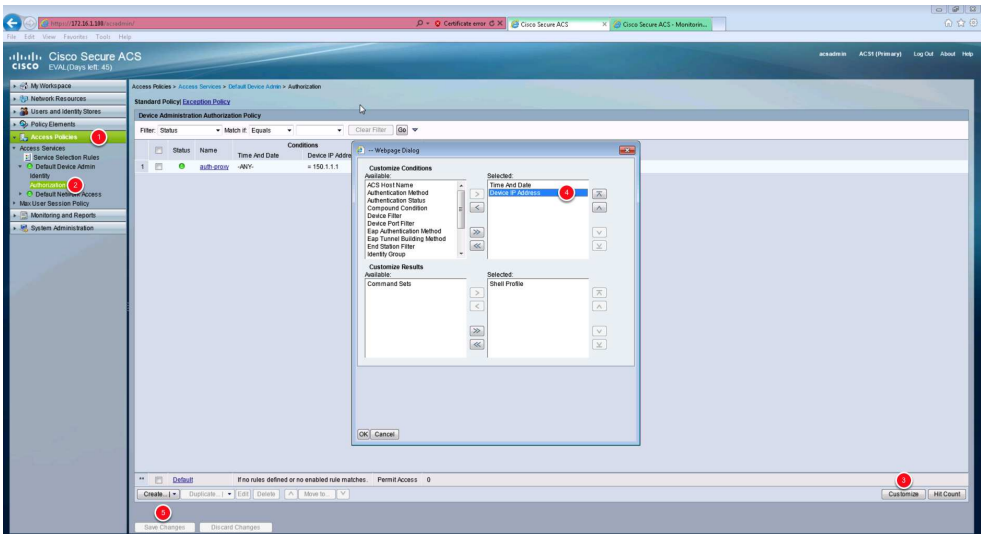
https://t.me/learningnets



Configure a shell profile for the proxy-acl to be downloaded:

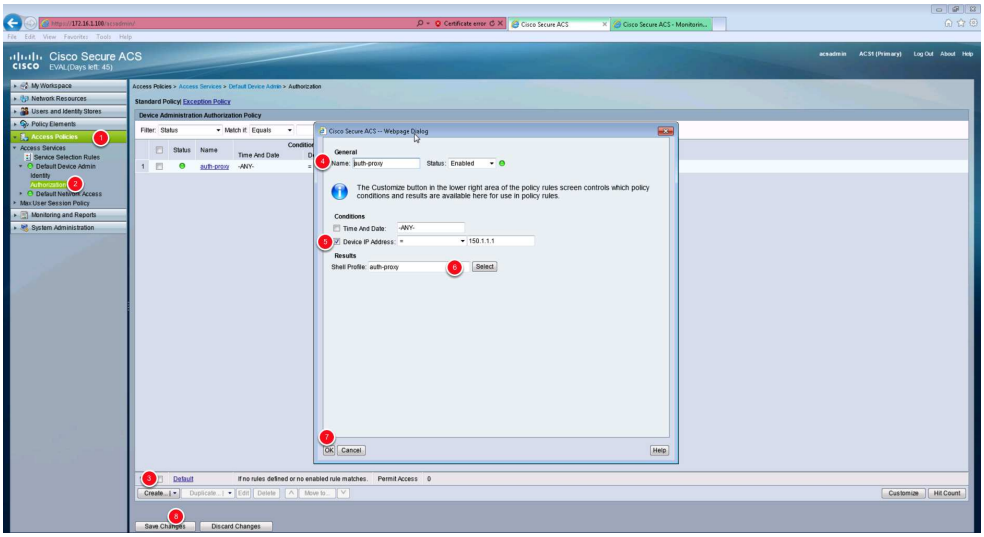


Add a condition to be used in the TACACS authorization access policy:



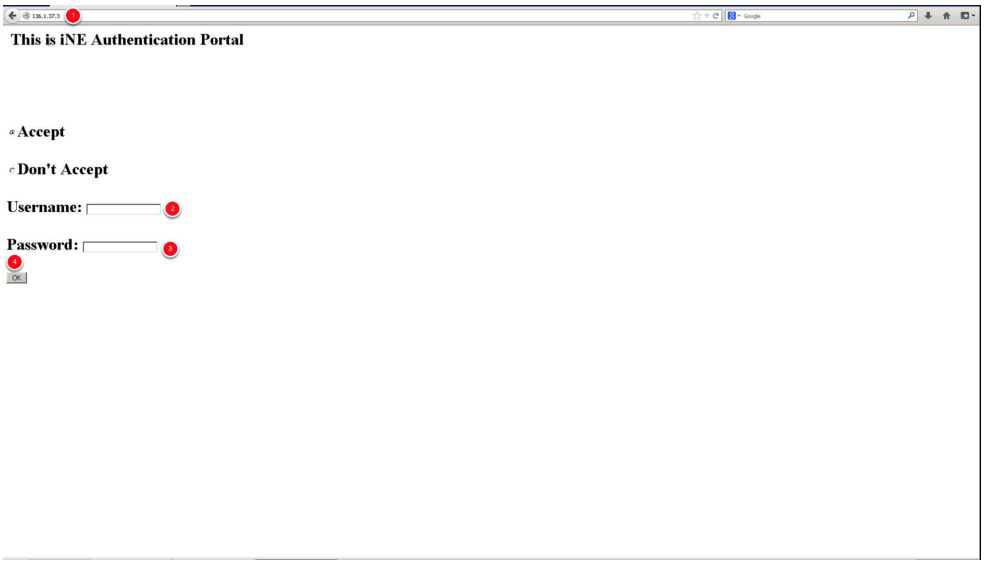
https://t.me/learningnets

Configure the TACACS authorization access policy and apply the configured shell profile:



Verification

Test the authentication proxy from the TEST-SRV-A and verify the consent option, as well as the banner in the authentication portal:



Verify that correct authorization is received from ACS server:

```
R1#show ip admission cache username auth-proxy
Authentication Proxy Cache

Authentication Method   : ProxyUser Name              : auth-proxy
Client IP               : 172.16.10.100
Client Port             : 52617
Timeout                 : 60
Time  Remaining         : 60              Connection state        : ESTAB
```

https://t.me/learningnets

```
EPM information : Authproxy
    Admission feature:  AUTHPROXY
         AAA Policies:   Proxy ACL:  permit icmp any any
Proxy ACL:  permit tcp any any

EOU information
--------------------------------------------------------------------------
Address         Interface              AuthType   Posture-Token Age(min)
--------------------------------------------------------------------------
EPM information : EOU

Existing Firewall Sessions Information:
!
!R1#show ip access-lists DEFAULT-ACL
Extended IP access list DEFAULT-ACL permit icmp host 172.16.10.100 any
permit tcp host 172.16.10.100 any (5 matches)

     10 deny ip any any (1622 matches)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

Installing ISE Certificates

Task

- Enroll ISE-1 with the CA server from TEST-SERVER-A, and use the certificate for both management and EAP.
- ISE-1 should use a CN of **ISE1.inelab.local**, Organization of **INE**, Organization Unit of **CCIESECURITY**, and RSA Key Pair of 2048 bit.
- Configure the CA server so that certificates are manually granted.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

By default, ISE uses self-signed certificates for HTTPS management, just like ACS. ISE has the same options as ACS; a CA-issued certificate can be used for HTTPS management, EAP authentications, or both. The basic steps of enrolling ISE with a CA are:

1. Generate a Certificate Signing Request (CSR).
2. Export the CSR and save it locally into a WordPad file.
3. Browse the CA to request a certificate.
4. Request a certificate using the advanced certificate request option.
5. Manually grant the certificate if required.

https://t.me/learningnets

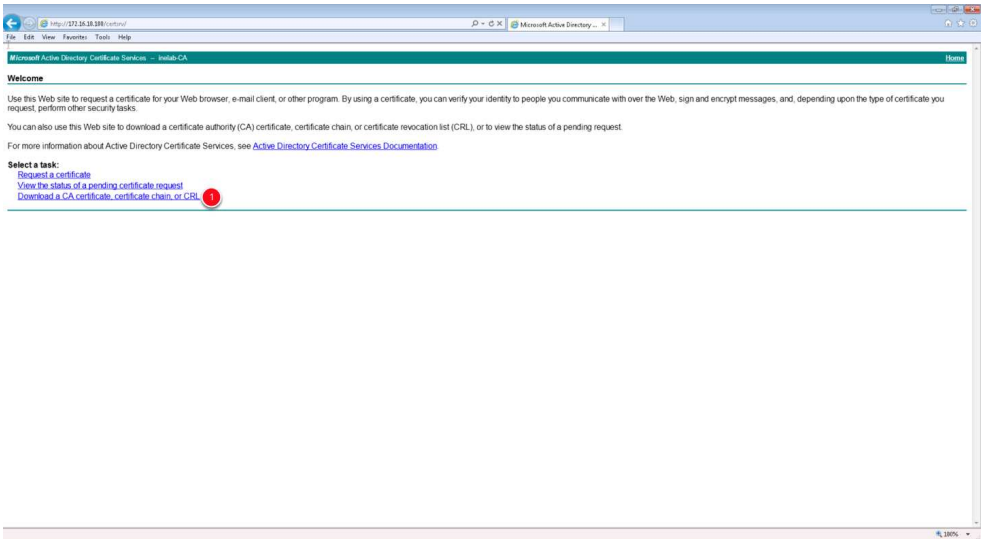
6. Download the certificate.
7. Import the certificate in ISE, binding it to the previous CSR.

To manually grant the certificate, CA settings must be modified before issuing the certificate request.

> Note that even after installing a certificate issued by a trusted CA, you may still see browser warnings if the FQDN or the IP address you are using in the browser for managing ACS does not match the CN or SAN values from the certificate; this is another check the browser does to consider the certificate valid.

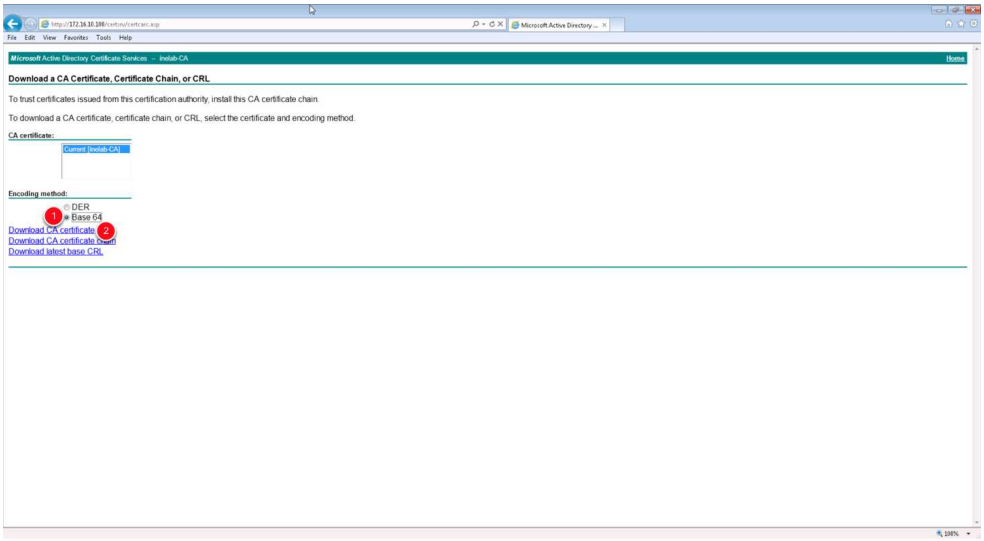
Configuration

Browse the CA server web enrollment page to download the CA certificate:

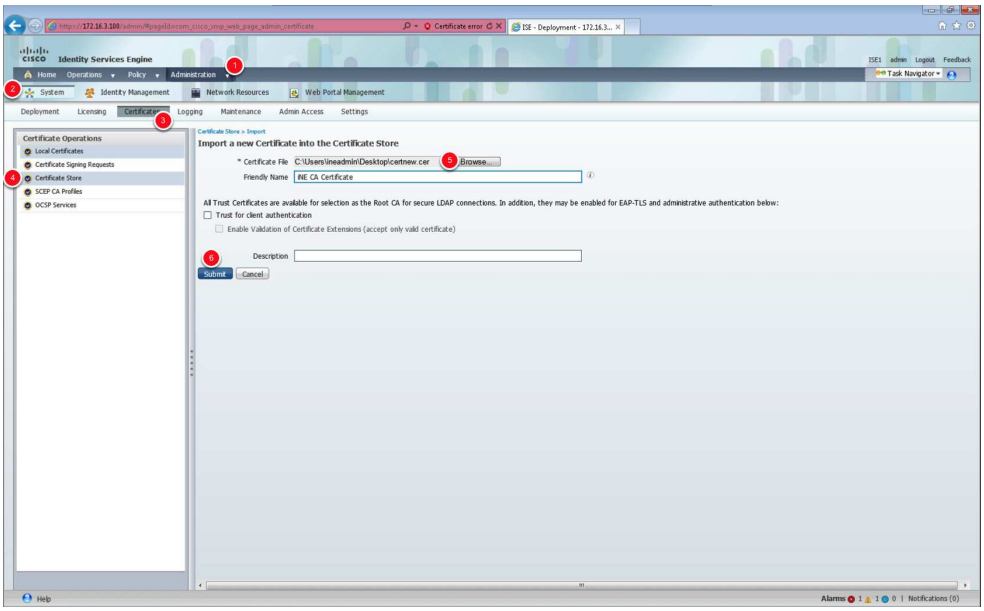


Select the **Base 64** format, download the CA certificate, and save the file locally on the PC:

https://t.me/learningnets

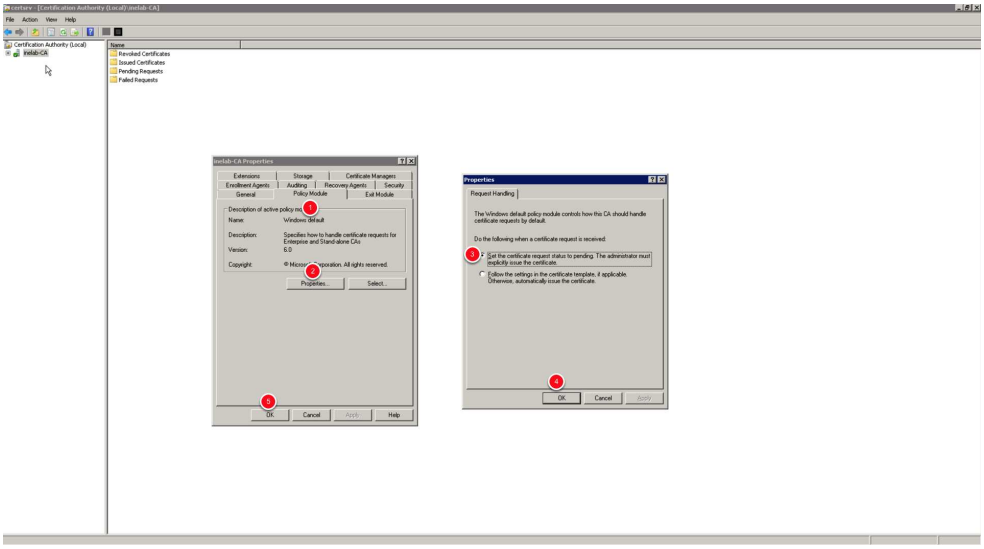


Browse to the **Certificate Store** section of the ISE and import the CA certificate; to trust client certificates issued by this CA (for example, for EAP-TLS handshake), select the **Trust for client authentication** check box:



Open the CA management console using the command **certsrv.msc** from Run, right-click, and select **Properties**. From here, configure the CA so that certificate are manually granted after request:

https://t.me/learningnets

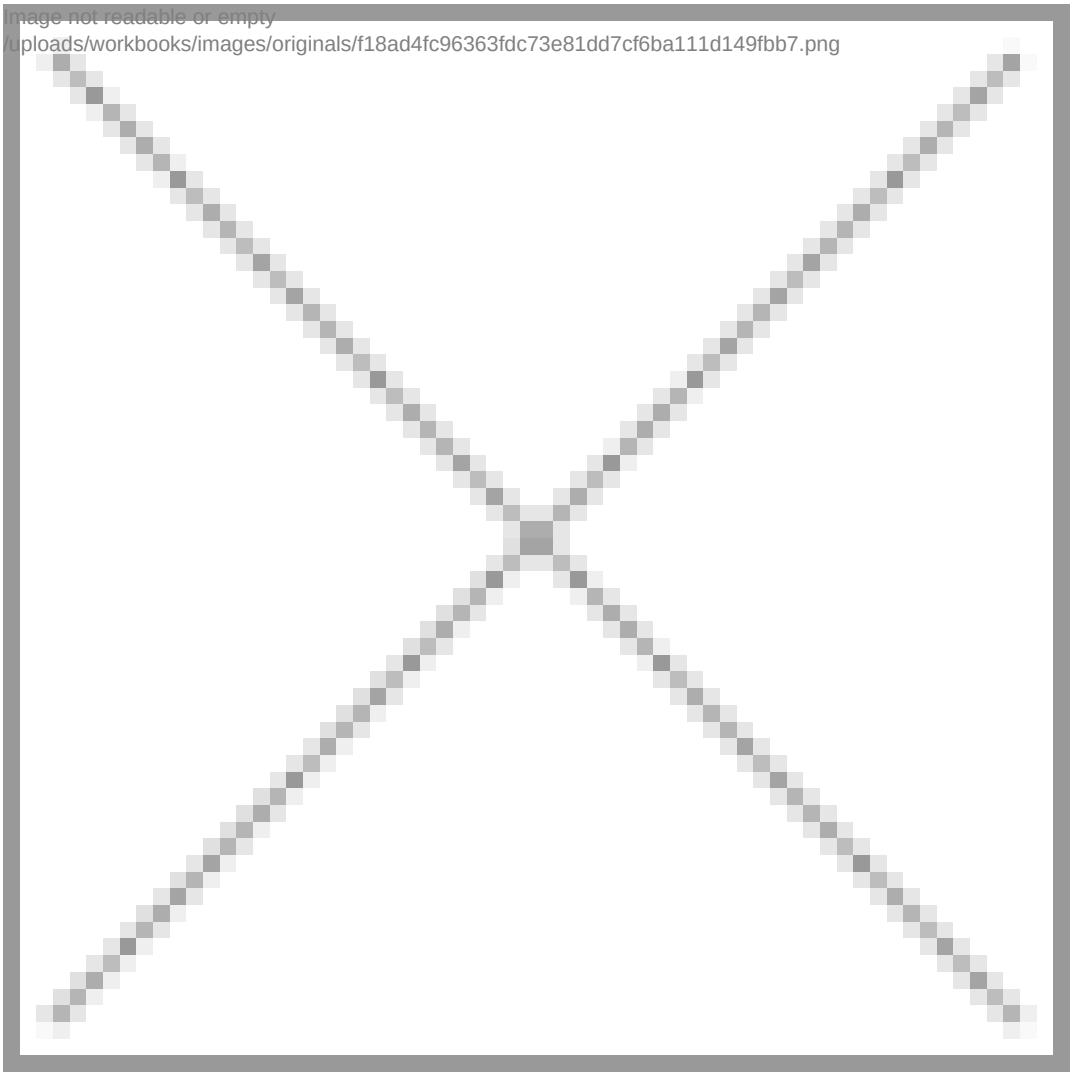


After settings have been changed, stop and start the CA service for these to take effect:

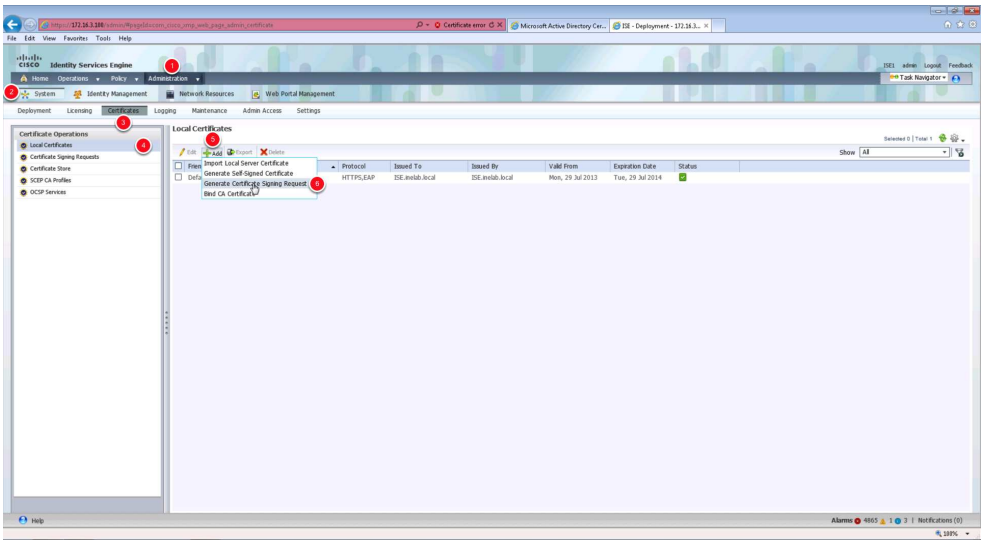
Image not readable or empty
/uploads/workbooks/images/originals/b7293b5dbdfe7d51c613e1c036daa4eee93c542a.png

https://t.me/learningnets

Image not readable or empty
/uploads/workbooks/images/originals/f18ad4fc96363fdc73e81dd7cf6ba111d149fbb7.png

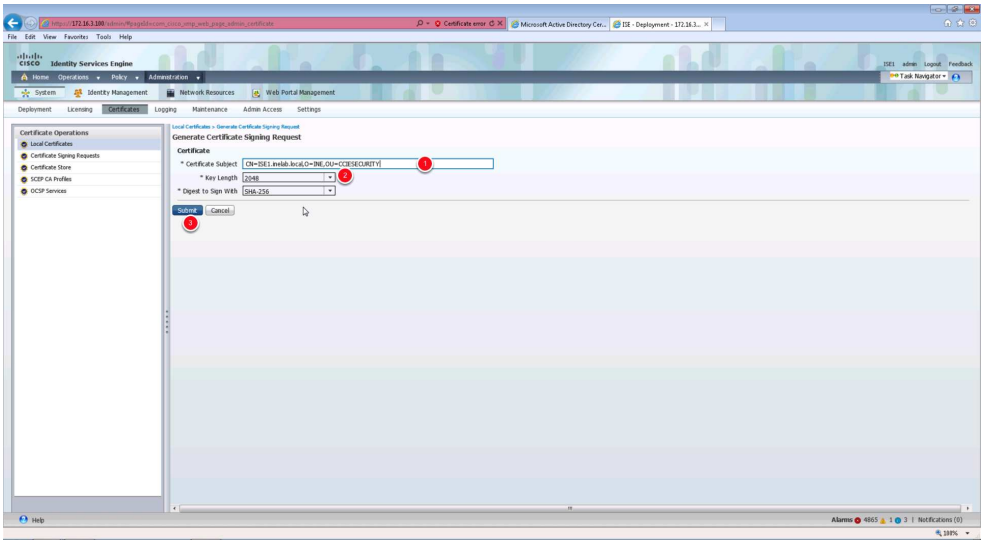


Browse ISE GUI to initiate a CSR request:

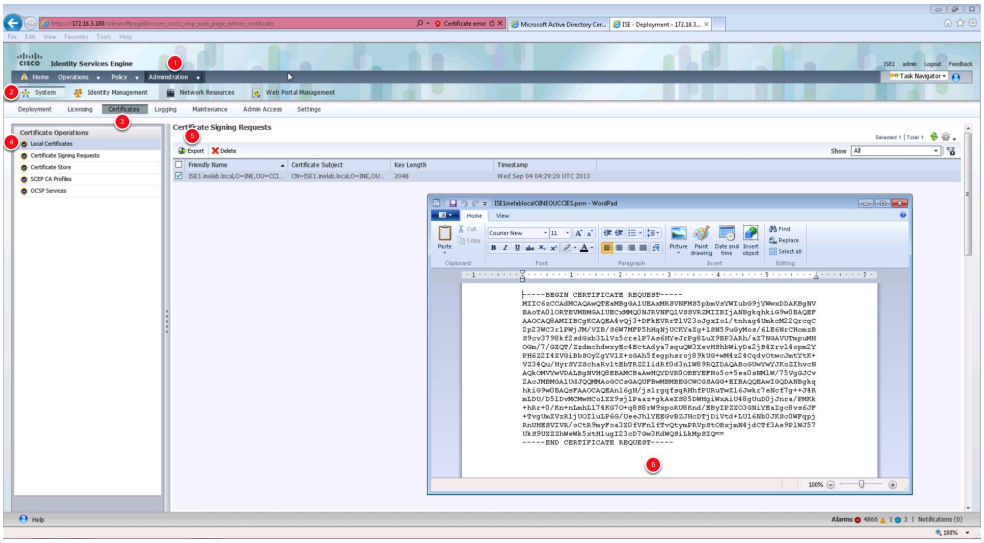


Configure the CSR data as requested by the task:

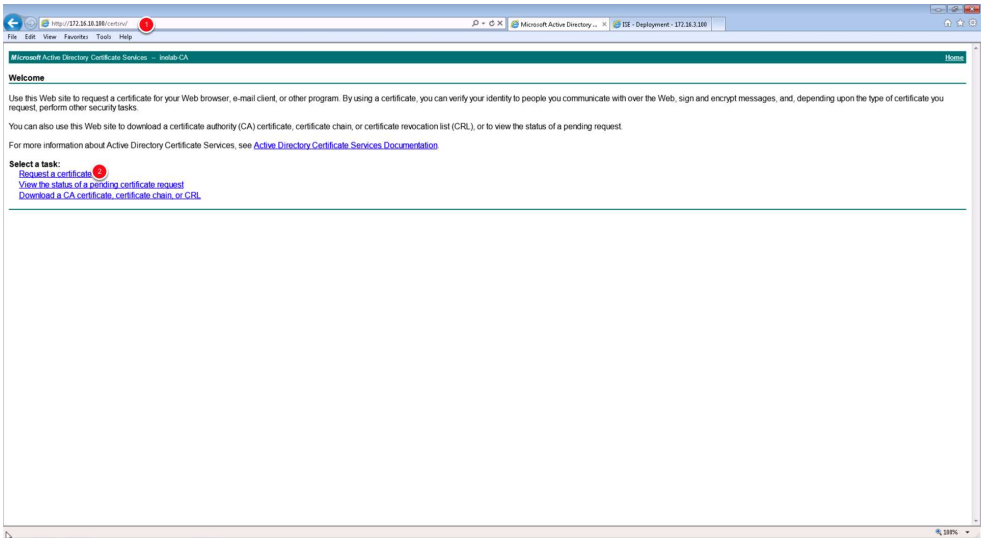
https://t.me/learningnets



Export the CSR data and save it into a WordPad file:

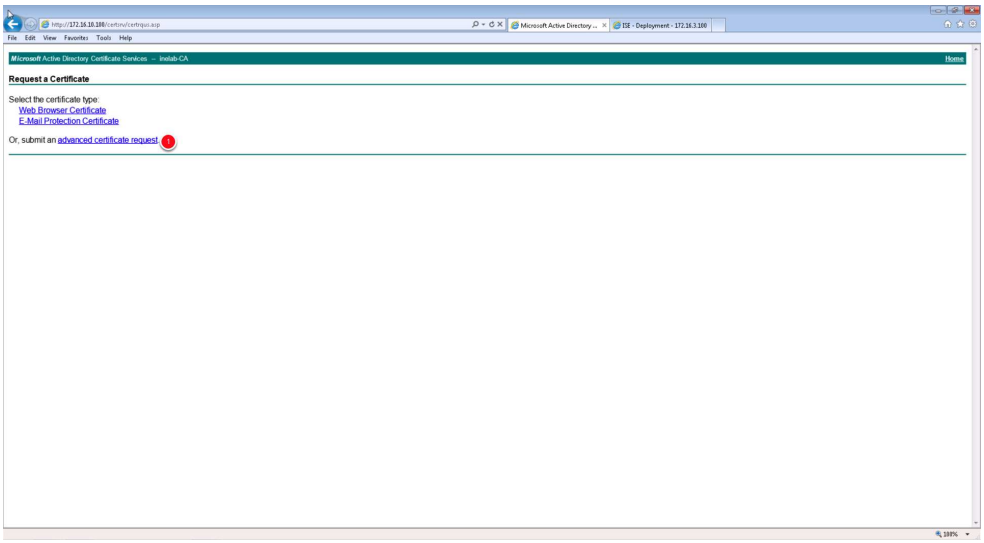


Browse to the CA server to request a certificate:

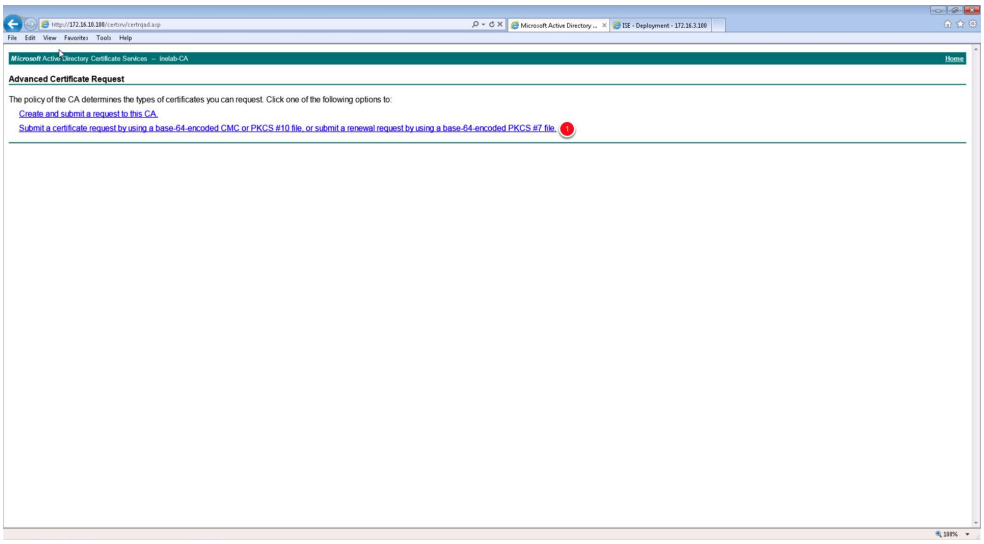


https://t.me/learningnets

Select the **Advanced Certificate Request** option to make use of the CSR data:

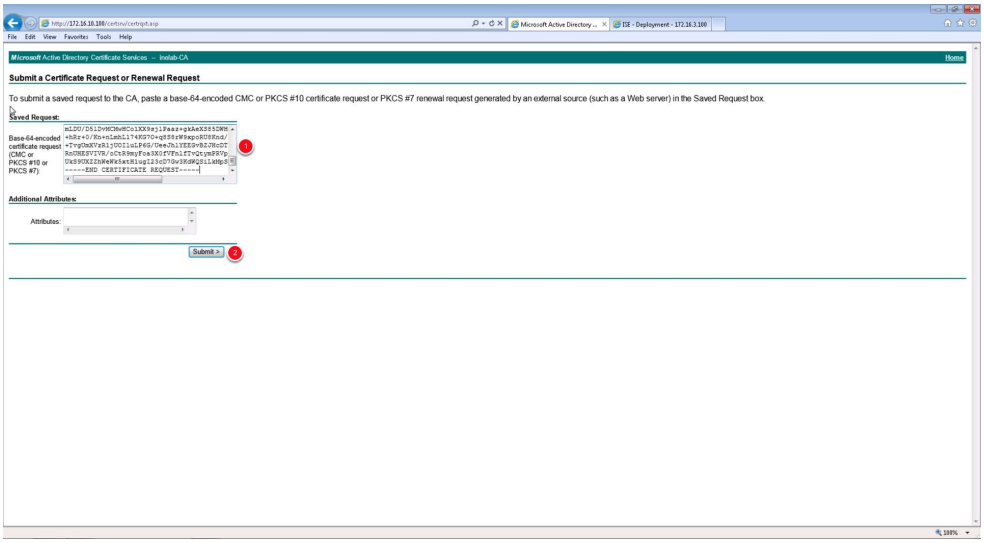


Select the second option to paste the CSR data:

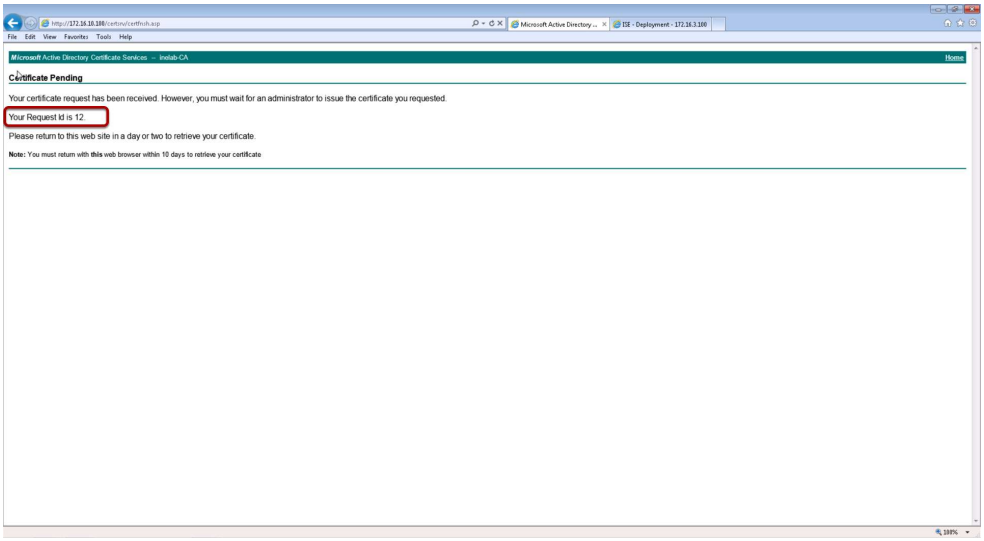


Paste the CSR data into the **Saved Reuqest** window:

https://t.me/learningnets

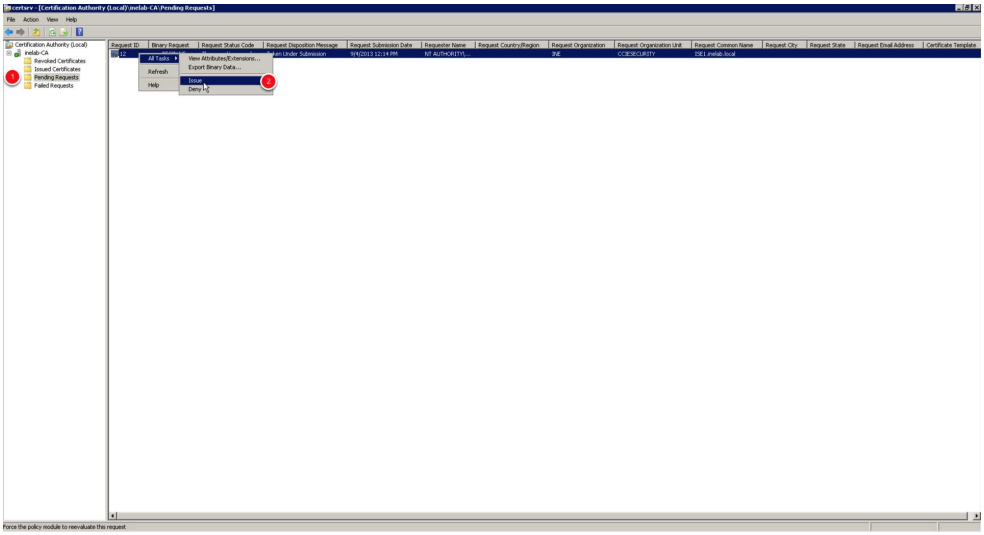


The presented screen instructs you, the CA administrator, to manually approve the certificate request:

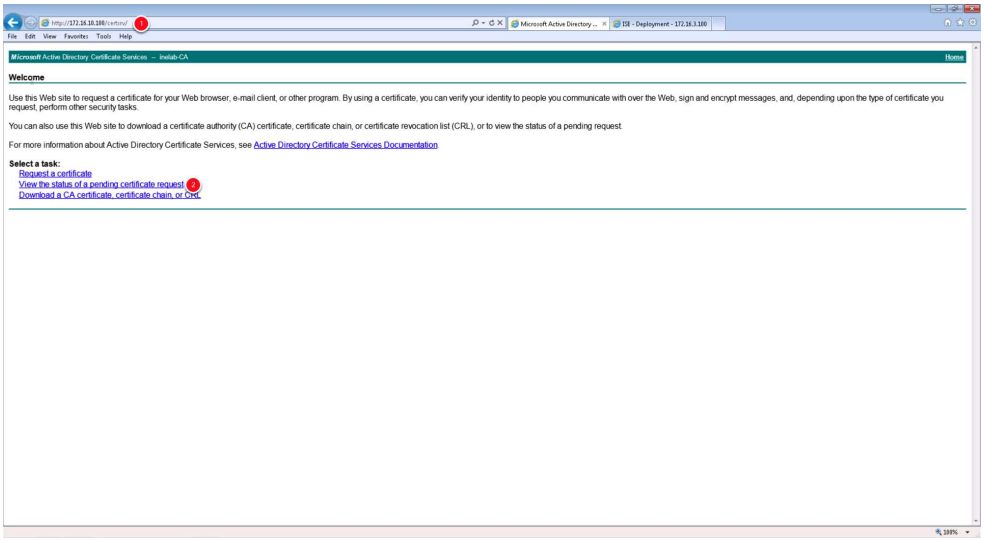


Open the CA management console using the command **certsrv.msc** from Run, go to **Pending Requests**, and grant the certificate:

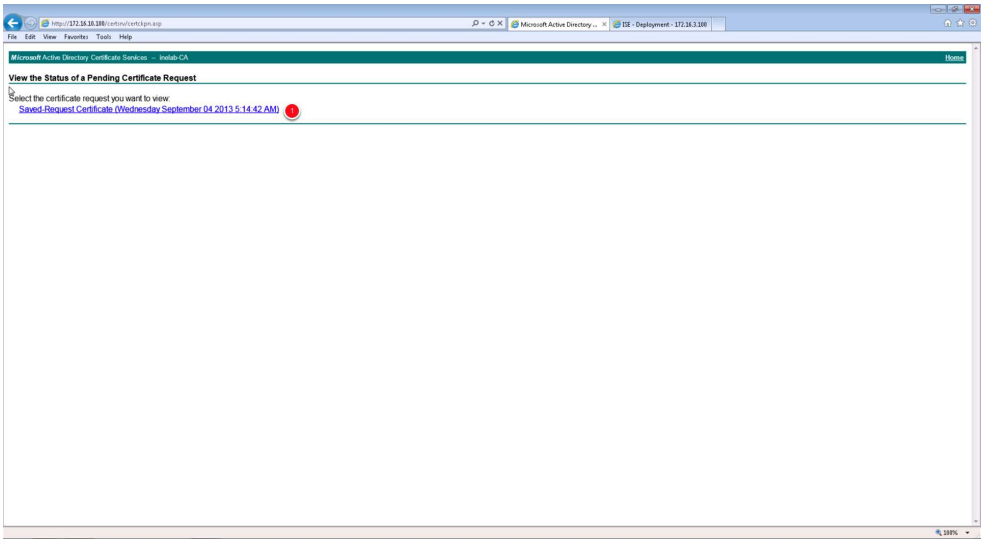
https://t.me/learningnets



Browse the CA server to retrieve the certificate:

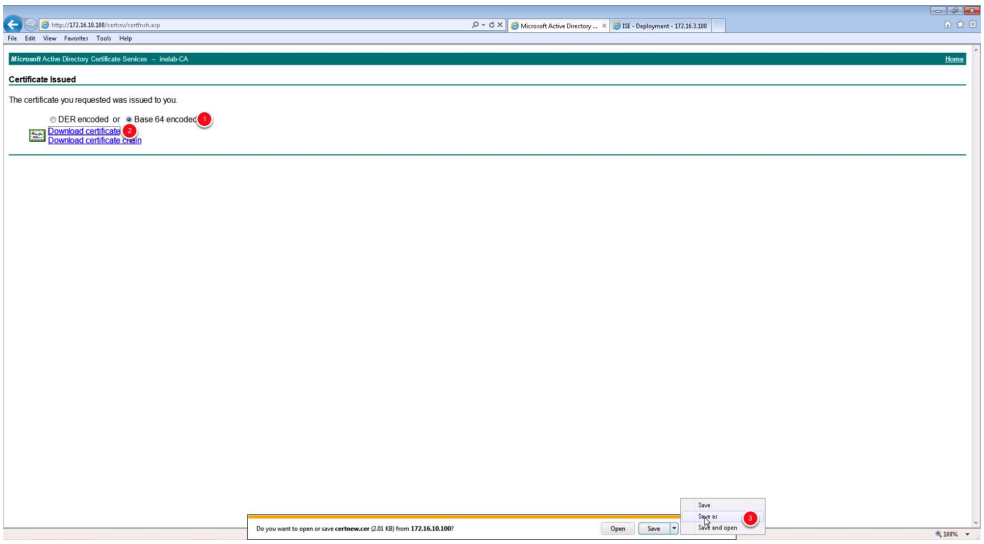


The certificate has been granted, so click to download it:

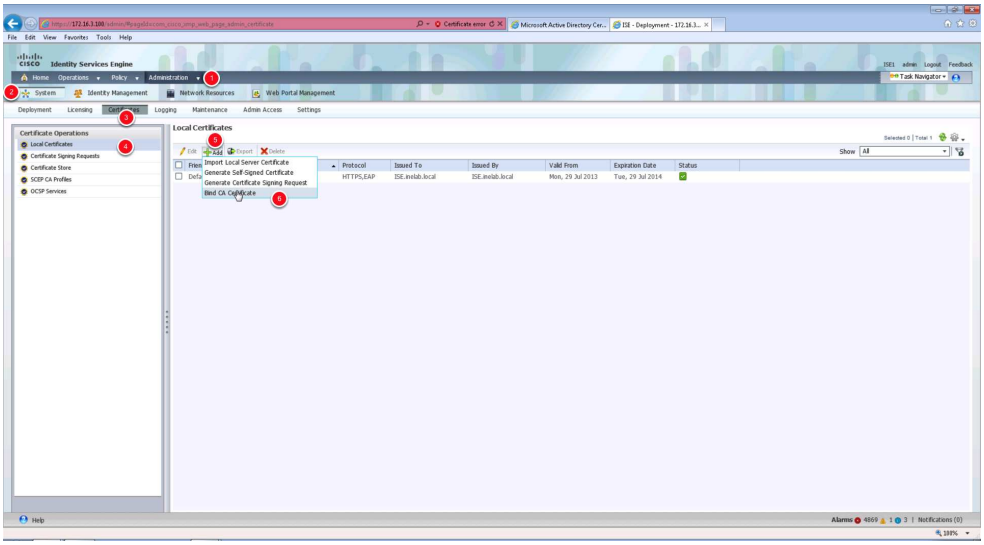


https://t.me/learningnets

Download the certificate in **Base 64** format and save it locally:

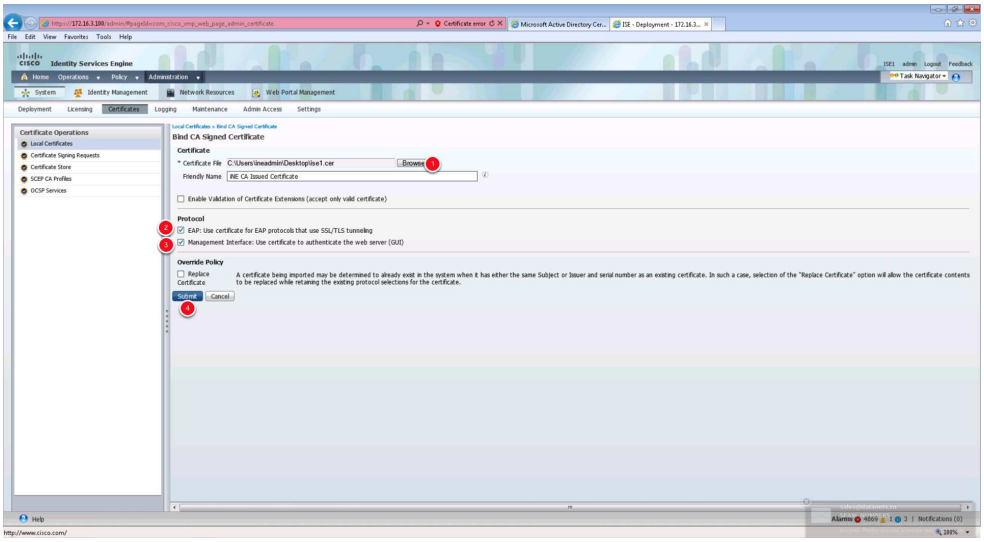


Install the certificate to ISE, binding it to the CSR:



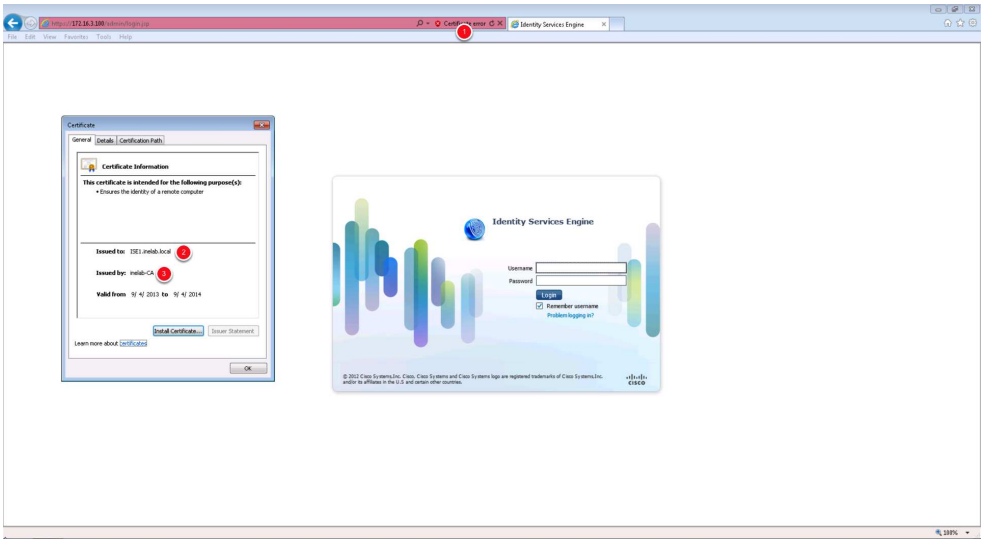
Browse to the saved certificate file and select to use it for both management and EAP:

https://t.me/learningnets



Verification

After certificate is installed, browse to ISE again and verify the certificate used for management:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

MAB Using ACS

Task

- Configure SW1 to authenticate R3's access on VLAN 37 using MAB. The authorization policy from the ACS server should be based on:
  - Service-Type
  - NAS-Port-Type
  - NAS-IP-Address
  - Authentication Method
- Configure SW2 to authenticate ASA3's access on VLAN 37 using MAB, but ensure that the switch does not send the **User-Password** attribute in the RADIUS Access-Request message. The authorization policy from the ACS server should be based on:
  - Service-Type
  - NAS-Port-Type
  - NAS-IP-Address
  - Authentication Method

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

For just MAB (Mac Authentication Bypass), enabling 802.1x globally on the switch

https://t.me/learningnets

with the command **dot1x system-auth-control** is not required, because MAB is not actually 802.1x, it's just a method to authenticate/authorize the port/MAC in one of the following cases:

- Connecting system does not have the 802.1x supplicant configured
- Connecting system does not support 802.1x, so it does not have a 802.1x supplicant
- Connecting system fails 802.1x authentication

When the switch uses MAB to authenticate/authorize the port, there are two possible implementations:

- Using the port-level **mab** command: The authentication method is PAP, so the username is sent in clear-text. The password is also sent but encrypted by RADIUS; RADIUS Attribute 31, **Calling-Station-Id**, which is the MAC address of the client, is also sent in clear-text.
- Using the port-level **mab eap** command: The authentication method is EAP-MD5, so the username is sent in clear-text. The password is hashed using CHAP-MD5; RADIUS Attribute 31, **Calling-Station-Id**, which is the MAC address of the client, is also sent in clear-text.

However, MAB-EAP does not really add any security to the authentication process; it can be used as a differentiator in the RADIUS server authorization policies (ACS or ISE), because whereas MAB uses a RADIUS **Service-Type** of **Call-Check**, MAB-EAP actually behaves like regular 802.1x authentication and uses a RADIUS **Service-Type** of **Framed**.

Cisco optimised the MAB process to have ACS use by default a **Host Lookup** mechanism for MAB authentication requests. By doing this, ACS/ISE ignores the **User-Name** and **User-Password** RADIUS attributes and thus the authentication is based on the **Calling-Station-Id** RADIUS attribute. Although these three attributes essentially have the same value, which is the MAC address, with the **Host Lookup** method there is no need to configure a username, which requires a password on the ACS/ISE side. ACS/ISE can be configured to do MAB as any other RADIUS server does, if the **Process Host Lookup** check box is cleared for the respective rule, in which case it is required to configure a username/password pair equal to the MAC address value.

Technically, MAB-EAP cannot use the same **Host Lookup** mechanism for two reasons:

- The RADIUS **Service-Type** attribute value is **Framed**, and **Host Lookup** works only for when attribute value is **Call-Check**.
- EAP-MD5 uses the CHAP protocol to create a hash of the password, so it actually

https://t.me/learningnets

requires the authentication to be username/password based.

However, Cisco implemented the **Host Lookup** also for MAB-EAP, just that it works differently: authentication is performed based on the **User-Name** RADIUS attribute and this will be compared against the internal endpoint store, so no need for username/password based authentication.

With MAB, the authorization received from the RADIUS server can have both a VLAN and a dACL.

Configuration

```
SW1:
aaa new-model
aaa authentication dot1x default group radius
!
interface FastEthernet1/0/3
 authentication port-control auto
 authentication host-mode single-host
 mab
!
radius-server host 172.16.1.100 key CISCO
ip radius source-interface Loopback0

SW2:

aaa new-model
aaa authentication dot1x default group radius
!
interface FastEthernet1/0/16
 authentication port-control auto
 authentication host-mode single-host
 mab eap
!
radius-server host 172.16.1.100 key CISCO
ip radius source-interface Vlan19
```

Identify the MAC address of R3 to be added to the ACS Internal Hosts Store; note that this MAC address is different depending on the rack being assigned:

```
SW1#show mac address-table interface fastEthernet1/0/3
      Mac Address Table
-------------------------------------------
```

https://t.me/learningnets

```
Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
37    001e.f779.4770    DYNAMIC     Fa1/0/3

Total Mac Addresses for this criterion: 1
```

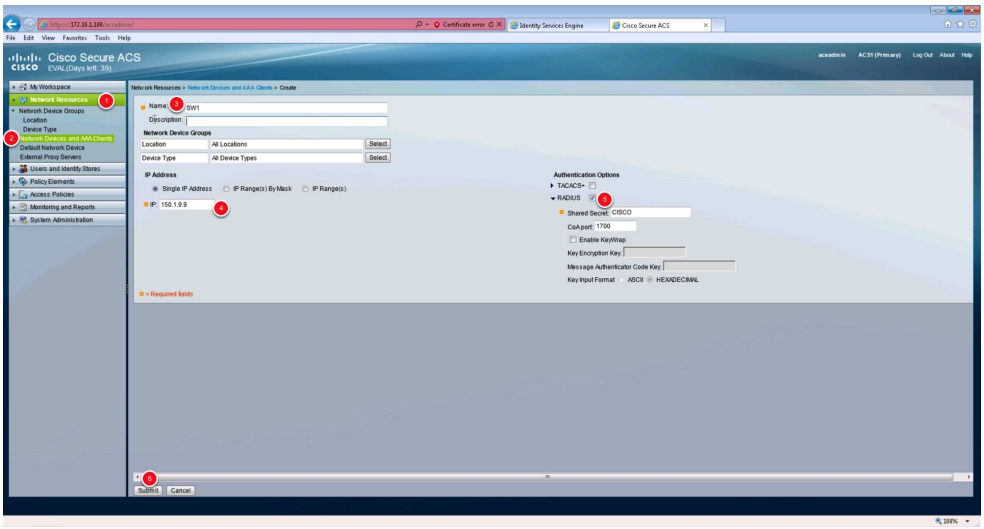
Identify the MAC address of ASA3 to be added to the ACS Internal Users Store; note that this MAC address is different depending on the rack being assigned:

```
SW2#show mac address-table interface fastEthernet1/0/16
      Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
37    a493.4caa.f508    STATIC      Fa1/0/16

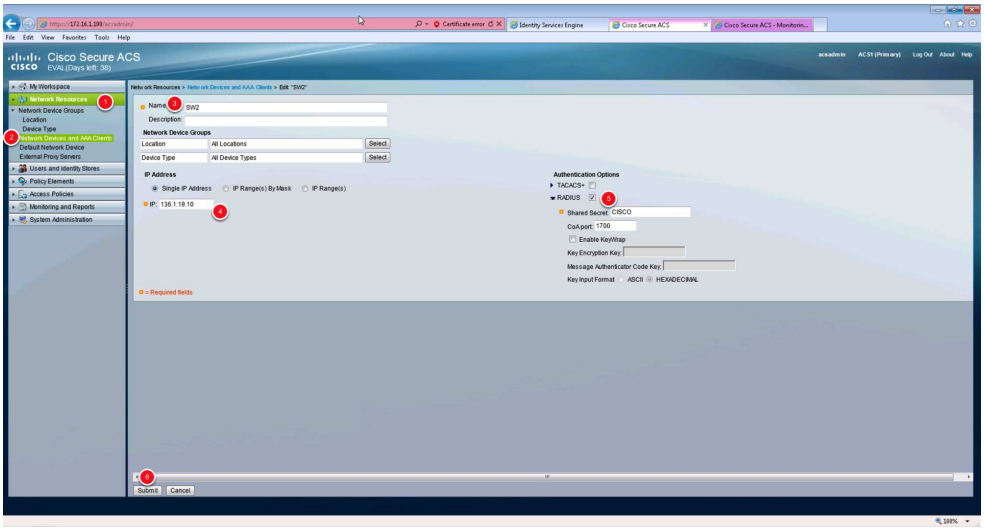
Total Mac Addresses for this criterion: 1
```

Add SW1 as a NAD into ACS, and enable it for RADIUS only:

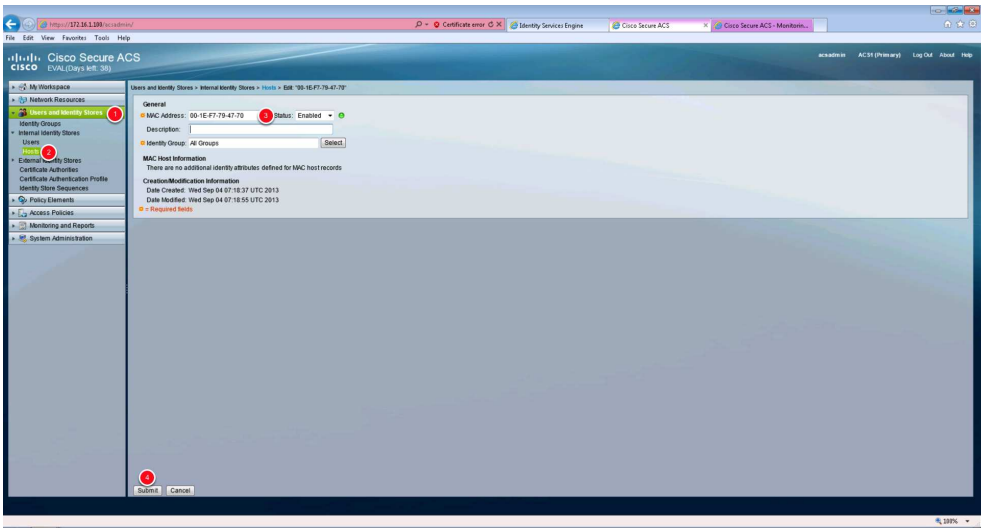


Add SW2 as a NAD into ACS, and enable it for RADIUS only:

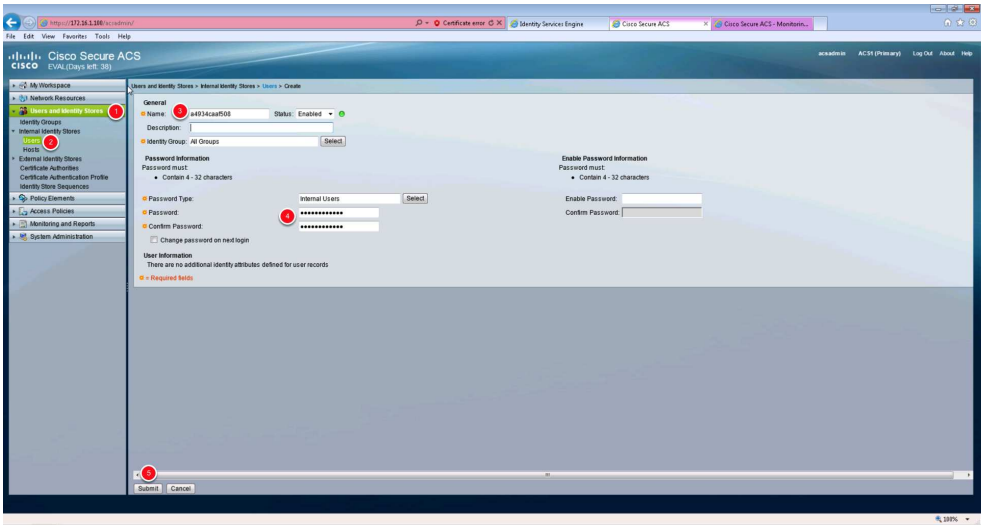
https://t.me/learningnets



Add R3's MAC address into the Internal Hosts Store:



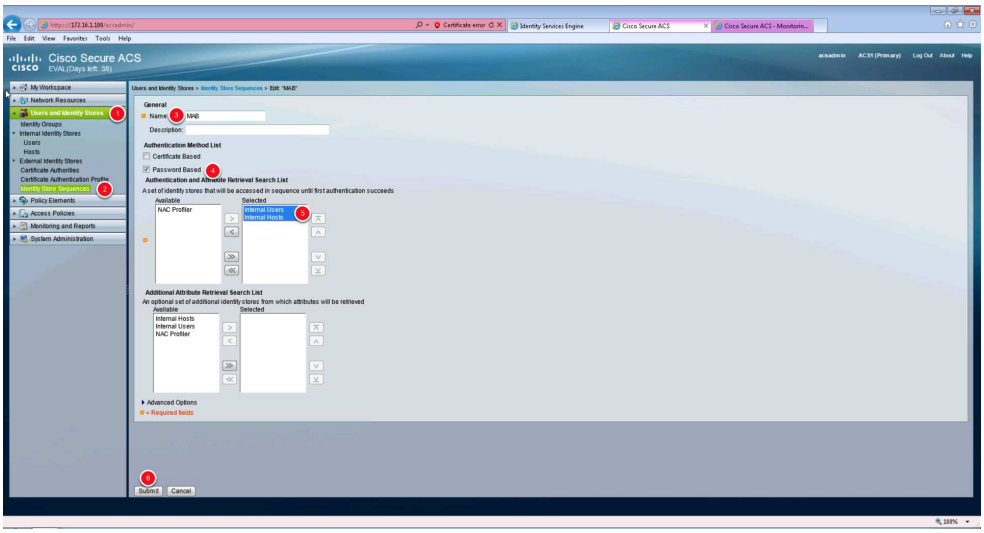
Add ASA3's MAC address as a username into the Internal Users Store; the password is also the MAC address value:



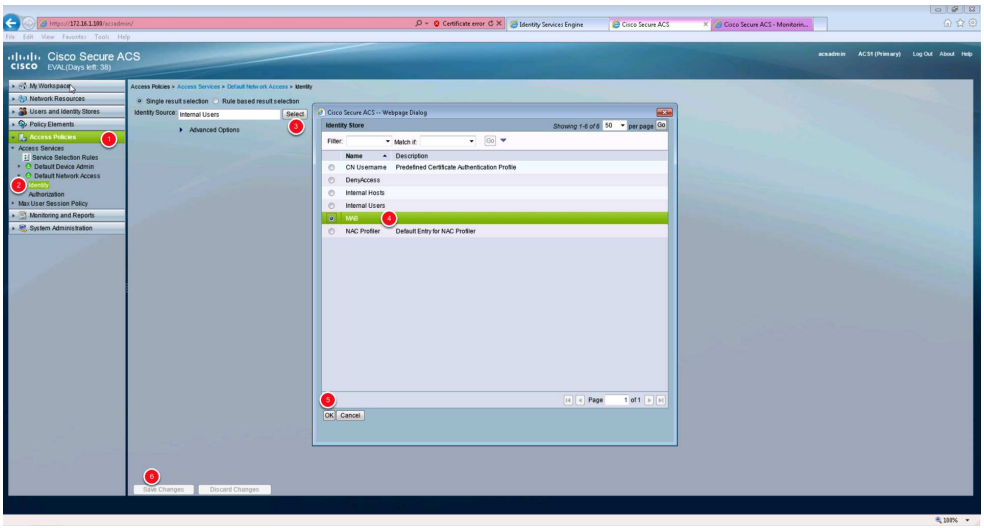
Configure an Identity Store Sequence that uses both Internal Hosts and Internal

https://t.me/learningnets

Users Stores:

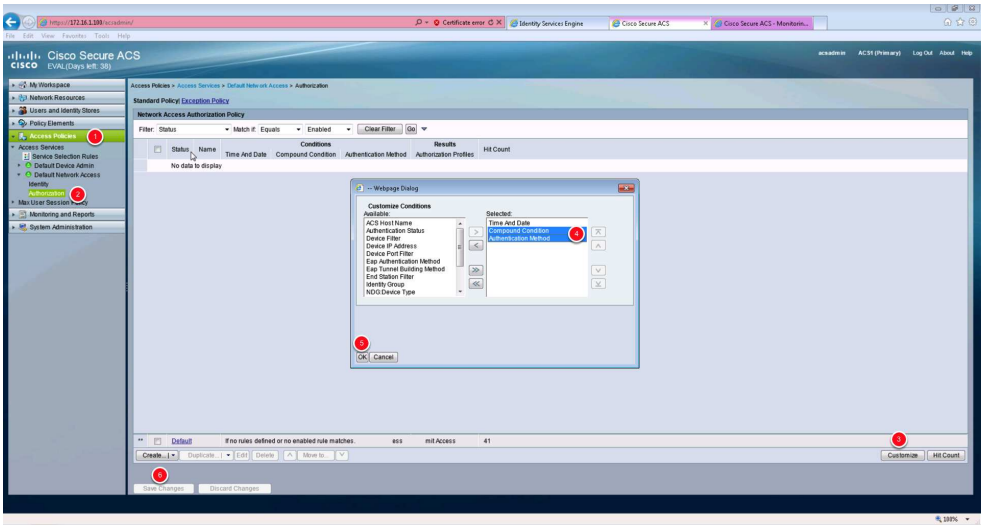


Configure the Default Network Access policy to use the configured Identity Store Sequence:

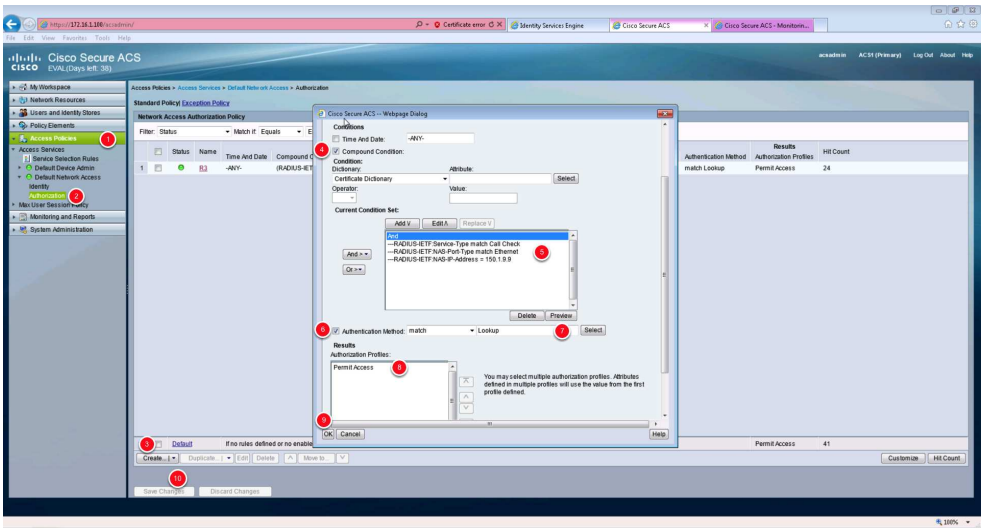


Add **Compound Condition** and **Authentication Method** as conditions for the **Default Network Access** authorization policies:

https://t.me/learningnets

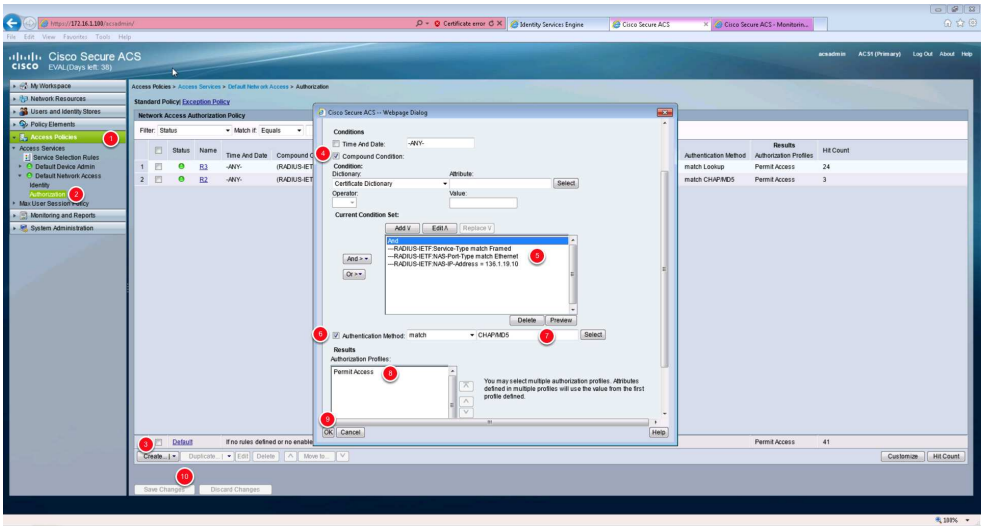


Configure the authorization policy for R3 to match the task requirements, which uses MAB:



Configure the authorization policy for ASA3 to match the task requirements, which uses MAB-EAP:

https://t.me/learningnets



Verification

Verify which ports have MAB enabled on both SW1 and SW2:

```
SW1#show mab all

MAB details for FastEthernet1/0/3
--------------------------------------Mac-Auth-Bypass            = Enabled
!
!SW2#show mab all
MAB details for FastEthernet1/0/16
--------------------------------------Mac-Auth-Bypass            = Enabled (EAP)
```

Verify successful authentication on both SW1 and SW2:

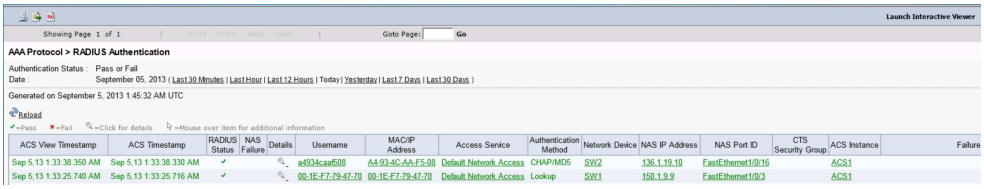
```
SW1#show authentication sessions interface fastEthernet1/0/3
        Interface:  FastEthernet1/0/3
      MAC Address:  001e.f779.4770
       IP Address:  Unknown
        User-Name:  00-1E-F7-79-47-70 Status:  Authz Success
           Domain:  DATA
  Security Policy:  Should Secure
  Security Status:  Unsecure Oper host mode:  single-host
 Oper control dir:  both Authorized By:  Authentication Server
       Vlan Group:  N/A
  Session timeout:  N/A
     Idle timeout:  N/A
Common Session ID:  8801130900000242EF79DE4
  Acct Session ID:  0x00000025
           Handle:  0x96000024
```

https://t.me/learningnets

```
Runnable methods list:
   Method   State mab        Authc Success
!
!SW2#show authentication sessions interface fastethernet1/0/16
        Interface:  FastEthernet1/0/16
      MAC Address:  a493.4caa.f508
       IP Address:  Unknown
        User-Name:  a4934caaf508 Status:  Authz Success
           Domain:  DATA
  Security Policy:  Should Secure
  Security Status:  Unsecure Oper host mode:  single-host
 Oper control dir:  both Authorized By:  Authentication Server
       Vlan Group:  N/A
  Session timeout:  N/A
     Idle timeout:  N/A
Common Session ID:  8801130A0000000B796FC254
  Acct Session ID:  0x0000000C
           Handle:  0xF100000B

Runnable methods list:
   Method   State mab        Authc Success
```

Verify RADIUS Authentication Logs in the ACS Server, and that **Hit Counts** are increasing for your configured Authorization Policies to ensure that these are matched, because if configured policies are not matched, the default rule is **Permit Access**; thus, it acts as a gateway of last resort, allowing access to the network.



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

MAB Using ISE

Task

- Configure SW1 to authenticate R3's access on VLAN 37 using MAB. The authorization policy from the ISE server should be based on:
  - Service-Type
  - NAS-Port-Type
  - NAS-Port-Id
  - Authentication Method
- Profile R3 as a **Cisco Device** and ensure that ISE profiler cannot change its identity group of **Cisco Device**.
- Configure SW2 to authenticate ASA3's access on VLAN 37 using MAB, but ensure that the switch does not send the **User-Password** attribute in the RADIUS Access-Request message. The authorization policy from the ISE server should be based on:
  - Service-Type
  - NAS-Port-Type
  - NAS-Port-Id
  - Authentication Method
- ASA3 is a member of the **Employee** user group.
- In both cases, the access VLAN should be received from ISE.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.

https://t.me/learningnets

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

For any authorization received from the RADIUS server (ACS or ISE), which usually is a VLAN or dACL, the **aaa authorization network** command is required; otherwise, the switch ignores the authorization attributes received.

IOS supports three methods for ACL authorization thorugh MAB or 802.1x:

- dACL (Dynamic Access-List):
  - Requires an initial ACL applied on the port before authentication (depending on IOS code version, a default one may be applied by the switch implicitly).
  - Requires `aaa authorization network default group` and `radius-server vsa send authentication` configured on on switch.
  - Verfied with `show authentication session interface [interface-name]` , `show ip access-list interface [interface-name]` and `show epm session ip [IP Address]` commands.
  - Requires `ip device tracking` command so ACL is applied by EPM.
  - It's Cisco proprietary.

- Filter-ID ACL
  - Requires ACL configured on switch and ACL name pushed from RADIUS server as a RADIUS attribute.
  - Requires `aaa authorization network default group` command configured on switch.
  - Requires Filter-ID attribute on RADIUS server applied in format of **ACL-NAME.in** or command `radius-server attribute 11 default direction inbound` configured on switch.
  - Verfied with `show authentication session interface [interface-name]` , `show ip access-list interface [interface-name]` and `show epm session ip [IP Address]` commands.
  - Requires `ip device tracking` command so ACL is applied by EPM.
  - Uses a standard IETF RADIUS attribute of 11.

https://t.me/learningnets

- Per-User ACL
  - First possible format configured on RADIUS server is **ip:inacl#=acl-number** or **ip:inacl#=acl-name** and requires ACL with given number or name configured on switch.
  - Second possible format configured on RADIUS server is **ip:inacl#n=acl-rule** , for example **ip:inacl#1=permit icmp any any** to allow all ICMP traffic.
  - Requires `aaa authorization network default group` command configured on switch.
  - Verfied with `show authentication session interface [interface-name]` , `show ip access-list interface [interface-name]` and `show epm session ip [IP Address]` commands.
  - Requires `ip device tracking` command so ACL is applied by EPM.
  - Uses Cisco AV Pair.

Configuration

SW1:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
interface FastEthernet1/0/3
 authentication port-control auto
 authentication host-mode single-host
 mab
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Loopback0
```

SW2:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
interface FastEthernet1/0/16
 authentication port-control auto
 authentication host-mode single-host
 mab eap
!
aaa group server radius ISE
```

https://t.me/learningnets

```
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Vlan19
```

Identify the MAC address of R3 in order to be added to the ISE Internal Hosts Store, note that this MAC address is different depending on the rack being assigned:

```
SW1#show mac address-table interface fastEthernet1/0/3
      Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  37    001e.f779.4770    DYNAMIC     Fa1/0/3

Total Mac Addresses for this criterion: 1
```

Identify the MAC address of ASA3 in order to be added to the ISE Internal Users Store, note that this MAC address is different depending on the rack being assigned:

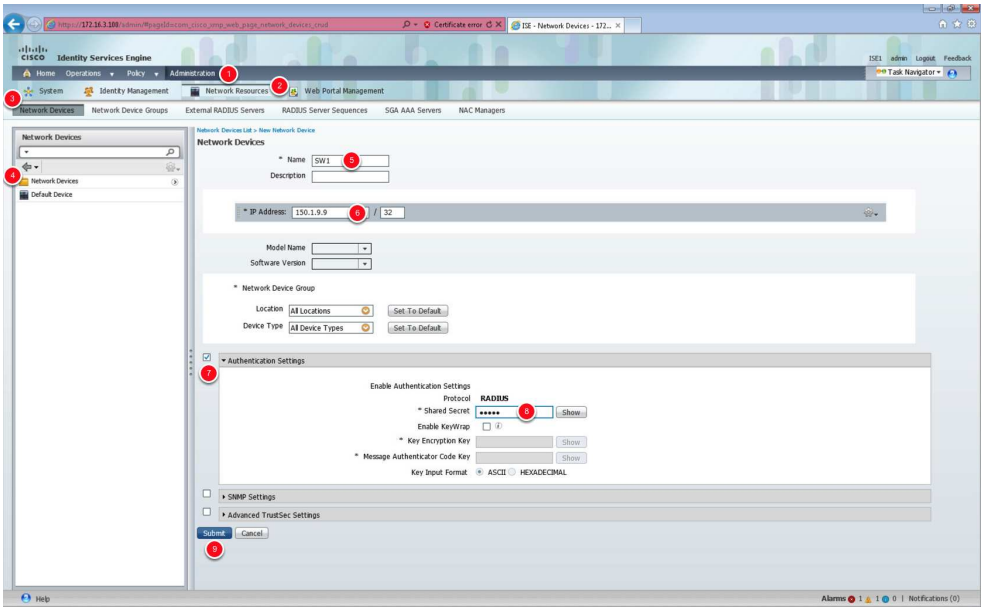
```
SW2#show mac address-table interface fastEthernet1/0/16
      Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  37    a493.4caa.f508    STATIC      Fa1/0/16

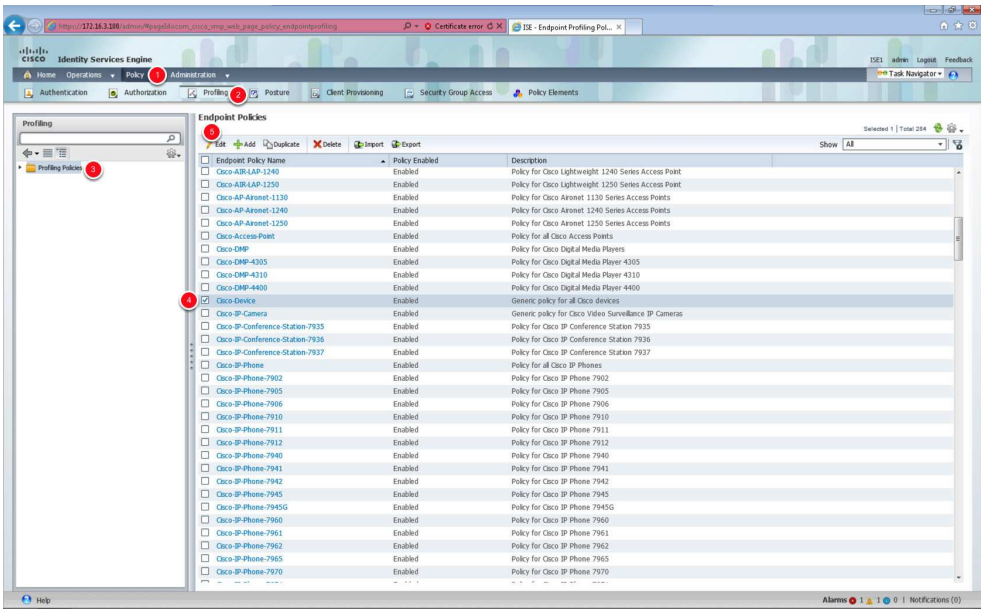
Total Mac Addresses for this criterion: 1
```

On ISE, first add SW1 and SW2 as NAD devices, example is given for SW1:

https://t.me/learningnets

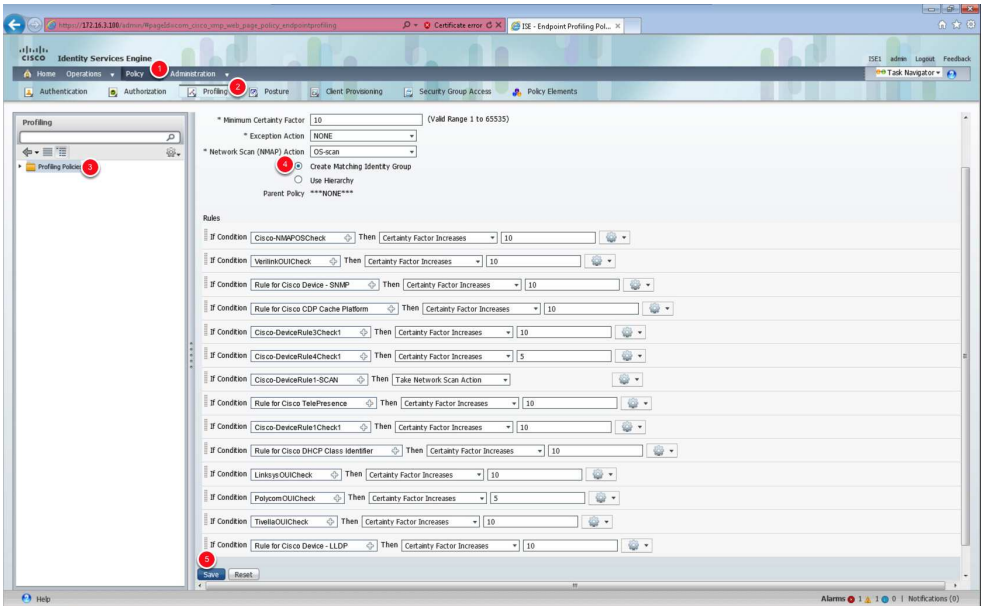


Navigate to the profiler and edit **Cisco Device** policy:

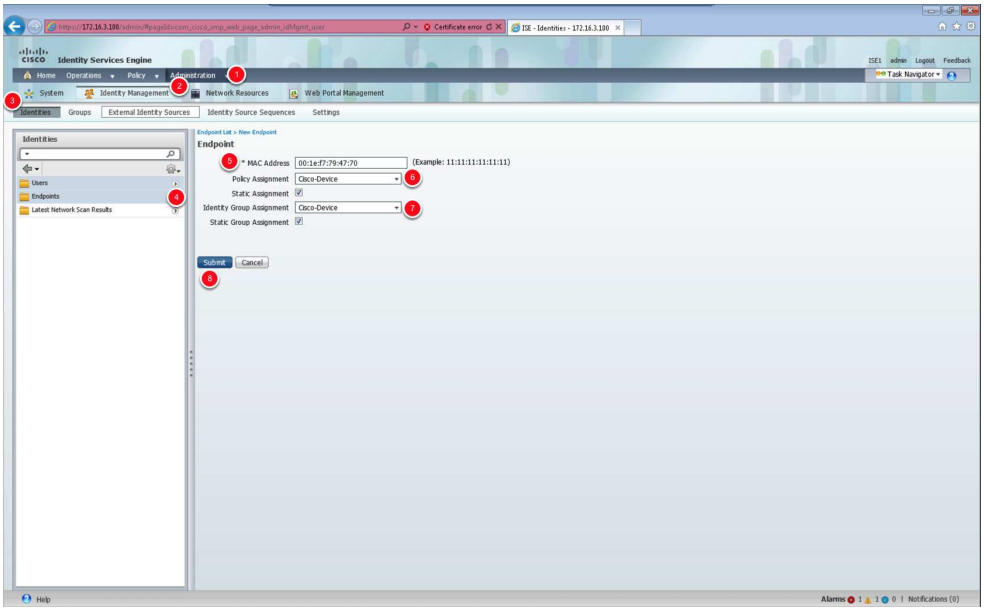


Edit the **Cisco Device** policy so that a matching **Identity Group** is created for profiled devices:

https://t.me/learningnets

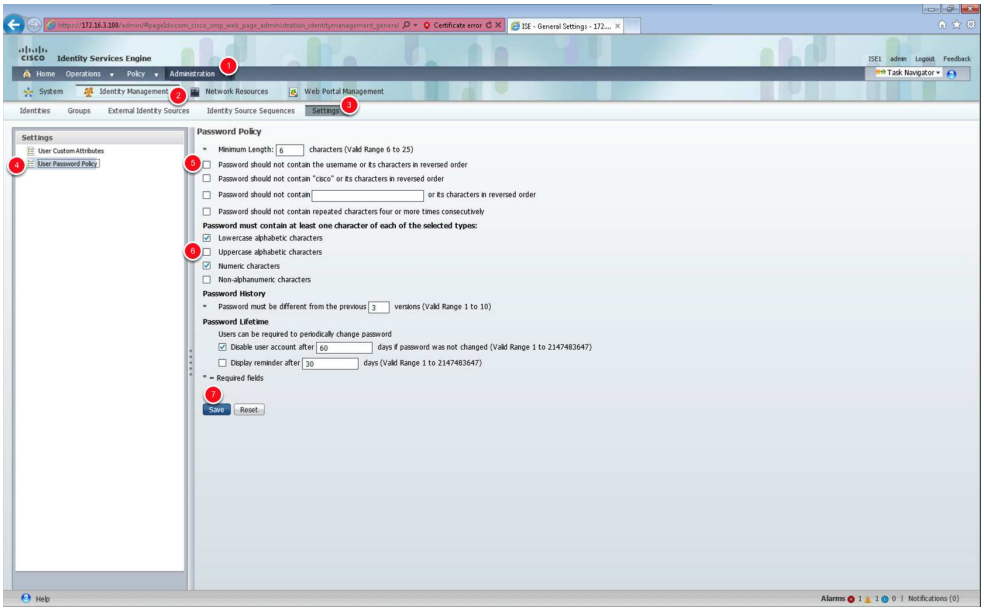


Add the MAC address of R3 in the internal Endpoint store, in order to be used for MAB:

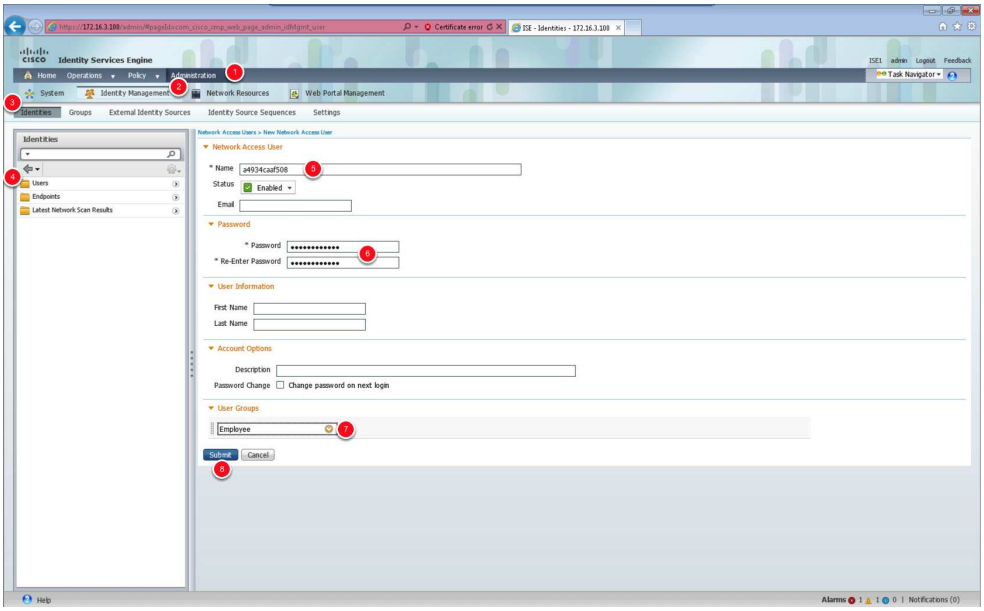


In order to add a user account with a password identical with the username, and without uppercase letters, the default user password policy needs to be changed:

https://t.me/learningnets

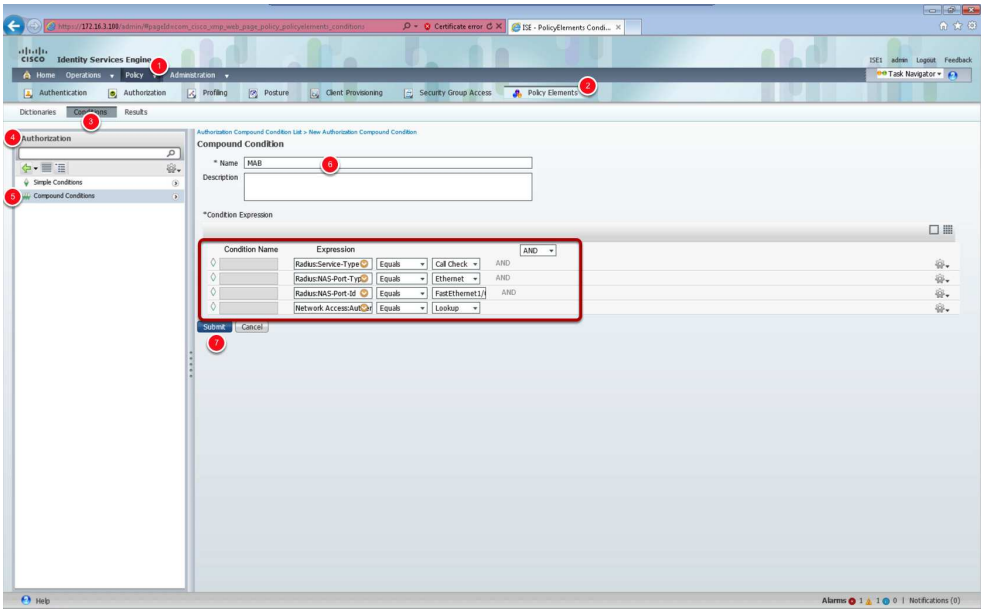


Add ASA3's MAC address into the internal User store, in order to be used for MAB-EAP, password is the MAC address:

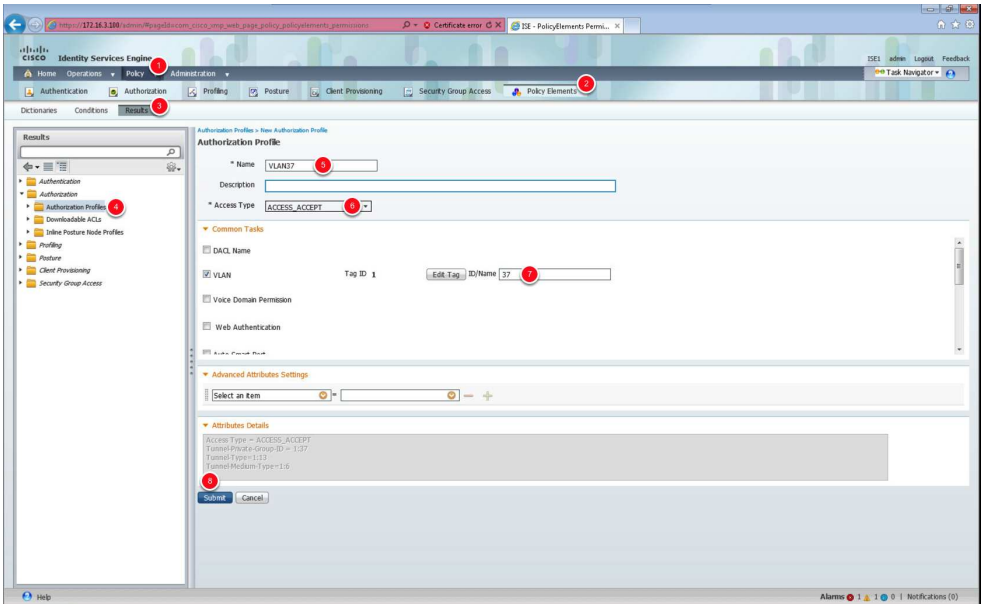


Configure the compound authorization condition for MAB, for MAB-EAP it will be edited inline in the authorization policy:

https://t.me/learningnets

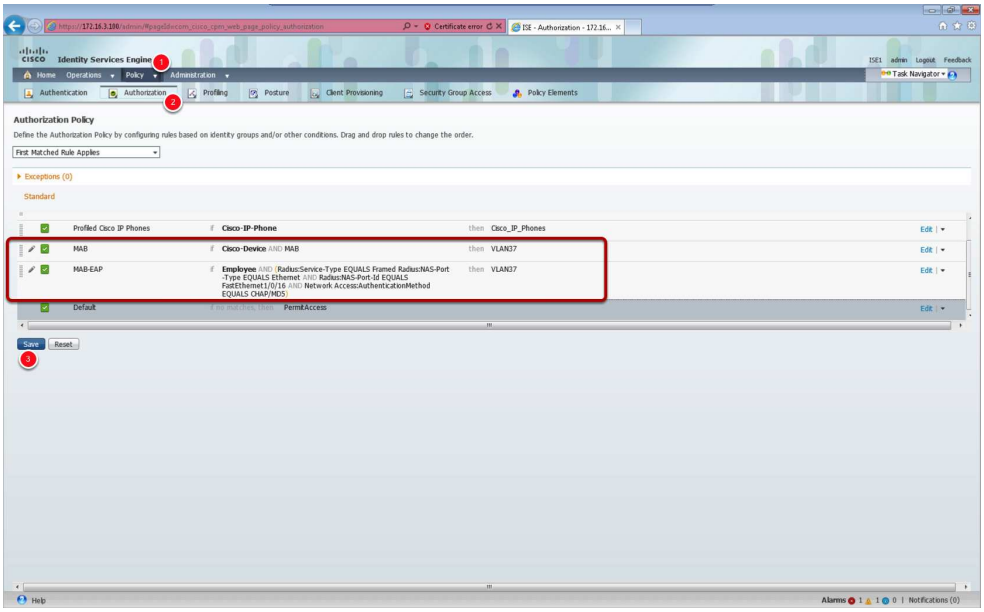


Configure the authorization profile in order to aloww access and assign the access VLAN:



The default authentication polices will allow both devices to authenticate, we need to configure the authorization policies as requested:

https://t.me/learningnets



Verification

Verify on which ports is MAB enabled and the MAB method (PAP or EAP):

```
SW1#show mab all
MAB details for FastEthernet1/0/3
--------------------------------------Mac-Auth-Bypass            = Enabled
!
!SW2#show mab all

MAB details for FastEthernet1/0/16
--------------------------------------Mac-Auth-Bypass            = Enabled (EAP)
```

Verify that ports are both authenticated and authorized (VLAN is received from ISE):

```
SW1#show authentication sessions interface fastethernet1/0/3
         Interface:  FastEthernet1/0/3
       MAC Address:  001e.f779.4770
        IP Address:  Unknown
         User-Name:  00-1E-F7-79-47-70 Status:  Authz Success
            Domain:  DATA
   Security Policy:  Should Secure
   Security Status:  Unsecure Oper host mode:  single-host
  Oper control dir:  both Authorized By:  Authentication Server
 Vlan Policy:  37
   Session timeout:  N/A
      Idle timeout:  N/A
Common Session ID:  000000000000001E053F6CB5
```

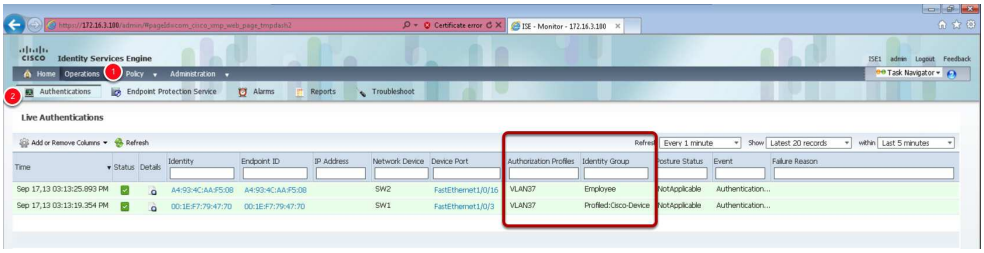
https://t.me/learningnets

```
  Acct Session ID:  0x0000001F
           Handle:  0x1400001E

Runnable methods list:
   Method   State mab        Authc Success
!
!SW2#show authentication sessions interface fastethernet1/0/16
        Interface:  FastEthernet1/0/16
      MAC Address:  a493.4caa.f508
       IP Address:  Unknown
        User-Name:  a4934caaf508 Status:  Authz Success
           Domain:  DATA
  Security Policy:  Should Secure
  Security Status:  Unsecure Oper host mode:  single-host
 Oper control dir:  both Authorized By:  Authentication Server
Vlan Policy:  37
  Session timeout:  N/A
     Idle timeout:  N/A
Common Session ID:  00000000000000002052323D4
  Acct Session ID:  0x00000003
           Handle:  0x9A000002

Runnable methods list:
   Method   State mab        Authc Success
```

Verify the ISE authentication logs to ensure the configured authorization policies are matched, as well as the Identity Groups:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

PEAP (EAP-MSCHAPv2) Using ACS and AD

Task

- Configure **TEST-PC-A** for 802.1x authentication using PEAP. Authenticate with **USER1** and password **1qaz@WSX** from **inelab.local** domain.
- For joining ACS into **inelab.local** domain, use **ldapuser** with the password **1qaz@WSX**.
- Authorize no more than one MAC address per domain (VOICE and DATA); otherwise, drop traffic and log.
- Your PC authorization rule from ACS must include the **FINANCIAL** user group from Active Directory.
- Ensure that your configuration authorizes the IP Phone as fast as possible.
- SW1 should retransmit the EAP request ID packets at a 20-second interval; if the ACS server is not reachable, it should authorize both domains (VOICE and DATA).

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Cisco ACS can integrate with external identity sources to validate credentials from various authentication requests, and to retrieve group information as well as other attributes that are associated with the user, and uses them in authorization policies. Integration with a Certificate Authority is also seen as an external identity source.

https://t.me/learningnets

For ACS to join an Active Directory domain, except for having IP connectivity with the Domain Controller, there are three additional requirements:

- Clocks of ACS and DC should be synchronized.
- Domain name configured in ACS CLI should match the joining domain.
- DNS server configured on ACS CLI should be the IP address of the Domain Controller.

At a minimum, considering that there might be a firewall between ACS and DC, the following services must be allowed:

Protocol	Port
LDAP	UDP/389
SMB1	TCP/445
KDC2	TCP/88
Global Catalog	TCP/3268 , 3269
KPASS	TCP/464
NTP	UDP/123
LDAP	TCP/389
LDAPS	TCP/636

There are two versions of PEAP, v0(Windows XP) and v1 (Windows 7). One major difference between these versions is that v1 supports the identity protection functionality, where the username or CN from the certificate is not sent before the secure tunnel is established between the supplicant and the authentication server. PEAP can use EAP-MSCHAPv2, EAP-GTC, and EAP-TLS as the inner authentication methods, this task uses EAP-MSCHAPv2 because a

https://t.me/learningnets

username/password authentication is requested. PEAP requires a certificate only on the server (ACS) side to build the secure tunnel. In real deployments, the client will be configured to validate the server's certificate, in which case ACS needs a certificate issued from an external CA (not the self-signed one) and the client will need to have the certificate of the CA imported and configured as trusted root.

Configuration

SW1:

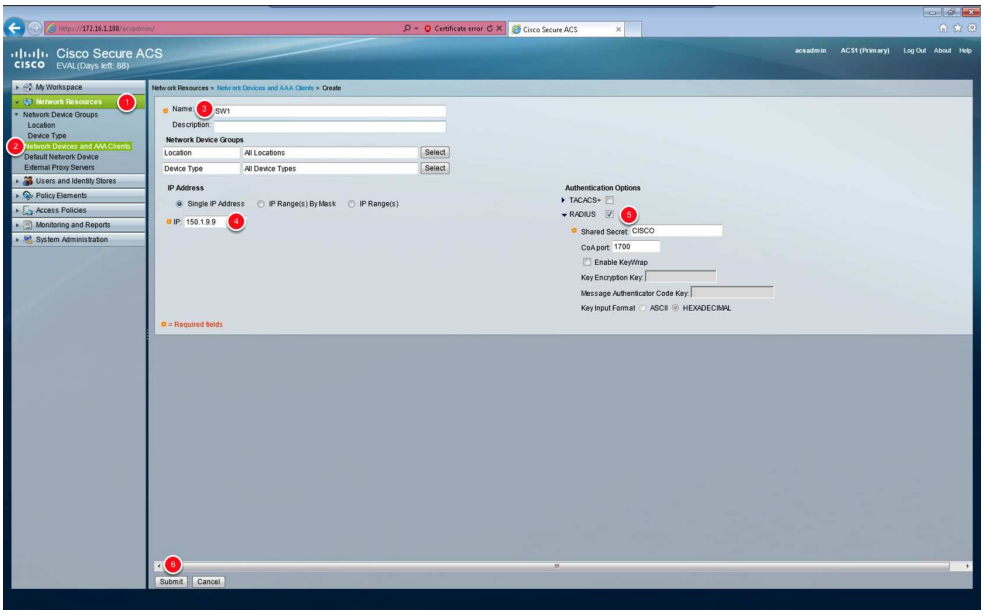
```
aaa new-model
aaa authentication dot1x default group ACS
aaa authorization network default group ACS
!
dot1x system-auth-control
!
interface FastEthernet1/0/5
 authentication host-mode multi-domain
 authentication port-control auto
 authentication order mab dot1x
 authentication violation restrict
 authentication event server dead action authorize voice
 authentication event server dead action authorize vlan 27
 dot1x pae authenticator
 dot1x timeout tx-period 20
 mab
!
aaa group server radius ACS
 server-private 172.16.1.100 key CISCO
 ip radius source-interface Loopback0
```

ACS is configured to synchronize time with the NTP server at R1. To join ACS into Active Directory, the time difference between ACS and Domain Controller should be within 5 minutes. Because Domain Controller (TEST-SRV-A) cannot be synchronized with R1, make sure to manually set the time on R1 or directly into ACS, if required, to match the time from TEST-SRV-A. Then wait until ACS synchronizes its time again with R1.

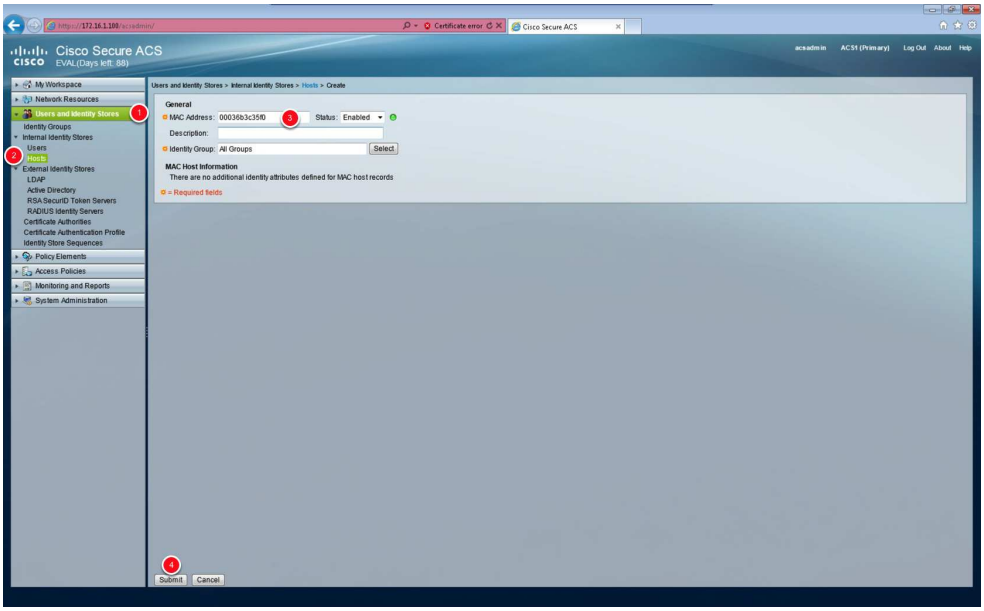
https://t.me/learningnets

With 802.1x enabled, for the switch to allow the IP Phone access to the VOICE VLAN, it needs to receive the **device-traffic-class=voice** VSA form the RADIUS server; to take it into consideration the **aaa authorization network default** command is required, as for any other authorization received from the RADIUS server.

On ACS, first add SW1 as a NAD, and enable it for RADIUS communication:

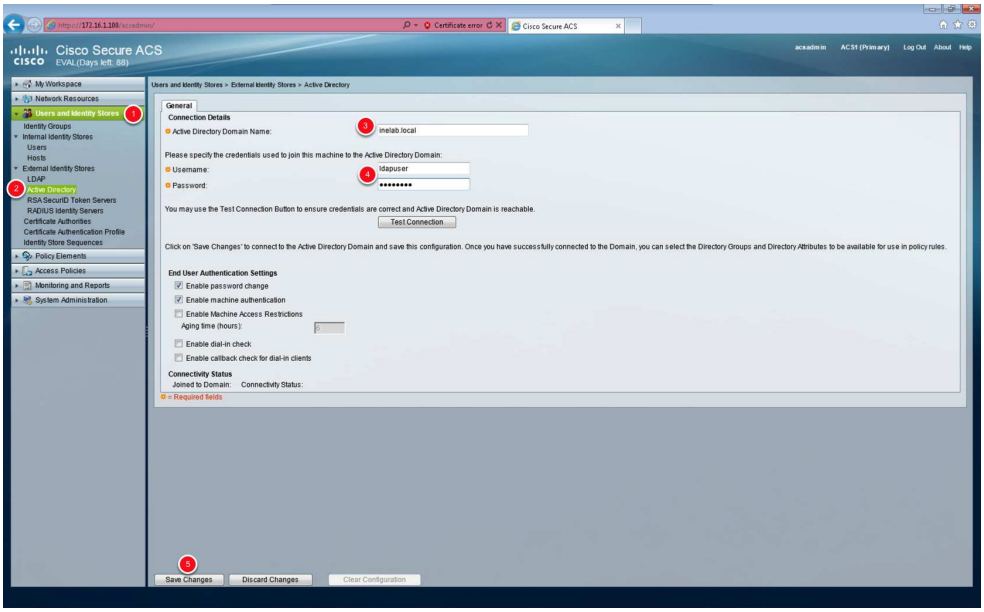


Add the MAC address of the IP Phone in the Internal Hosts Store to be used for MAB; this may differ based on the rack being assigned:

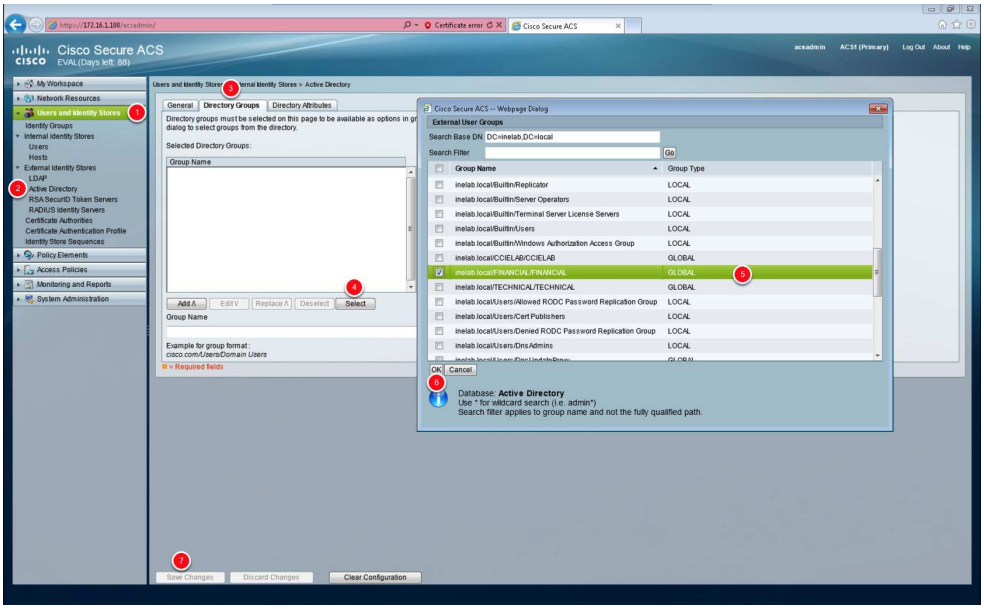


Join ACS into AD using the given credentials:

https://t.me/learningnets

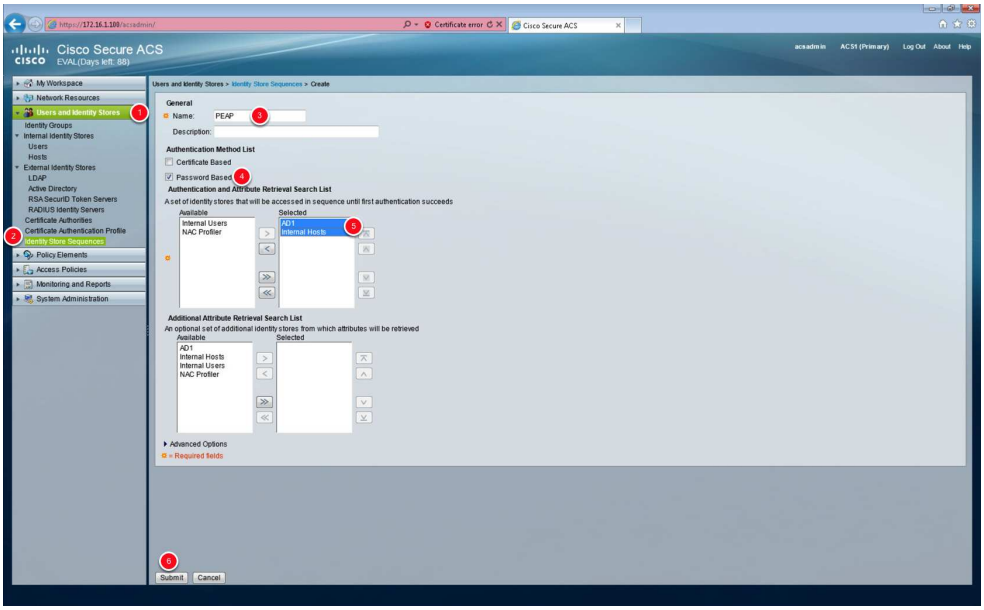


Add the **FINANCIAL** user group so that it can be referenced into authorization policies:

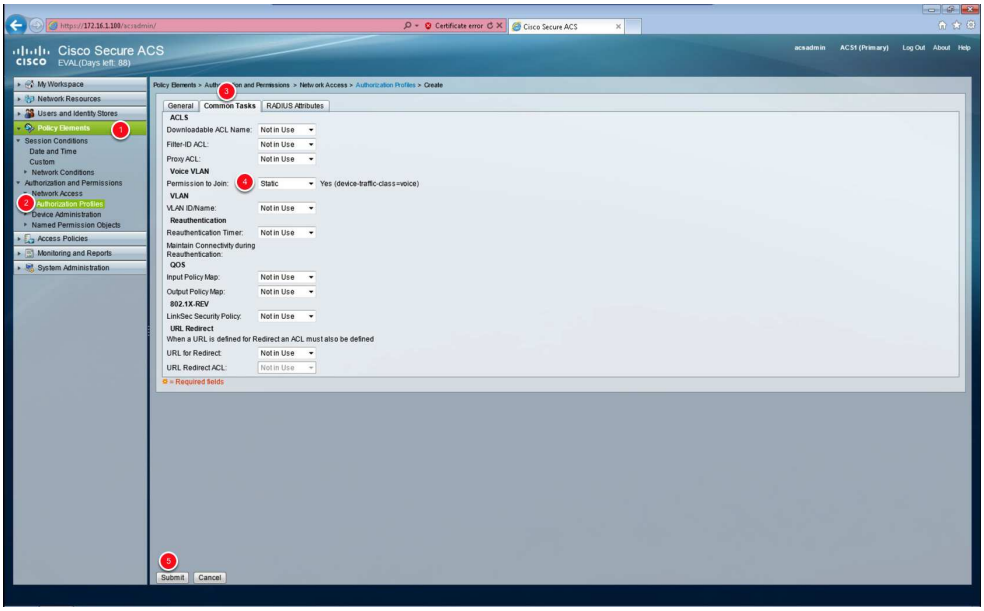


Add a new **Identity Store Sequence** to contain both AD and the Internal Hosts Store:

https://t.me/learningnets

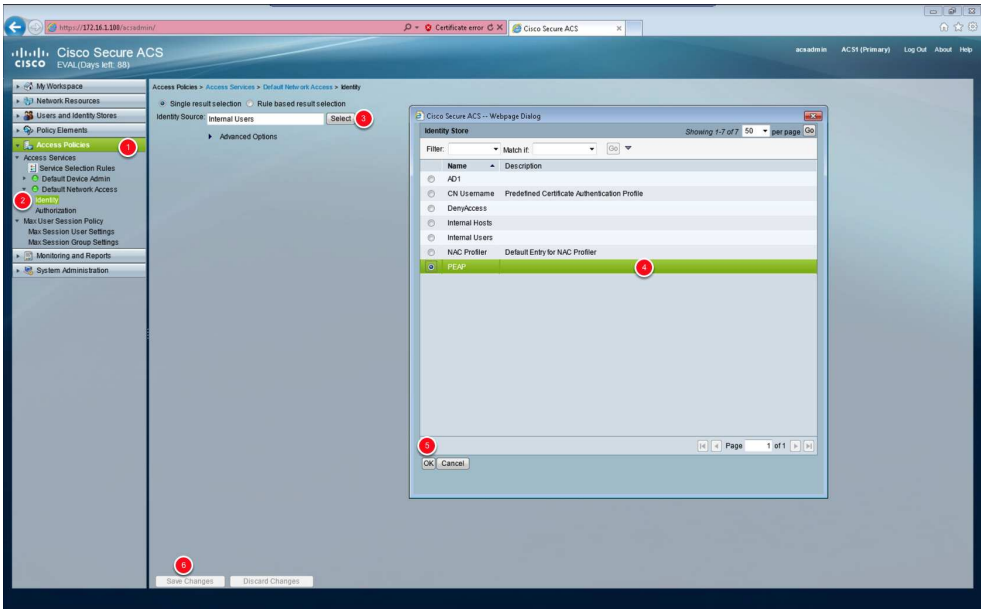


Configure an authorization policy for the voice vlan:

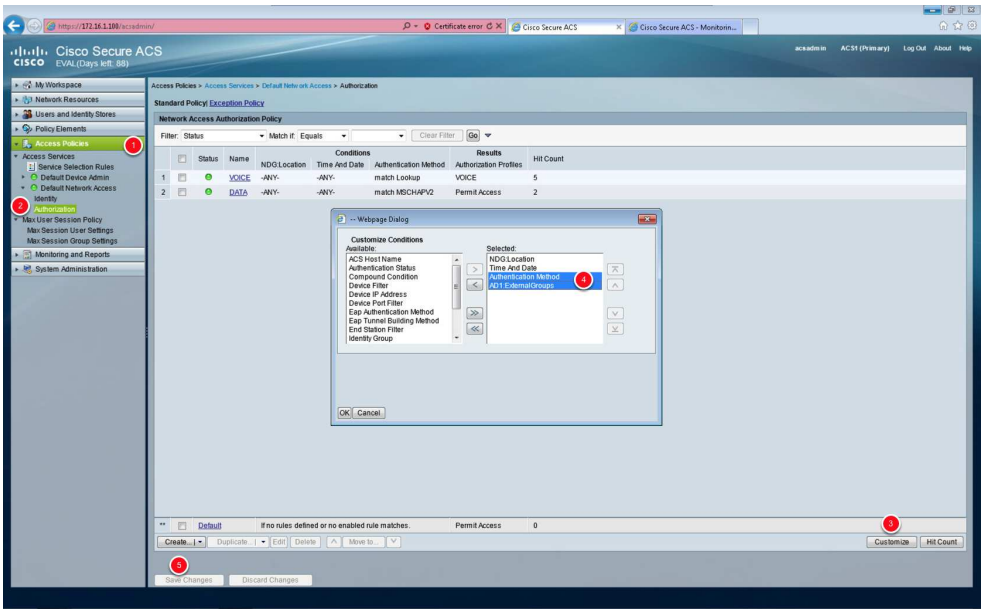


Configure the created **Identity Store Sequence** to be used by the RADIUS authorization policies:

https://t.me/learningnets

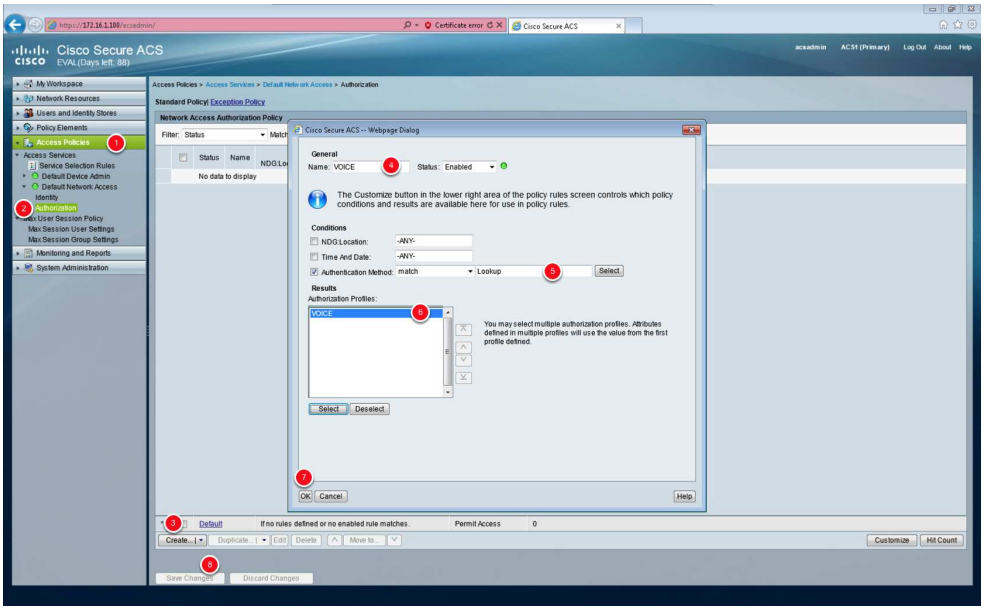


Add required conditions for the RADIUS authorization policies:

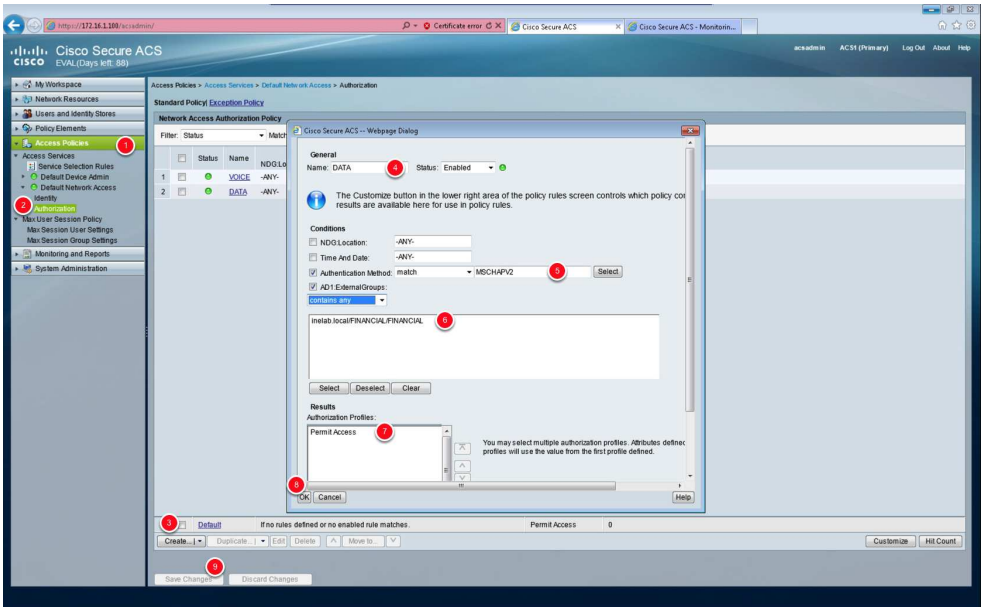


Configure the authorization policy for the VOICE domain:

https://t.me/learningnets



Configure the authorization policy for the DATA domain:

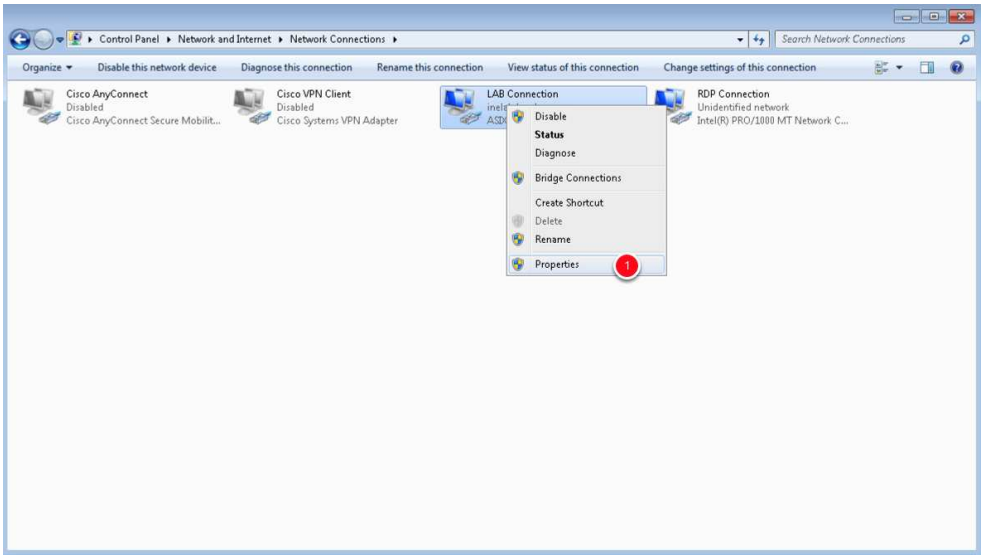


If both Windows native supplicant and AnyConnect supplicant are active, AnyConnect has precedence.

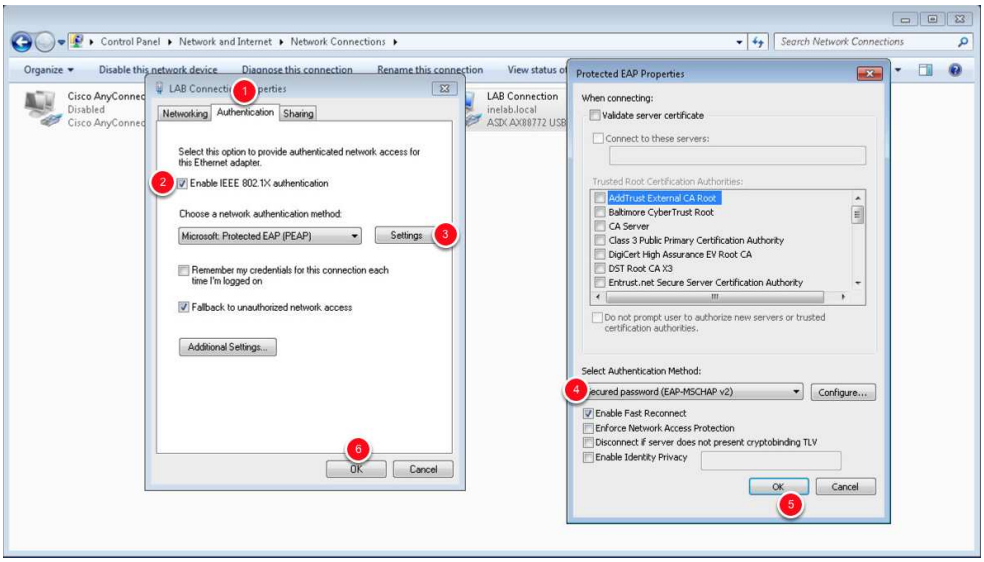
Option 1 Windows Native Supplicant

Configure the native Windows supplicant for the **LAB** Ethernet interface:

https://t.me/learningnets

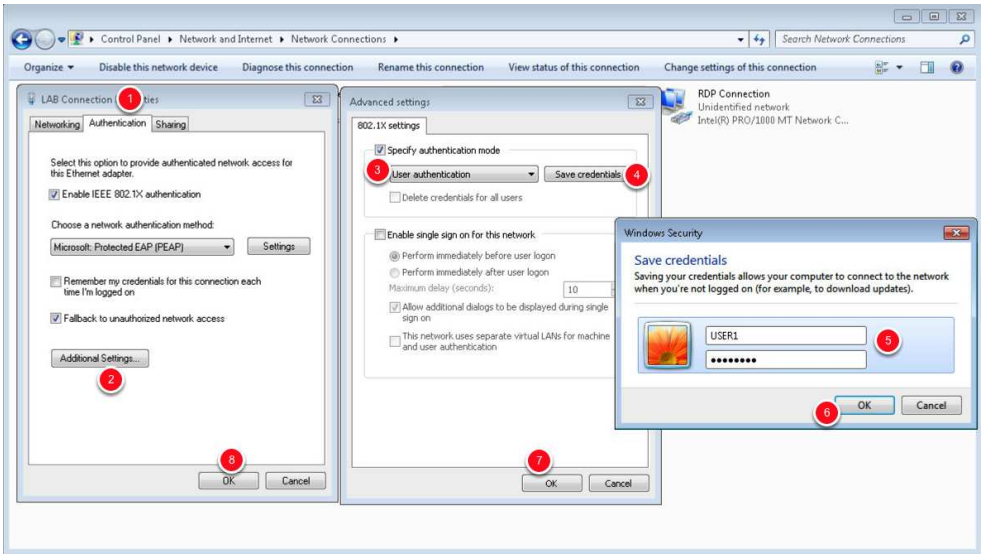


Enable 802.1x for PEAP(EAP-MSCHAPv2). If the **Authentication** tab is not available, ensure that the **Wired AutoConfig** service is started:



Preconfigure the username/password credentials to be used so that Windows does not prompt you for it; if the client machine had been integrated into AD as well, this step would not be needed because the Windows login credentials would have been used.

https://t.me/learningnets



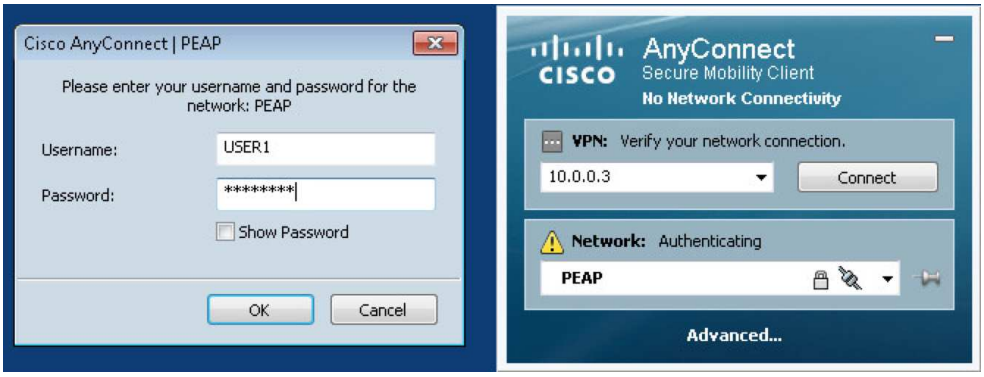
Option 2 AnyConnect Supplicant

Enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConnect 802.1x supplicant service. To access the Windows services, type **services.msc** from command prompt or Run. Open AnyConnect and navigate to the **Advanced** tab to configure a 802.1x profile:



Open AnyConnect and apply the configured profile; you'll be asked for the username and password when the EAP Identity Request is received from the switch. Also, you'll be asked to trust the ACS self-signed certificate.

https://t.me/learningnets



Verification

Verify that both domains are successfully authorized, IP Phone using MAB and PC using 802.1x:

```
SW1#show authentication sessions interface fastEthernet1/0/5
        Interface:  FastEthernet1/0/5
      MAC Address:  48f8.b32e.2423
       IP Address:  Unknown User-Name:  USER1
Status:  Authz Success
Domain:  DATA
  Security Policy:  Should Secure
  Security Status:  Unsecure
   Oper host mode:  multi-domain
 Oper control dir:  both Authorized By:  Authentication Server
       Vlan Group:  N/A
  Session timeout:  N/A
     Idle timeout:  N/A
Common Session ID:  000000000000000A0144CE46
  Acct Session ID:  0x0000000D
           Handle:  0xC500000A

Runnable methods list:
       Method   State
       mab      Failed over dot1x    Authc Success
```

```
        Interface:  FastEthernet1/0/5
      MAC Address:  0003.6b3c.35f0
       IP Address:  Unknown User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  VOICE
  Security Policy:  Should Secure
  Security Status:  Unsecure
```

https://t.me/learningnets

```
   Oper host mode:  multi-domain
 Oper control dir:  both Authorized By:  Authentication Server
  Session timeout:  N/A
     Idle timeout:  N/A
Common Session ID:  000000000000000B0144E807
  Acct Session ID:  0x0000000E
           Handle:  0xC100000B

Runnable methods list:
       Method   State mab       Authc Success

       dot1x    Not run
```

Verify dot1x configuration:

```
SW1#show dot1x interface fastEthernet1/0/5
Dot1x Info for FastEthernet1/0/5
----------------------------------
PAE                       = AUTHENTICATOR
PortControl               = AUTO
ControlDirection          = Both HostMode                  = MULTI_DOMAIN
QuietPeriod               = 60
ServerTimeout             = 0
SuppTimeout               = 30
ReAuthMax                 = 2
MaxReq                    = 2 TxPeriod                  = 20
```

Verify the list of authorized clients:

```
SW1#show authentication interface fastEthernet1/0/5

Client list:
Interface  MAC Address     Method   Domain   Status         Session ID
Fa1/0/5    48f8.b32e.2423  dot1x    DATA     Authz Success  000000000000000A0144CE46
 Fa1/0/5    0003.6b3c.35f0  mab      VOICE    Authz Success  000000000000000B0144E807


Available methods list:
  Handle  Priority  Name
    3        0      dot1x
    2        1      mab
Runnable methods list:
  Handle  Priority  Name
    2        0      mab
```

https://t.me/learningnets

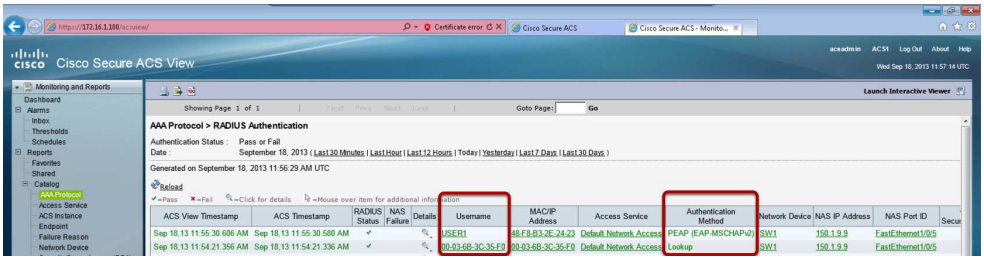
```
    3        1       dot1x
```

When MAC addresses are authorized, it will appear as static on the MAC database:

```
SW1#show mac address-table interface fastEthernet1/0/5
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----   27    48f8.b32e.2423 STATIC
      Fa1/0/5    47    0003.6b3c.35f0 STATIC
      Fa1/0/5
Total Mac Addresses for this criterion: 2
```

Verify the RADIUS Authentication logs from the ACS server:



To test that the switch authorizes both domains in case the RADIUS server is not reachable (critical VLAN), configure an inexistent RADIUS server and flap the port to trigger the authentication process again:

```
SW1:

aaa group server radius ACS
 no server-private 172.16.1.100 auth-port 1645 acct-port 1646 key CISCO
 server-private 172.16.1.101 auth-port 1645 acct-port 1646 key CISCO
!
interface FastEthernet1/0/5
 shutdown
 no shutdown
```

Messages similar to the following should be logged by the switch indicating the expected behavior:

https://t.me/learningnets

%AUTHMGR-7-RESULT: Authentication result 'server dead' from 'mab' for client (48f8.b32e.2423) on Interface Fa1/0/5 AuditSessionID 00000000000001501951E08
%AUTHMGR-7-RESULT: Authentication result 'server dead' from 'mab' for client (0003.6b3c.35f0) on Interface Fa1/0/5 AuditSessionID 00000000000000160195378E

Verify the authorization policy:

```
SW1#show authentication sessions interface fastEthernet1/0/5

            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  Unknown
            User-Name:  48f8b32e2423 Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-host
     Oper control dir:  both Authorized By:  Critical Auth
Vlan Policy:  27
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000001501951E08
      Acct Session ID:  0x00000018
               Handle:  0xD3000015

Runnable methods list:
       Method   State mab        Authc Failed
dot1x      Not run

Critical Authorization is in effect for domain(s) DATA and VOICE
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

PEAP (EAP-MSCHAPv2) Using ISE and AD

Task

- Configure **TEST-PC-A** for 802.1x authentication using PEAP; authenticate with **USER1** and password **1qaz@WSX** from **inelab.local** domain.
- For joining ISE into **inelab.local** domain, use **ldapuser** with a password of **1qaz@WSX**.
- Your configuration should take into consideration that multiple PCs may be attached to the IP Phone through a switch, but each one should be authenticated.
- Your PC authorization rule from ISE must include the **sAMAccountName** attribute from Active Directory.
- Configure ISE so that the IP Phone matches the default authorization policy of **Profiled Cisco IP Phones**.
- Ensure that your configuration authorizes the IP Phone as fast as possible.
- Configure SW1 to log on successful dACL download and enforcement.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Cisco ISE can integrate with external identity sources to validate credentials from various authentication requests, and to retrieve group information as well as other attributes that are associated with the user, and uses them in authorization policies.

https://t.me/learningnets

Integration with a Certificate Authority is also seen as an external identity source. For ISE to join an Active Directory domain, with the exception of IP connectivity with the Domain Controller, there are three additional requirements:

- The clocks of ISE and DC should be synchronized.
- The domain name configured in ISE CLI should match the joining domain.
- The DNS server configured on ISE CLI should be the IP address of the Domain Controller.

At a minimum, because there might be a firewall between ISE and DC, the following services must be allowed:

Protocol	Port
LDAP	UDP/389
SMB1	TCP/445
KDC2	TCP/88
Global Catalog	TCP/3268 , 3269
KPASS	TCP/464
NTP	UDP/123
LDAP	TCP/389
LDAPS	TCP/636

There are two versions of PEAP, v0(Windows XP) and v1 (Windows 7). One major difference between these versions is that v1 supports the identity protection functionality, where the username or CN from the certificate are not sent before the secure tunnel is established between the supplicant and the authentication server. PEAP can use EAP-MSCHAPv2, EAP-GTC, and EAP-TLS as the inner

https://t.me/learningnets

authentication methods; this task uses EAP-MSCHAPv2 because a username/password authentication is requested. PEAP requires a certificate only on the server (ISE) side to build the secure tunnel. In real deployments, the client will be configured to validate the server's certificate, in which case ISE needs a certificate issued from an external CA (not the self-signed one) and the client will need to have the certificate of the CA imported and configured as trusted root.

Configuration

SW1:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
dot1x system-auth-control
!
interface FastEthernet1/0/5
 authentication host-mode multi-auth
 authentication port-control auto
 authentication order mab dot1x
 dot1x pae authenticator
 mab
!
radius-server vsa send authentication
epm logging
ip device tracking
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Loopback0
```

ISE is configured to synchronize time with the NTP server at R1. To join ISE into Active Directory, the time difference between ISE and the Domain Controller should be within 5 minutes. Because the Domain Controller (TEST-SRV-A) cannot be synchronized with R1, make sure to manually set the time on R1 or directly into ISE, if required, to match the time from TEST-SRV-A. Then wait until ISE synchronizes its time again with R1.

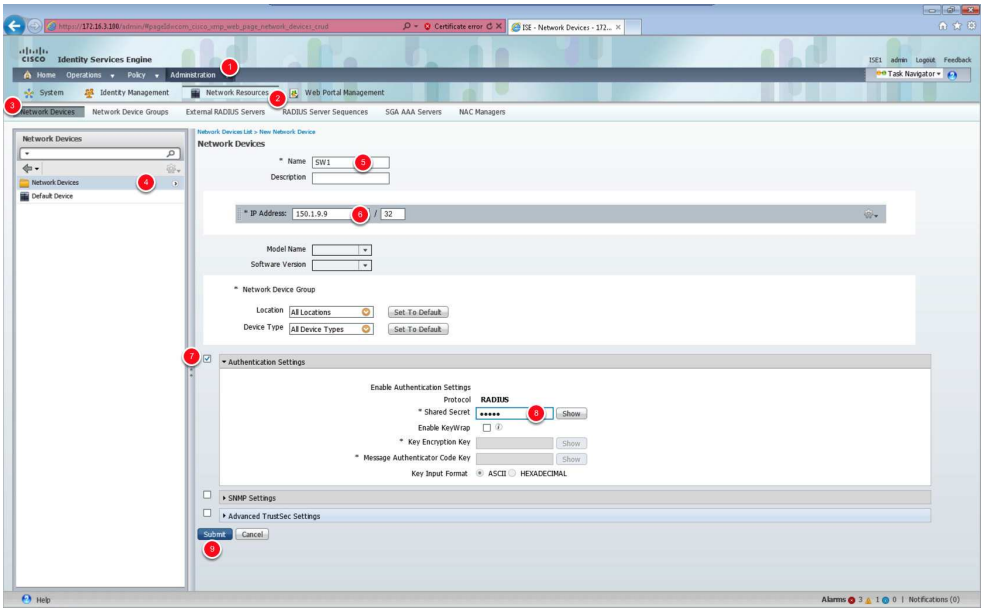
With 802.1x enabled, in order for the switch to allow the IP Phone access to the VOICE VLAN, it must receive the **device-traffic-class=voice** VSA from the RADIUS server; to take it into consideration, the `aaa authorization network default` command is required, as for any other authorization received from the RADIUS

https://t.me/learningnets

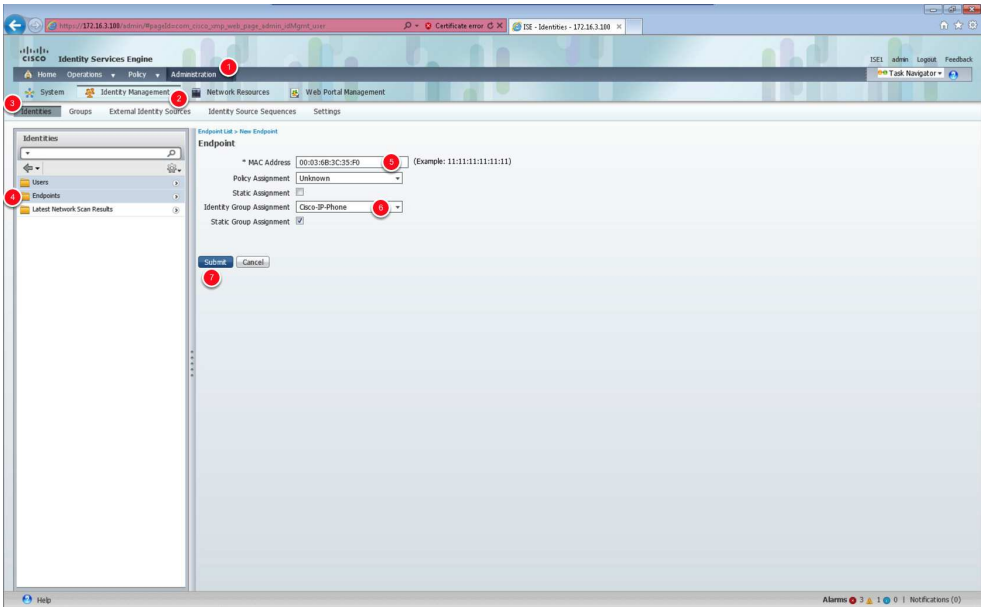
server.

The command `radius-server vsa send authentication` is required so that upon receiving the name of the dACL, the switch can request to download the access-list content. Both authorization profiles used have a dACL. The command `ip device tracking` is required so that the switch can actually apply the dACL on the port after it first learns the connecting device IP address from ARP packets, which is used to substitute the source IP address from the default of **any** in the dACL.

Add SW1 as a NAD device into ISE:

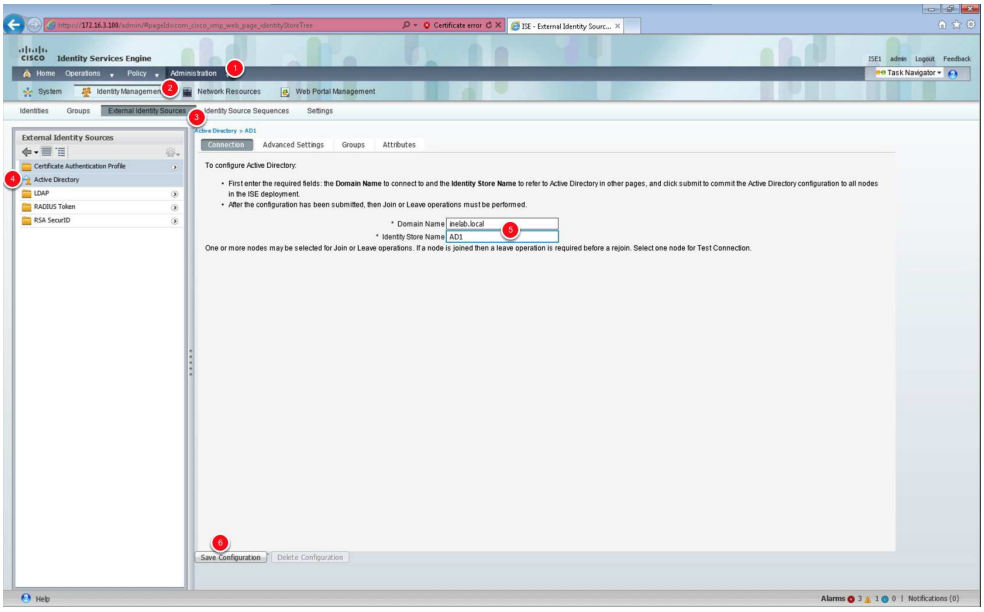


Add the MAC address of the IP Phone in the **Cisco-IP-Phone** Identity Group to be used by MAB, and match the default authorization rule for IP Phones. This rule automatically authorizes the voice VLAN:

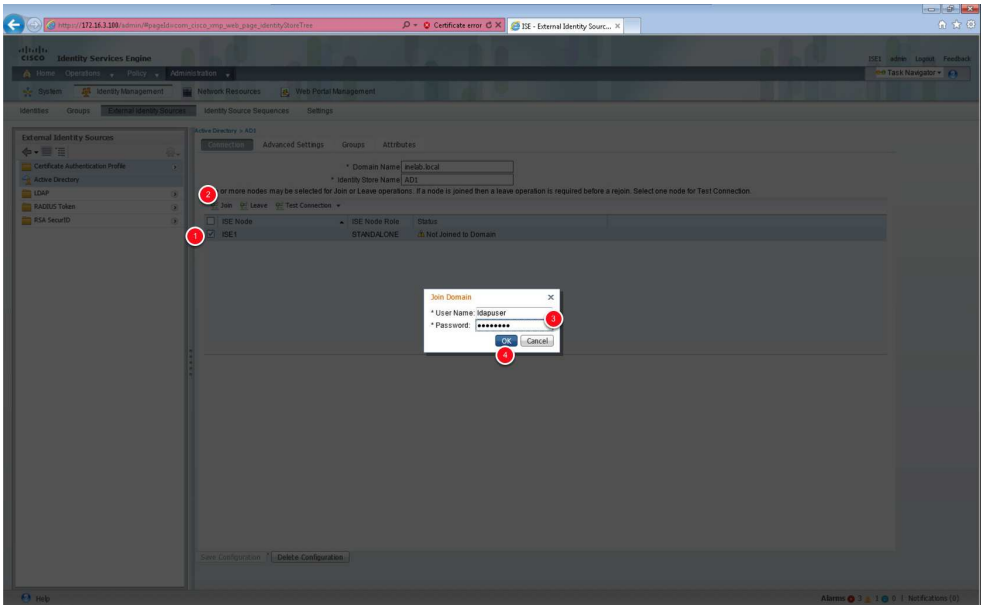


https://t.me/learningnets

Configure the Active Directory domain to be joined:

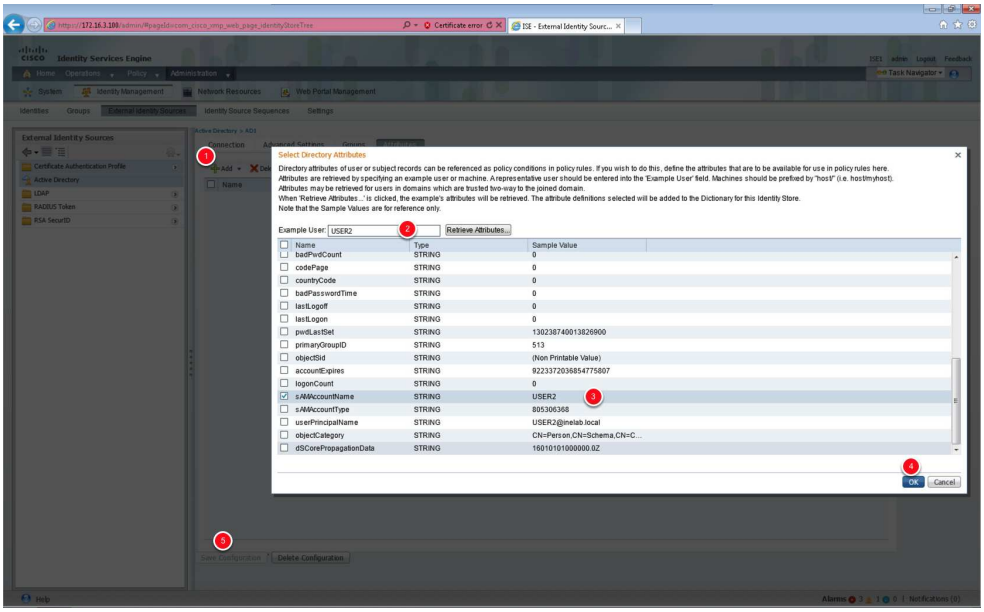


Join ISE into Active Directory using the given credentials:

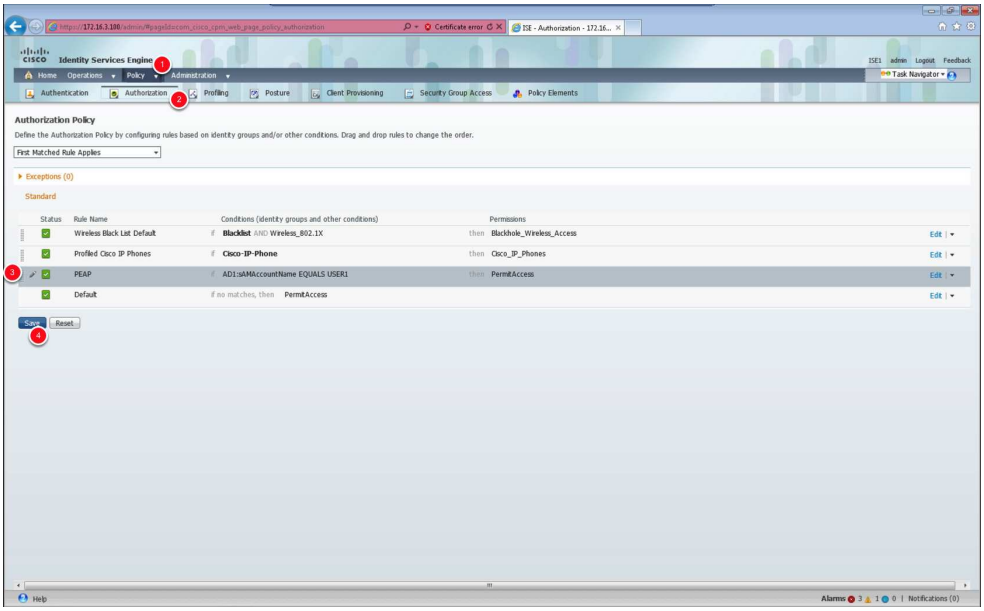


Add the **sAMAccountName** to be used as a condition in the authorization policies:

https://t.me/learningnets

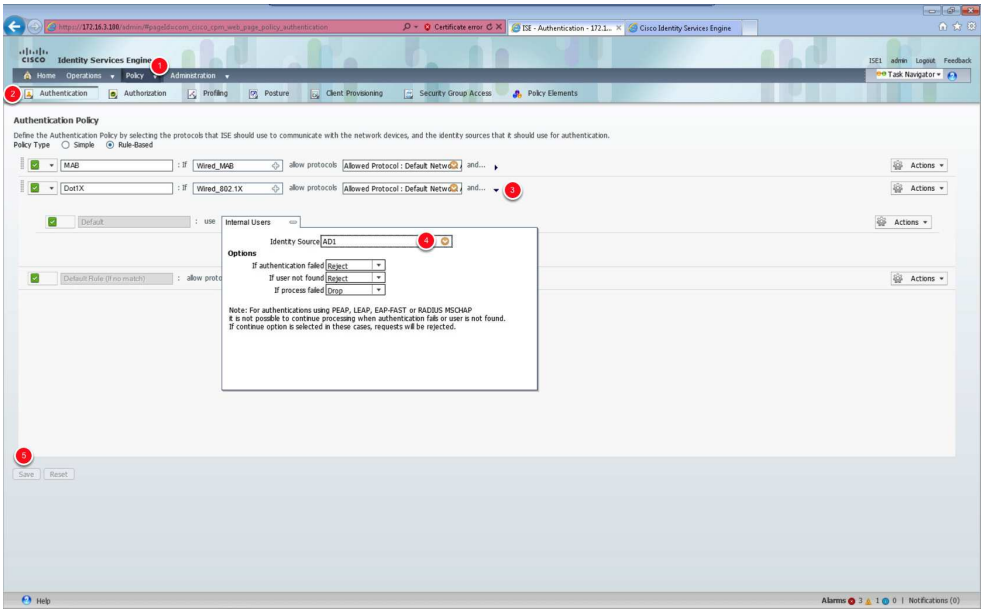


Default authentication rules will permit successful MAB authentication and authorization. Configure the authorization policy for TEST-PC-A:



Configure the default 802.1x authentication policy to use **inelab.local** AD for user authentication:

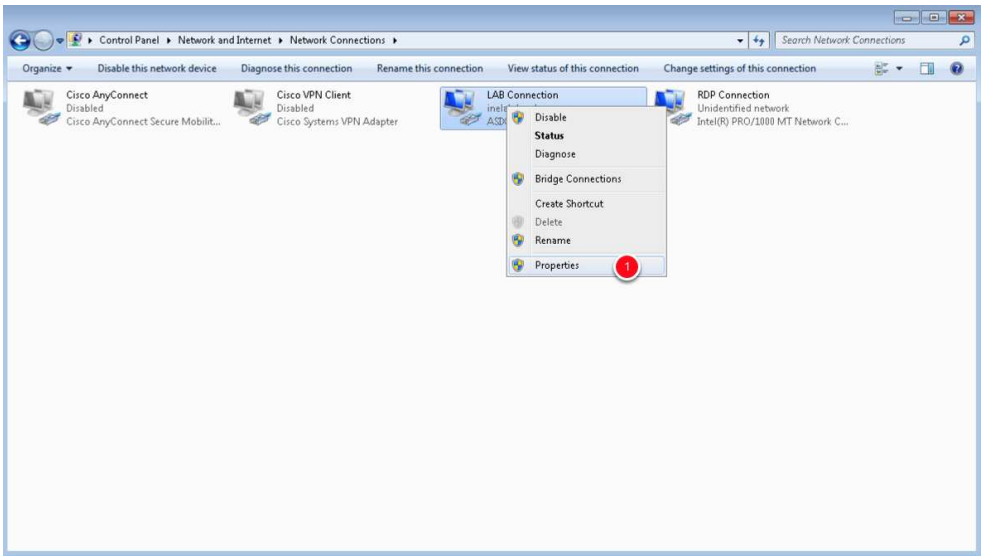
https://t.me/learningnets



If both Windows native supplicant and AnyConnect supplicant are active, AnyConnect has precedence.

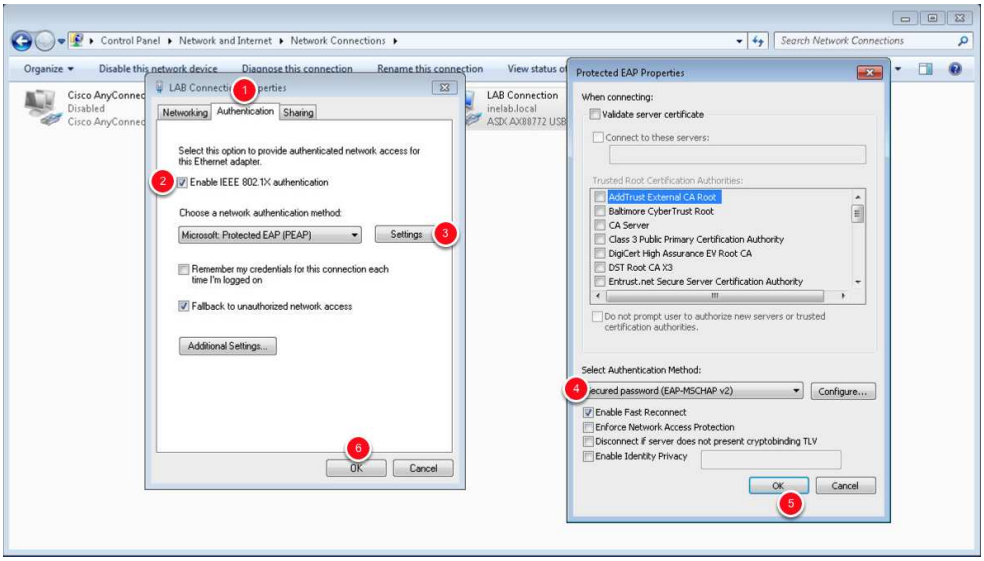
Option 1 Windows Native Supplicant

Configure the native Windows supplicant for the **LAB** Ethernet interface:

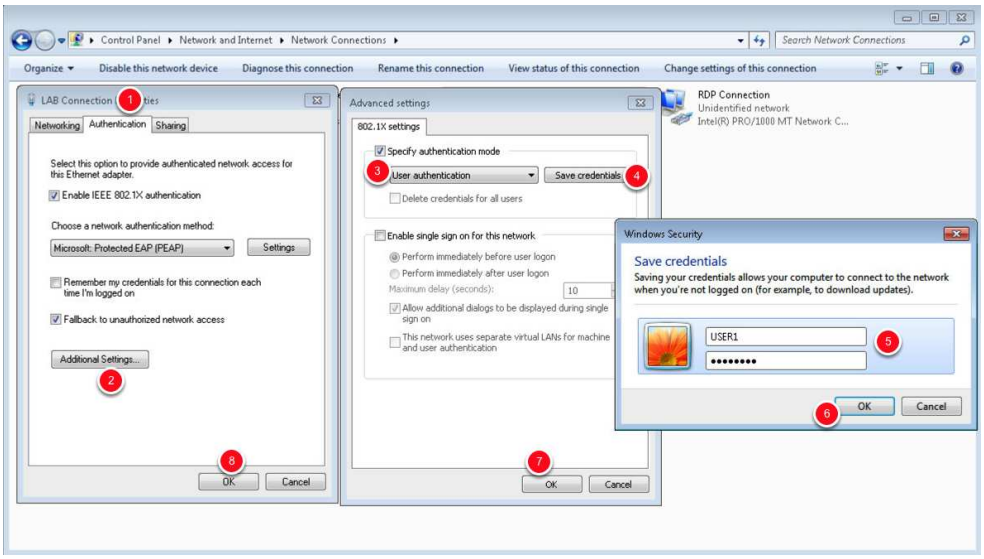


Enable 802.1x for PEAP(EAP-MSCHAPv2); if the **Authentication** tab is not available, ensure that the **Wired AutoConfig** service is started:

https://t.me/learningnets



Preconfigure the username/password credentials to be used so that Windows does not prompt you for it; if the client machine had been integrated into AD as well, this step would not be needed because the Windows login credentials would have been used:



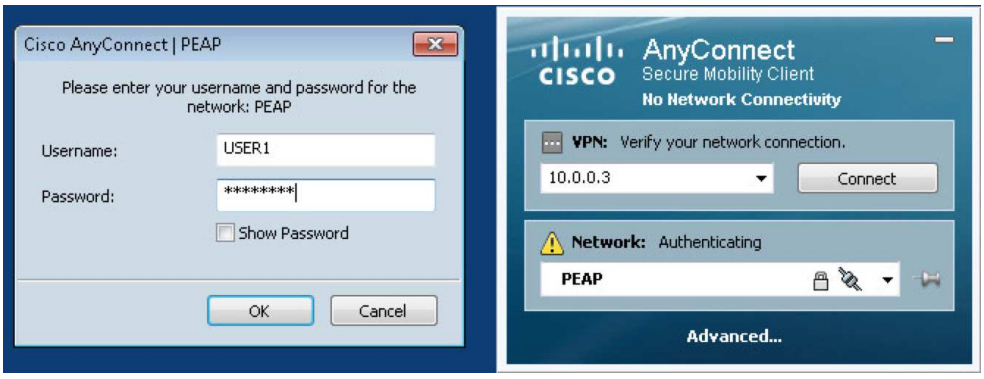
Option 2 AnyConnect Supplicant

First, enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConnect 802.1x supplicant service. To access the Windows services, type **services.msc** from the command prompt or Run. Open AnyConnect and navigate to the **Advanced** tab to configure an 802.1x profile:

https://t.me/learningnets



Open AnyConnect and apply the configured profile; you'll be asked for the username and password when the EAP Identity Request is received from the switch. Also, you'll be asked to trust the ISE self-signed certificate.



Verification

Verify that both the IP Phone and the TEST-PC-A are authorized:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423 IP Address:  136.1.27.51
 User-Name:  USER1
 Status:  Authz Success
 Domain:  DATA
       Security Policy:  Should Secure
       Security Status:  Unsecure Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
          Vlan Group:  N/A
       Session timeout:  N/A
          Idle timeout:  N/A
```

https://t.me/learningnets

```
    Common Session ID:  000000000000000E01046D8B
      Acct Session ID:  0x00000014
              Handle:  0x8400000E

Runnable methods list:
       Method   State
       mab      Failed over dot1x     Authc Success



----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  136.1.47.51 User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  VOICE
      Security Policy:  Should Secure
      Security Status:  Unsecure Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
ACS ACL:  xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000000F0104875D
      Acct Session ID:  0x00000013
              Handle:  0x3D00000F

Runnable methods list:
       Method   State mab       Authc Success

       dot1x    Not run
```

Verify the list of authorized clients:

```
SW1#show authentication sessions


 Interface  MAC Address     Method   Domain   Status         Session ID
 Fa1/0/5    48f8.b32e.2423  dot1x    DATA     Authz Success  000000000000000E01046D8B
 Fa1/0/5    0003.6b3c.35f0  mab      VOICE    Authz Success  000000000000000F0104875D
```

When MAC addresses are authorized, it will appear as static on the MAC database:

```
SW1#show mac address-table interface fastEthernet1/0/5
          Mac Address Table
-------------------------------------------
```

https://t.me/learningnets

```
Vlan    Mac Address       Type        Ports
----    -----------       --------    -----   27    48f8.b32e.2423 STATIC
      Fa1/0/5    47    0003.6b3c.35f0 STATIC
      Fa1/0/5
Total Mac Addresses for this criterion: 2
```

Verify the downloaded ACL from ISE, and that the source IP address is replaced by IP device tracking:

```
SW1#show ip access-lists xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2

Extended IP access list xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2 (per-user) 10 permit ip any any
!SW1#show ip access-lists interface fastEthernet1/0/5
     permit ip host 136.1.27.51
 any      permit ip host 136.1.47.51
 any
```

Verify the DHCP binding and IP device tracking databases:

https://t.me/learningnets

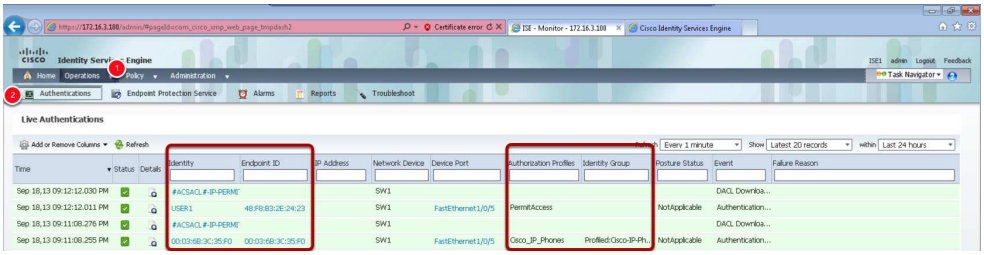
```
R1#show ip dhcp binding

Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name136.1.27.51
        0148.f8b3.2e24.23       Mar 02 1993 04:45 AM    Automatic136.1.47.51
        0100.036b.3c35.f0       Mar 02 1993 04:56 AM    Automatic
!
!SW1#show ip device tracking interface fastEthernet1/0/5

IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
IP Device Tracking Probe Delay Interval = 0
---------------------------------------------------------------------------
  IP Address      MAC Address    Vlan  Interface                STATE
---------------------------------------------------------------------------
136.1.47.51      0003.6b3c.35f0  47    FastEthernet1/0/5         ACTIVE
136.1.27.51      48f8.b32e.2423  27    FastEthernet1/0/5         ACTIVE


Total number interfaces enabled: 1
Enabled interfaces:
  Fa1/0/5
```

Verify the Authentication logs from the ISE server; observe that for both TEST-PC and IP Phone a dACL was received:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

PEAP (EAP-TLS) Using ACS

Task

- Configure **TEST-PC-A** for 802.1x authentication using PEAP; both the supplicant and authentication server should use certificates issued by the CA from **TEST-SRV-A**.
- Authorize no more than one MAC address per domain (VOICE and DATA); otherwise, drop traffic.
- SW1 should first use MAB, but allow **TEST-PC-A** to pass 802.1x authentication even if MAB is successful.
- If the TEST-PC fails authentication or its 802.1x supplicant is not configured, SW1 should still authorize it in VLAN 27.
- SW1 should use a pre-authentication ACL denying ICMP traffic; upon successful authentication, all traffic should be allowed without receiving any ACLs from the ACS server.
- ACS should select the **Identity Source** based on the authentication method; configure two such rules.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

To authorize one MAC address for each domain (VOICE and DATA), the port must run in **multi-domain** mode. The task makes use of the **Guest VLAN** (the client does

https://t.me/learningnets

not have an 802.1x supplicant active and configured, so it does not respond to EAP Identity Request messages from the switch) and **Authentication Fail VLAN** (the client has an 802.1x supplicant active and configured, but it fails both authentication and MAB). These two fallback mechanisms are not supported when the port runs in **multi-auth** mode.

If a pre-authentication access-list is configured on the port, hosts will be allowed to pass traffic before authentication as dictated by the access-list. After successful authentication, if no dACL is received from the RADIUS server, the same policy will remain active by default. This behavior can be overridden by using the **Open Directive** functionality via the global command `epm access-control open` :

- If the RADIUS server returns any form of ACL, it will be applied on the port and restrict traffic accordingly.
- If the RADIUS server does not return an ACL, the switch automatically creates a `permit ip host any` entry for the authenticated host, overriding the pre-authentication ACL.

For this functionality to work, `ip device tracking` must be configured as well; otherwise, the switch will not know the IP address of the authenticated host to replace the **host** value with its IP address, and it will fail to apply it. This functionality can optionally be configured on the RADIUS server and received in the authorization process as well, using the AV Pair **authz-directive** with a value of **open**.

By default, the authentication priority follows the authentication order, which by default is MAB, followed by 802.1x. With these default settings, a client must fail MAB to be allowed to authenticate with 802.1x. Change the default priority settings, even if client successfully authenticated via MAB; if the switch receives an EAP Start message from the client, it allows it to authenticate via 802.1x and possibly gain another authorization.

Configuration

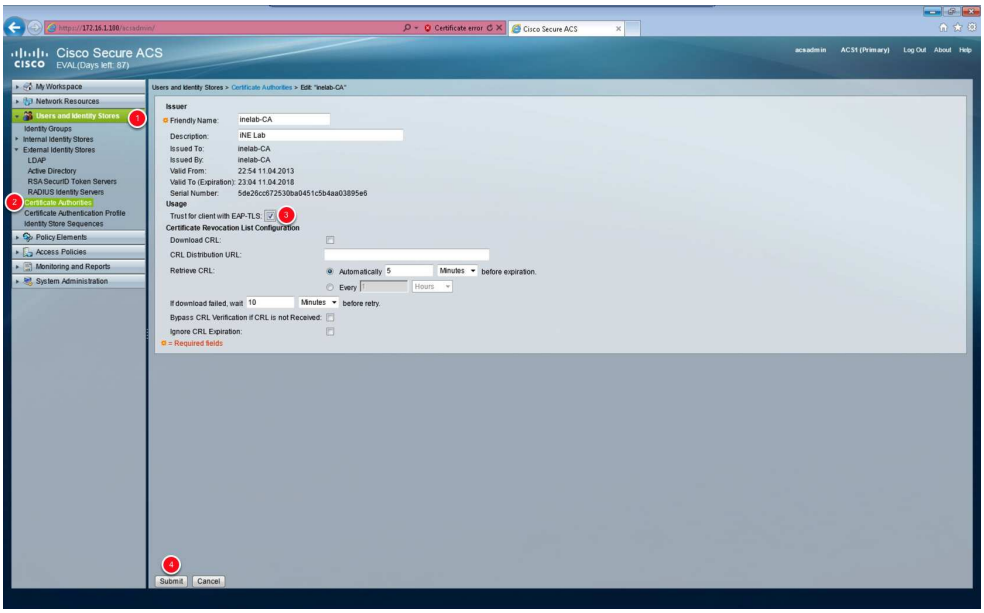
SW1:

```
aaa new-model
aaa authentication dot1x default group ACS
aaa authorization network default group ACS
!
dot1x system-auth-control
epm access-control open
ip device tracking
!
```

https://t.me/learningnets

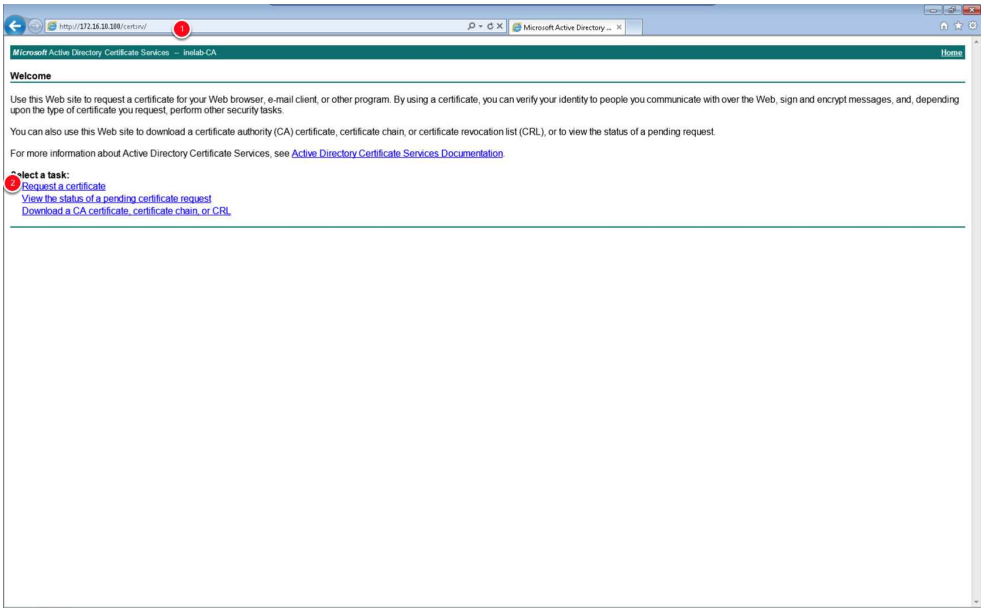
```
ip access-list extended PRE-AUTH
 deny icmp any any
 permit ip any any
!
interface FastEthernet1/0/5
 ip access-group PRE-AUTH in
 authentication host-mode multi-domain
 authentication port-control auto
 authentication order mab dot1x
 authentication priority dot1x mab
 authentication violation protect
 authentication event fail action authorize vlan 27
 authentication event no-response action authorize vlan 27
 dot1x pae authenticator
 mab
!
aaa group server radius ACS
 server-private 172.16.1.100 key CISCO
 ip radius source-interface Loopback0
```

First, request and install a certificate from the CA server for ACS. See the Installing ACS Certificates lab . When installing the certificate, you can specify to use it only for EAP, and leave the self-signed certificate to be used for management via HTTPS. Also make sure to mark the CA certificate as trusted for EAP-TLS, so that client certificates issued by this CA are trusted by ACS and the trust relationship required for EAP-TLS can be established:

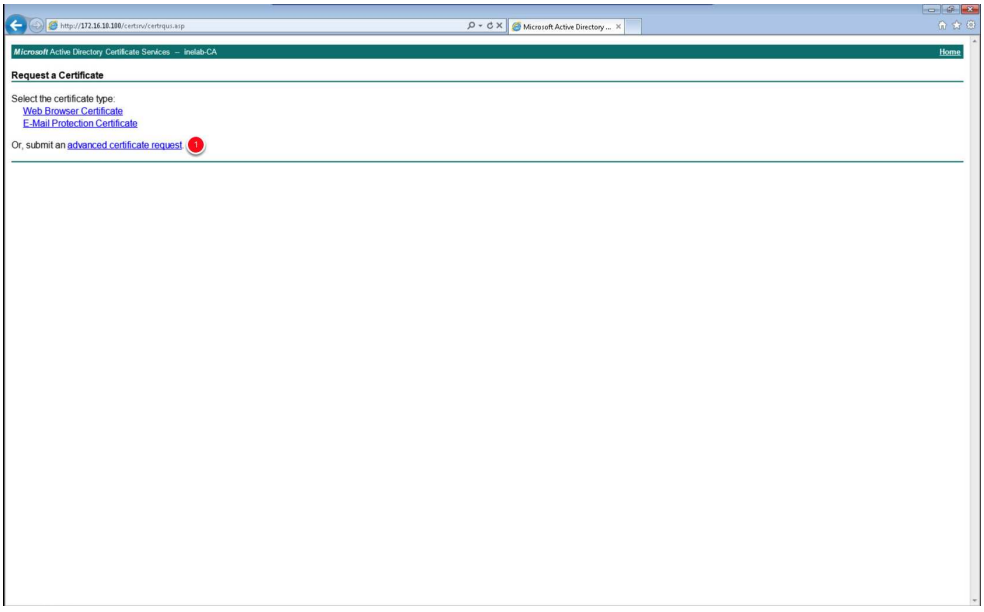


https://t.me/learningnets

Request a client certificate for **TEST-PC-A** from the CA server; begin by browsing to the CA URL of **http://172.16.10.100/certsrv**:

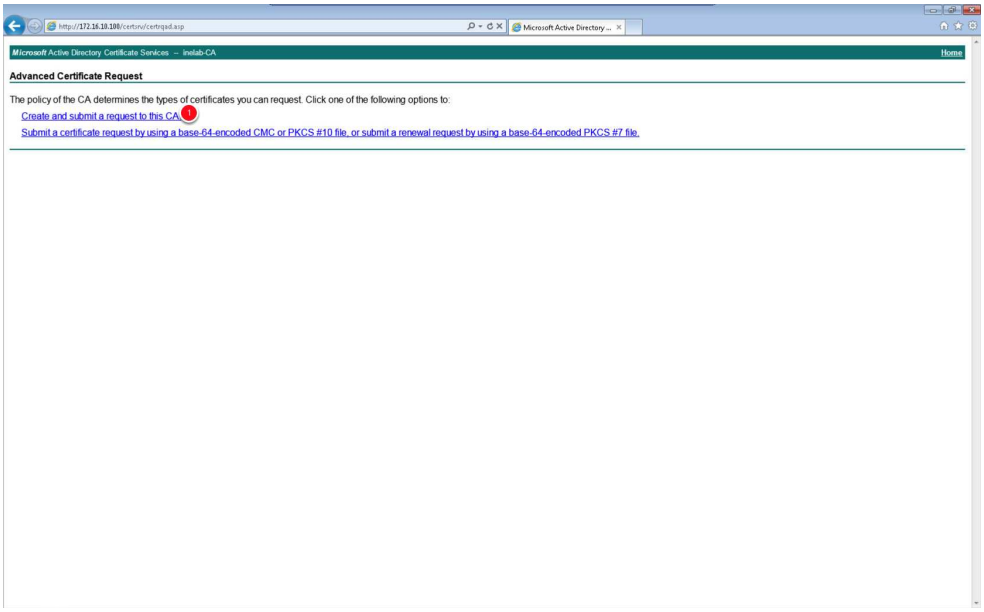


To request a client certificate usable for 802.1x, select the **Advanced Service Request** option; this may differ if you are using your own CA infrastructure:

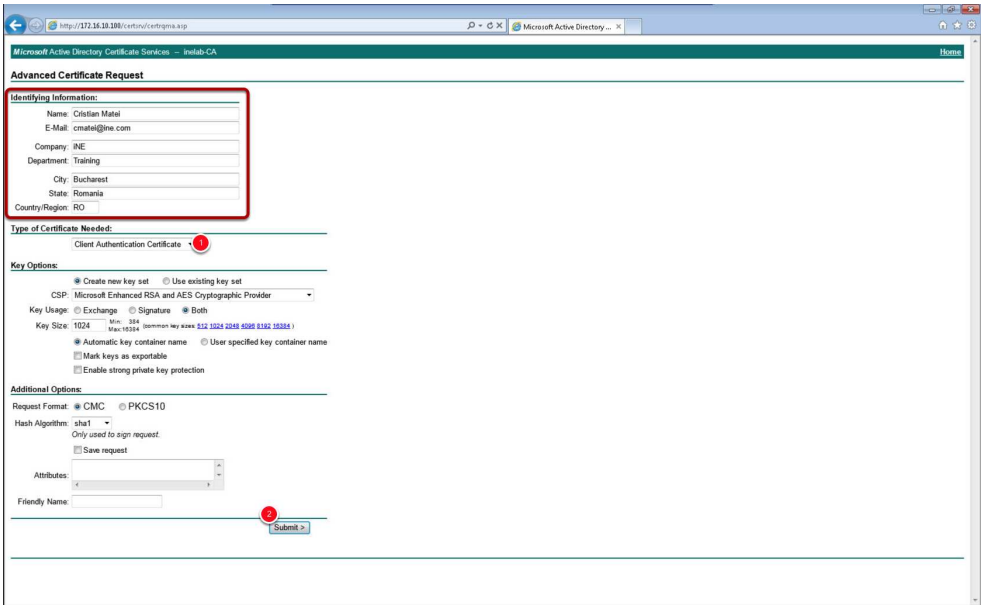


We need to create the CSR data and submit it as well, so select the first option; if presented with a warning, select **Yes**. When enrolling the ACS into CA, the second option was selected because the CSR data was generated by ACS itself:

https://t.me/learningnets

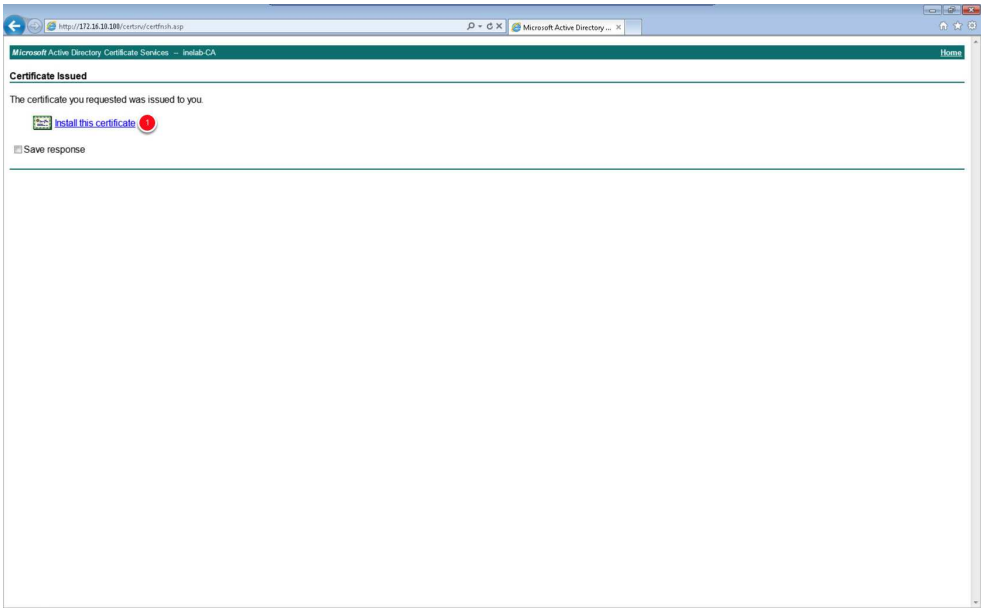


Complete the CSR data with any values because no restrictions are imposed by the task, and click **Submit**:

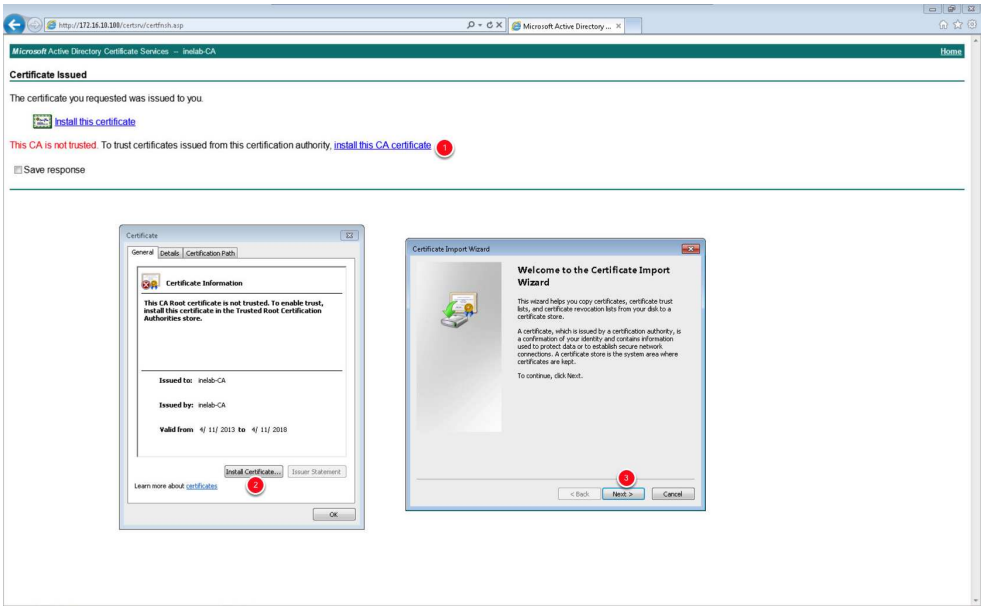


Because the CA is configured to automatically grant certificates, after clocking **Submit** you'll be presented with the certificate; click to install it in the local User Repository:

https://t.me/learningnets

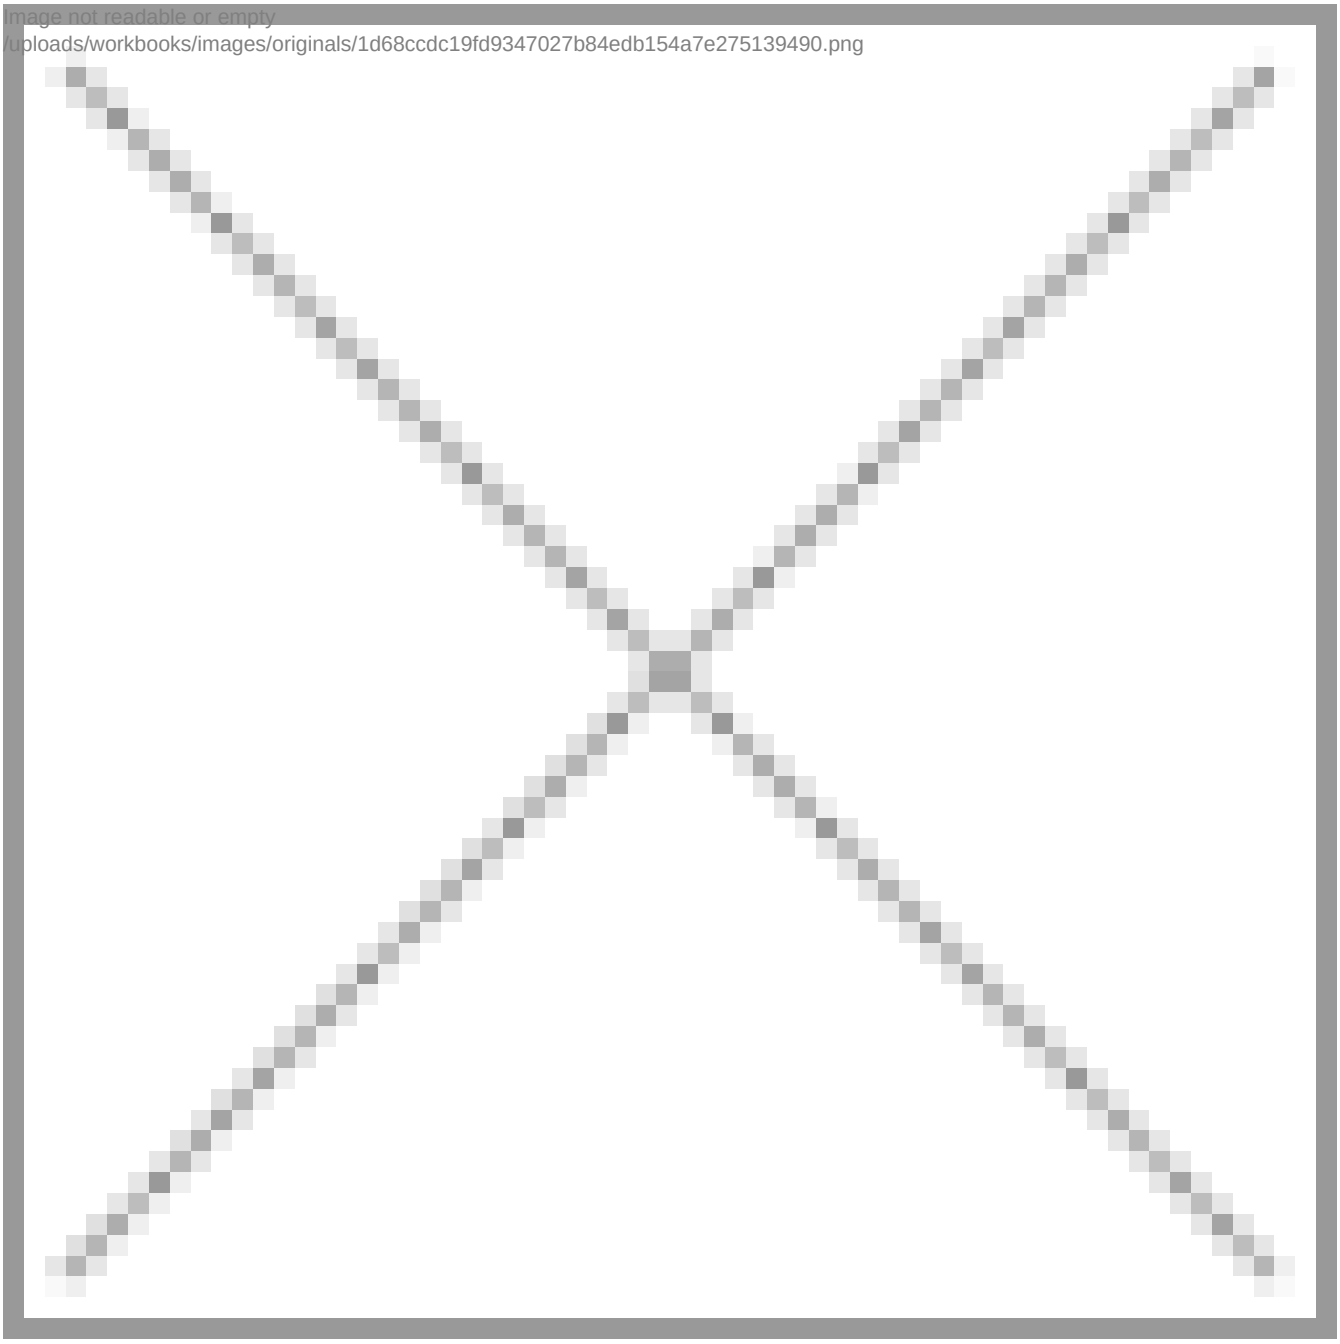


Because the CA is an internal private one, the CA certificate must be added to the **Trusted Root Certification Authorities** before the issued certificate can be installed, so you'll be presented with a screen asking for this:



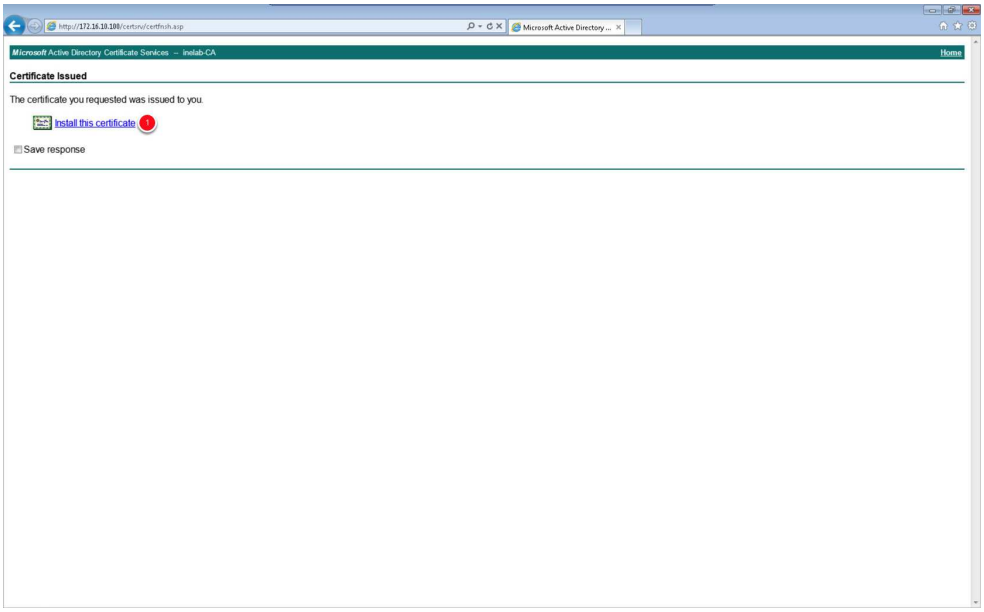
You must manually select installing the CA certificate into the **Trusted Root Certification Authorities**; then click **Next** until the end of the wizard and wait for a successful import message:

https://t.me/learningnets

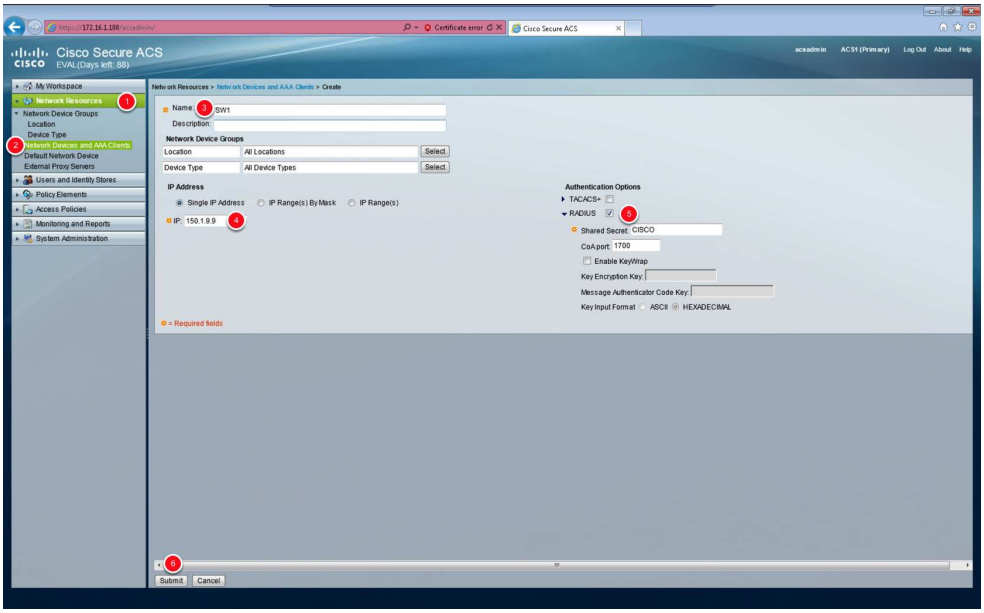


After the CA certificate is installed, click again to install the client certificate:

https://t.me/learningnets

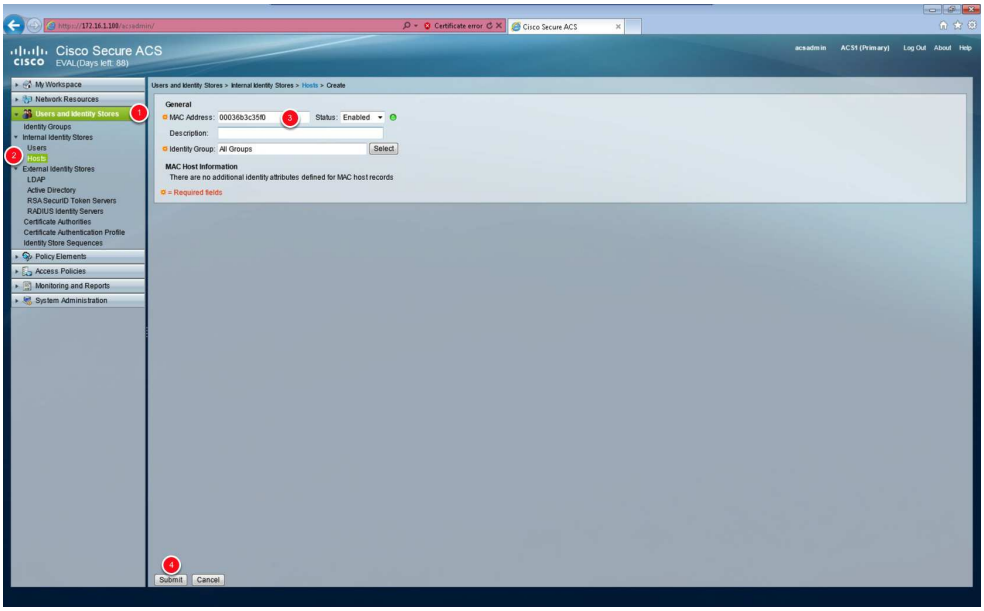


Add SW1 as a NAD into ACS, and configure it for RADIUS communication:

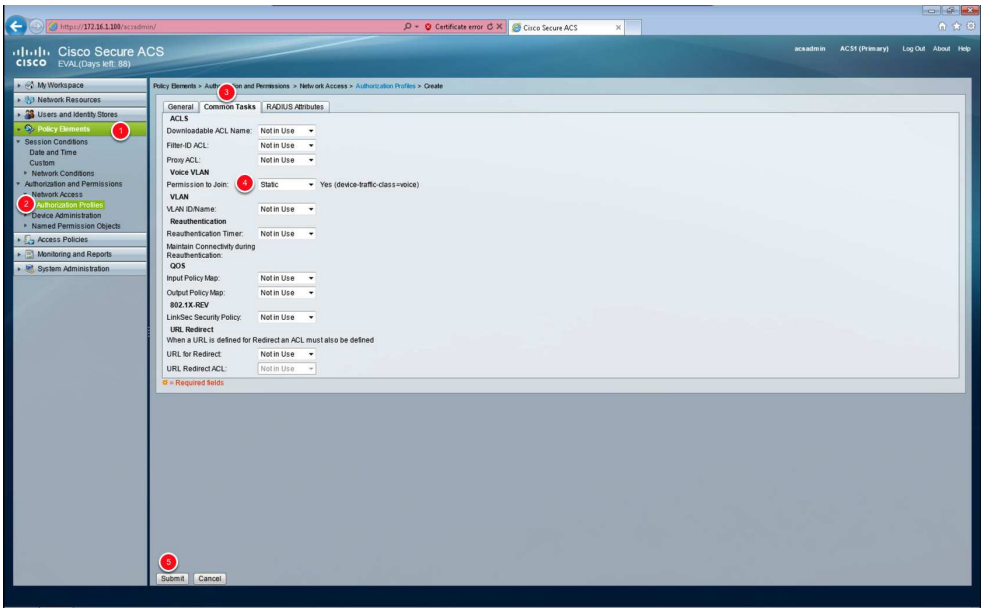


Add the MAC address of the IP Phone in the Internal Hosts Store to be used for MAB; this may differ based on the rack being assigned:

https://t.me/learningnets

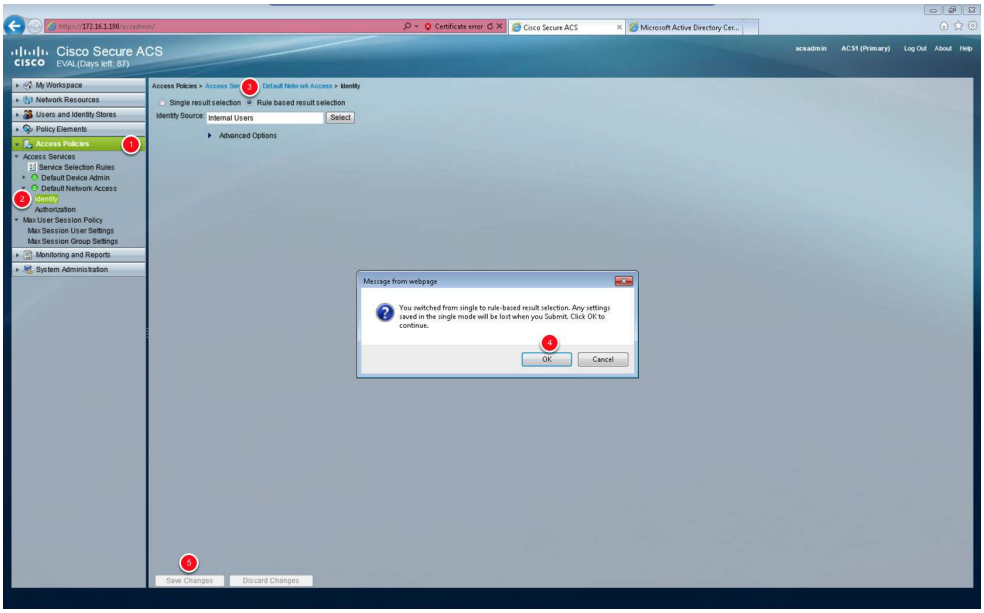


Configure an authorization profile for the voice vlan:

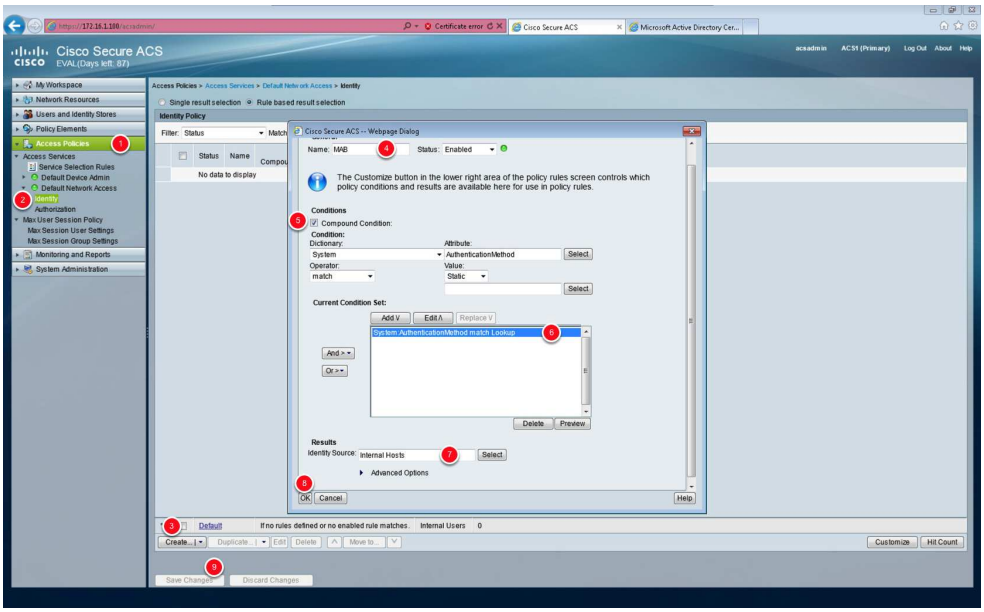


Configure the Identity Source selection to be rule based:

https://t.me/learningnets

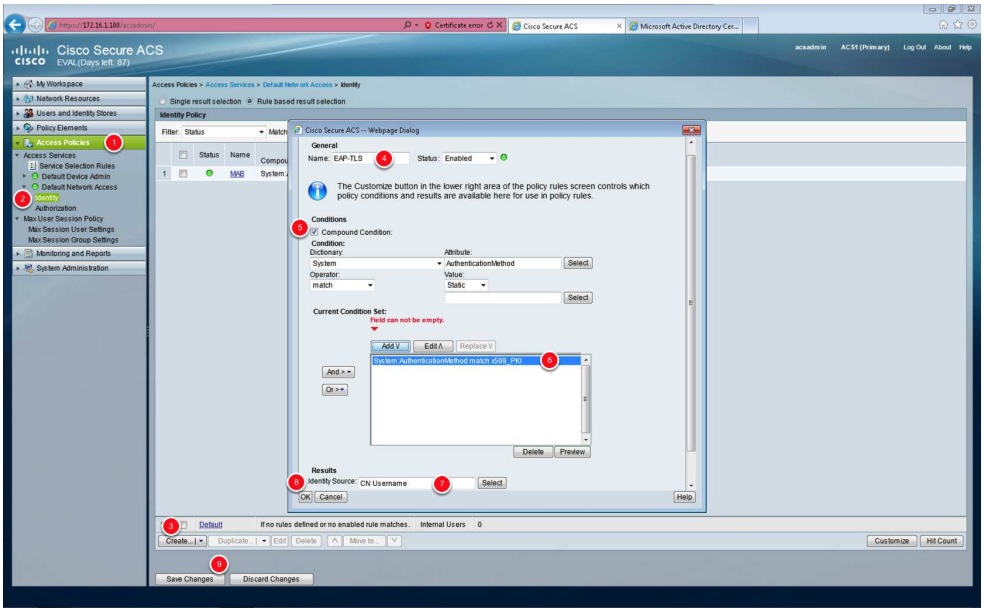


Configure the Identity Source selection for MAB:

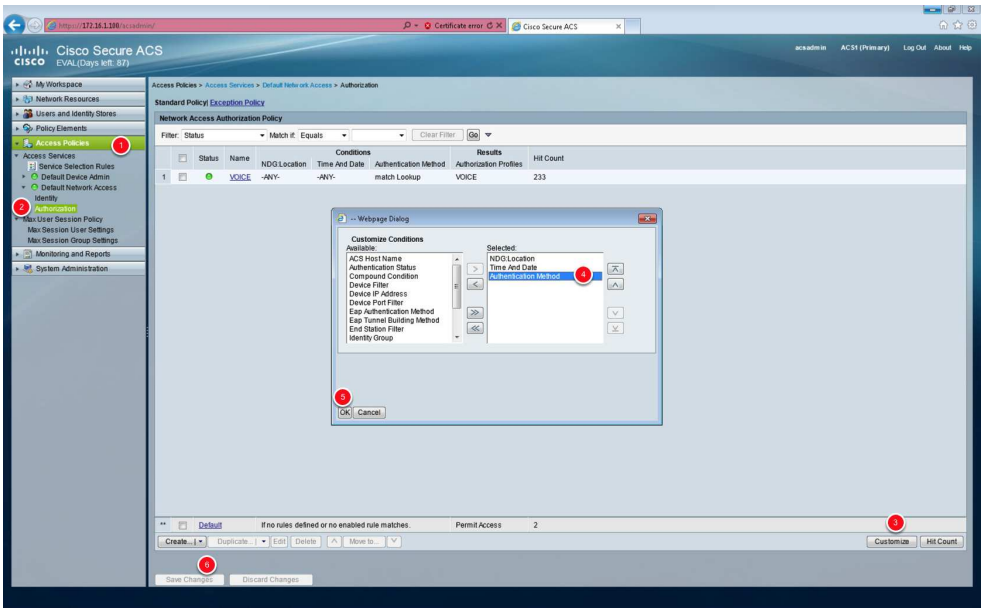


Configure the Identity Source selection for PEAP(EAP-TLS):

https://t.me/learningnets

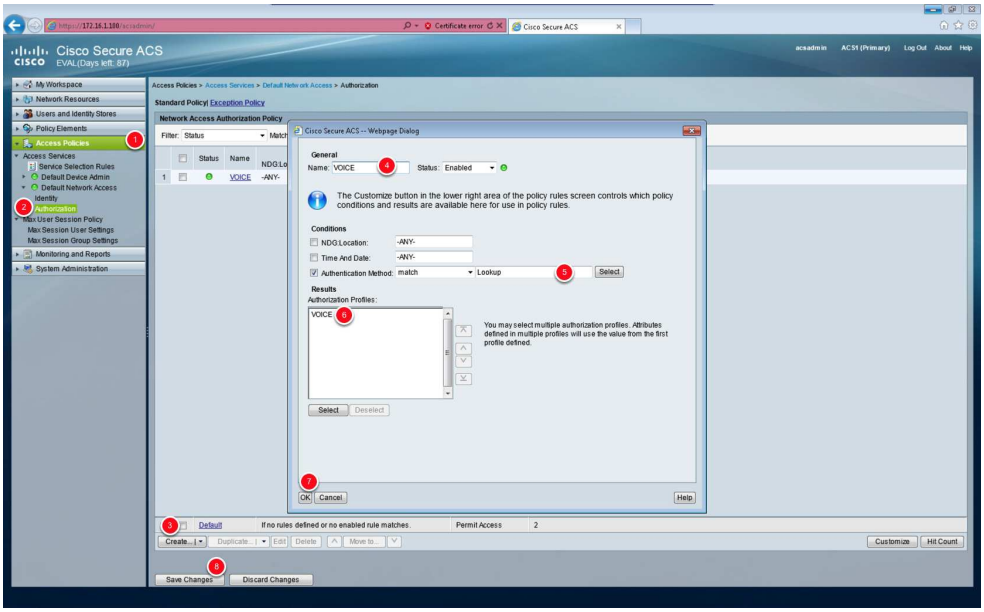


Add the **Authentication Method** as a condition for authorization policies:

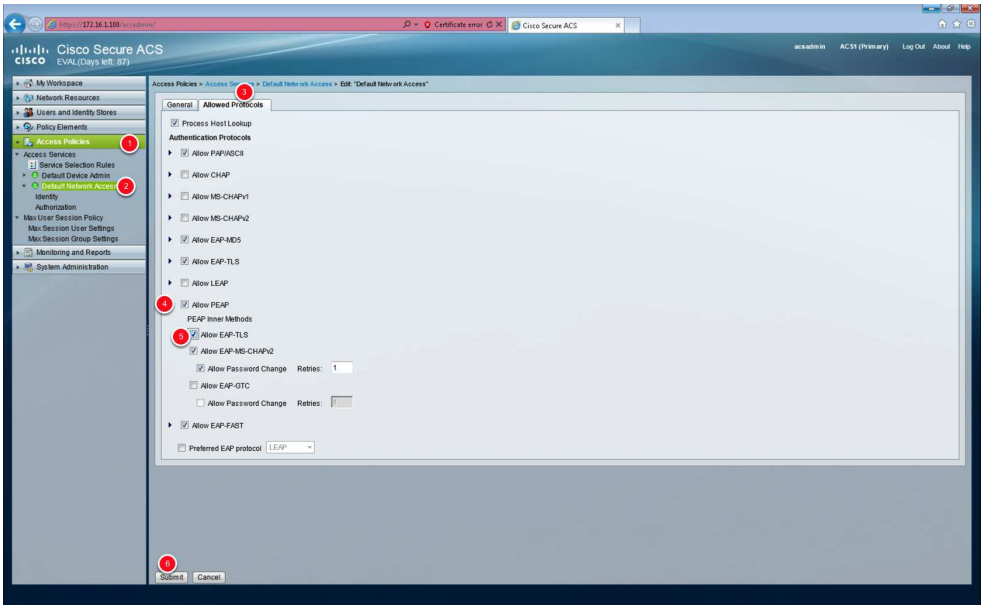


Configure an authorization policy for the VOICE domain; for DATA domain, the default authorization rule of **Permit Access** will be matched:

https://t.me/learningnets



Configure EAP-TLS as the allowed inner method for PEAP:

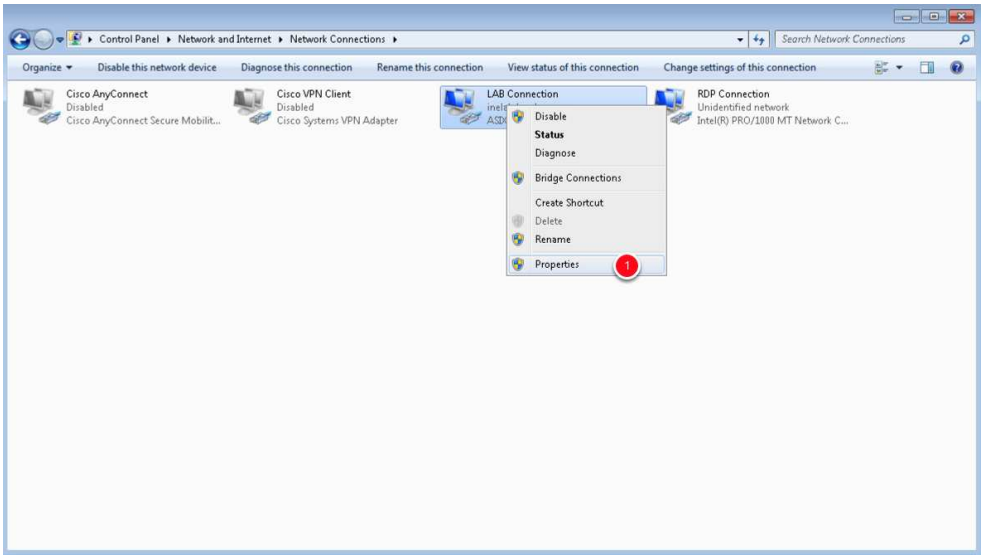


If both Windows native supplicant and AnyConnect supplicant are active, AnyConnect has precedence.

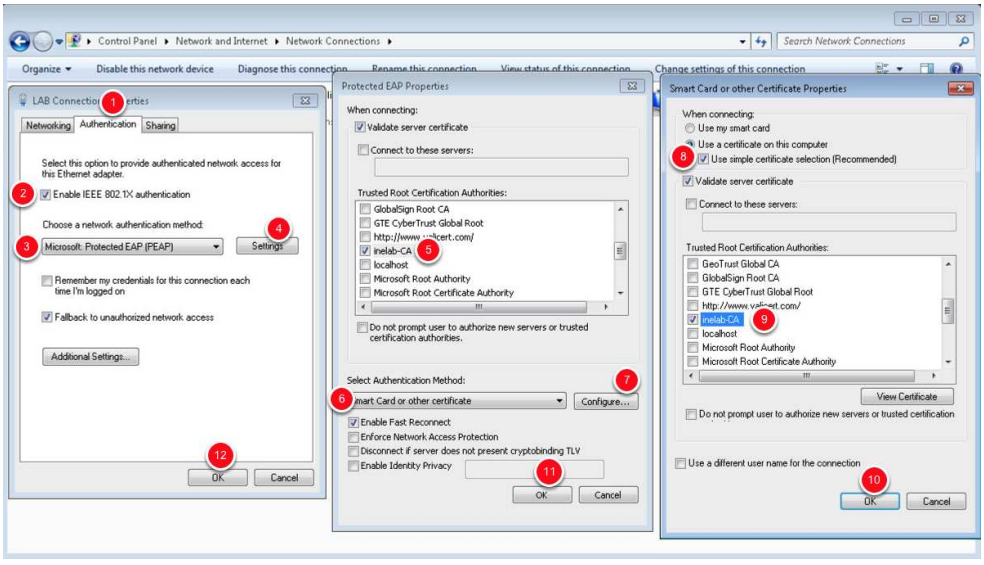
Option 1 Windows Native Supplicant

Configure the native Windows supplicant for the **LAB** Ethernet interface:

https://t.me/learningnets

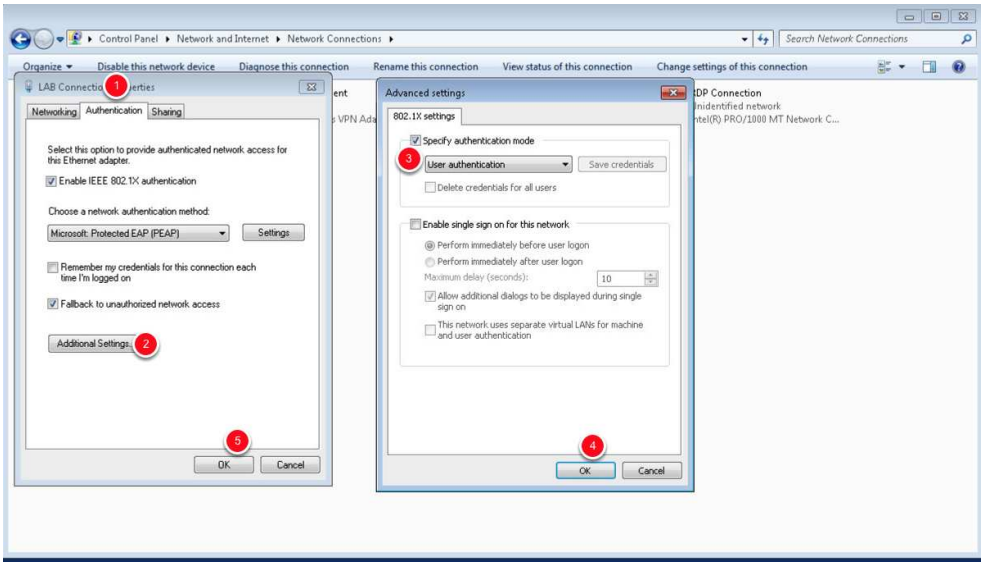


Enable 802.1x for PEAP(EAP-TLS); if the **Authentication** tab is not available, ensure that the **Wired AutoConfig** service is started. To avoid being prompted to trust the ACS certificate if using the **Validate Server Certificate** option, specify which Root CA should be trusted for the ACS issued certificate:



Ensure that user authentication is enabled for PEAP(EAP-TLS) so that a user certificate is selected, rather than a computer certificate, which in this case it does not even exist:

https://t.me/learningnets



Option 2 AnyConnect Supplicant

First, enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConnect 802.1x supplicant service. To access the Windows services, type **services.msc** from the command prompt or Run. Open AnyConnect and navigate to the **Advanced** tab to configure an 802.1x profile for EAP-TLS:



Open AnyConnect and apply the configured profile; you'll be asked to trust the ACS certificate issued by the CA on TEST-SRV-A.

Verification

https://t.me/learningnets

Verify that both domains are successfully authorized, IP Phone using MAB and PC using 802.1x:

https://t.me/learningnets

```
SW1#show authentication sessions interface fastEthernet1/0/5

            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  Unknown User-Name:  Cristian Matei
Status:  Authz Success
Domain:  DATA

      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both Authorized By:  Authentication Server

           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000012A01FC8555
      Acct Session ID:  0x00000221
               Handle:  0x9C00012A

Runnable methods list:
       Method   State
       mab      Failed over dot1x    Authc Success



----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  VOICE

      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both Authorized By:  Authentication Server

      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000007501B26CFC
      Acct Session ID:  0x000000B9
               Handle:  0xE2000075

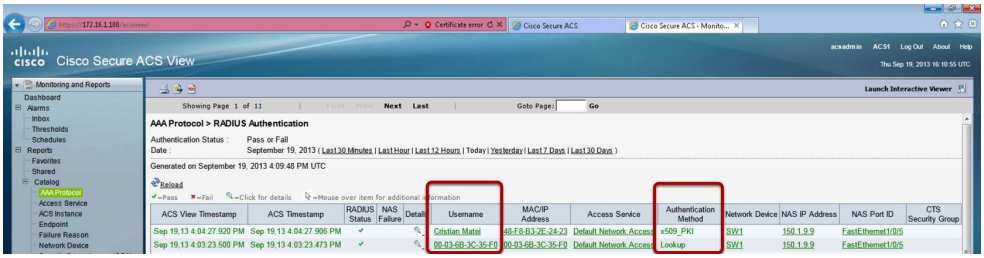
Runnable methods list:
```

https://t.me/learningnets

```
        Method    State mab       Authc Success

        dot1x     Not run
```

Verify the RADIUS Authentication logs:



Verify the list of authorized clients:

```
SW1#show authentication sessions


Interface  MAC Address     Method    Domain    Status          Session ID
Fa1/0/5    48f8.b32e.2423  dot1x     DATA      Authz Success   0000000000000012C0200050CE
Fa1/0/5    0003.6b3c.35f0  mab       VOICE     Authz Success   0000000000000012B02002E23
```

Verify the device tracking databse, and that the pre-authentication access-list was overriden after successful authentication:

```
SW1#show ip device tracking all

IP Device Tracking = Enabled
IP Device Tracking Probe Count = 3
IP Device Tracking Probe Interval = 30
IP Device Tracking Probe Delay Interval = 0
---------------------------------------------------------------------------
  IP Address      MAC Address    Vlan  Interface                STATE
---------------------------------------------------------------------------
136.1.47.51      0003.6b3c.35f0  47    FastEthernet1/0/5        ACTIVE
136.1.27.51      48f8.b32e.2423  27    FastEthernet1/0/5        ACTIVE

Total number interfaces enabled: 1
Enabled interfaces:
  Fa1/0/5
!
!
SW1#show ip access-lists interface fastEthernet1/0/5
```

https://t.me/learningnets

```
permit ip host 136.1.27.51 any
 permit ip host 136.1.47.51 any
```

Invalidate the 802.1x configuration so that the supplicant is active but authentication fails, so the switch will authorize the port on the fail VLAN, which is still VLAN 27. Note that because of the pre-authentication ACL, with the authentication fail action, TEST-PC-A will not be able to generate ICMP traffic:

```
SW1#show authentication sessions interface fastethernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  Unknown User-Name:  anonymous
Status:  Authz Success
Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  single-host
     Oper control dir:  both Authorized By:  Auth Fail Vlan
 Operational HostMode:  SINGLE_HOST
          Vlan Policy:  27
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000A301B65282
      Acct Session ID:  0x00000114
               Handle:  0x8E0000A3

Runnable methods list:
       Method   State
       mab      Failed over dot1x    Authc Failed




----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown
            User-Name:  00-03-6B-3C-35-F0
               Status:  Authz Success
               Domain:  VOICE
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both
        Authorized By:  Authentication Server
```

https://t.me/learningnets

```
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000007501B26CFC
      Acct Session ID:  0x000000B9
               Handle:  0xE2000075

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```

The following messages will be logged by SW1, indicating fallback on the authentication fail VLAN:

```
%AUTHMGR-7-RESULT: Authentication result 'fail' from 'dot1x' for client (48f8.b32e.2423) on Interface Fa1/0/5 AuditS
%AUTHMGR-5-SUCCESS: Authorization succeeded for client (48f8.b32e.2423) on Interface Fa1/0/5 AuditSessionID 00000000
%DOT1X-5-RESULT_OVERRIDE: Authentication result overridden for client (48f8.b32e.2423) on Interface Fa1/0/5 AuditSes
```

Disable both native Windows supplicant and AnyConnect supplicant; SW1 will fallback on guest VLAN:

```
SW1#show authentication sessions interface fastethernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  Unknown
           IP Address:  Unknown User-Name:  UNRESPONSIVE
 Status:  Authz Success
 Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-host
     Oper control dir:  both Authorized By:  Guest Vlan
         Vlan Policy:  27
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000C601CEA85C
      Acct Session ID:  0x00000159
               Handle:  0x6F0000C6

Runnable methods list:
       Method   State
       mab      Failed over dot1x    Failed over
```

https://t.me/learningnets

```
----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown
            User-Name:  00-03-6B-3C-35-F0
               Status:  Authz Success
               Domain:  VOICE
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both
        Authorized By:  Authentication Server
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000007501B26CFC
      Acct Session ID:  0x000000B9
               Handle:  0xE2000075

Runnable methods list:
       Method   State
       mab      Authc Success
       dot1x    Not run
```

The following messages will be logged by SW1, indicating fallback on the guest VLAN:

```
%AUTHMGR-7-NOMOREMETHODS: Exhausted all authentication methods for client (48f8.b32e.2423) on Interface Fa1/0/5 Audi
%AUTHMGR-5-VLANASSIGN: VLAN 27 assigned to Interface Fa1/0/5 AuditSessionID 00000000000000C601CEA85C
%AUTHMGR-5-SUCCESS: Authorization succeeded for client (Unknown MAC) on Interface Fa1/0/5 AuditSessionID 0000000000
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

Installing ACS Certificates

Task

- TEST-SERVER-A is configured as a CA and automatically grants certificates.
- Enroll ACS-1 with the CA server and use the certificate for both management and EAP.
- ACS-1 should use a CN of **ACS1.inelab.local**, and RSA Key Pair of 1024 bit.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

By default, ACS has a self-signed certificate that is used for HTTPS management. The Self-Signed certificate cannot be used for EAP authentications because ACS cannot act as a CA (Certificate Authority), and EAP requires certificates issued by trusted CAs. A public or private CA-issued certificate can be installed on ACS for use with HTTPS management connection as well as for EAP authentications. The benefit of using an external public/private Certificate Authority includes scalability and trust. By default, web browsers and operating systems come installed with trusted root certificates. If you use a certificate authority that's in the trusted store, you will not be presented with any browser warnings when managing ACS. If you use an internal CA server, you must install the Root Certificate on the client managing the ACS.

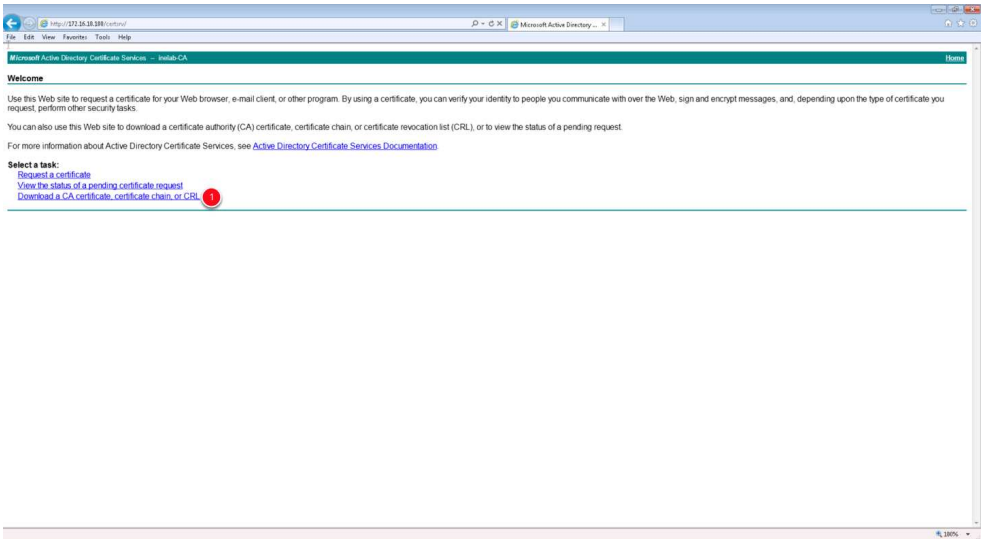
https://t.me/learningnets

A more detailed overview for a PKI infrastructure is presented in the VPN Section. The basic process of enrolling ACS with a CA is to generate a certificate signing request, submit the request to the CA, and download and install the issued identity certificate along with the CA Server root certificate.

> Note that even after installing a certificate issued by a trusted CA, you may still see browser warnings if the FQDN or the IP address you are using in the browser for managing ACS does not match the CN or SAN values from the certificate; this is another check the browser does to consider the certificate valid.

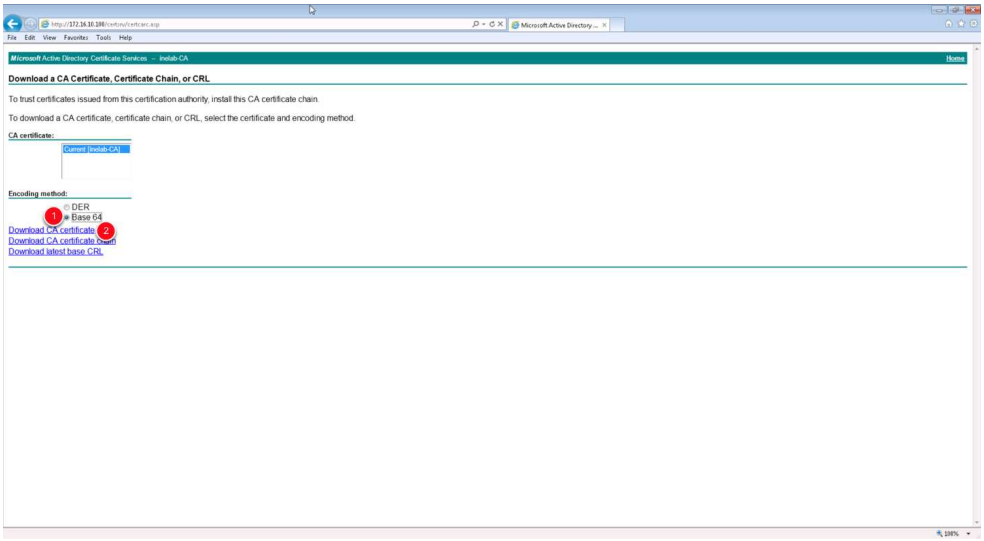
Configuration

Browse the CA server web enrollment page to download the CA certificate:

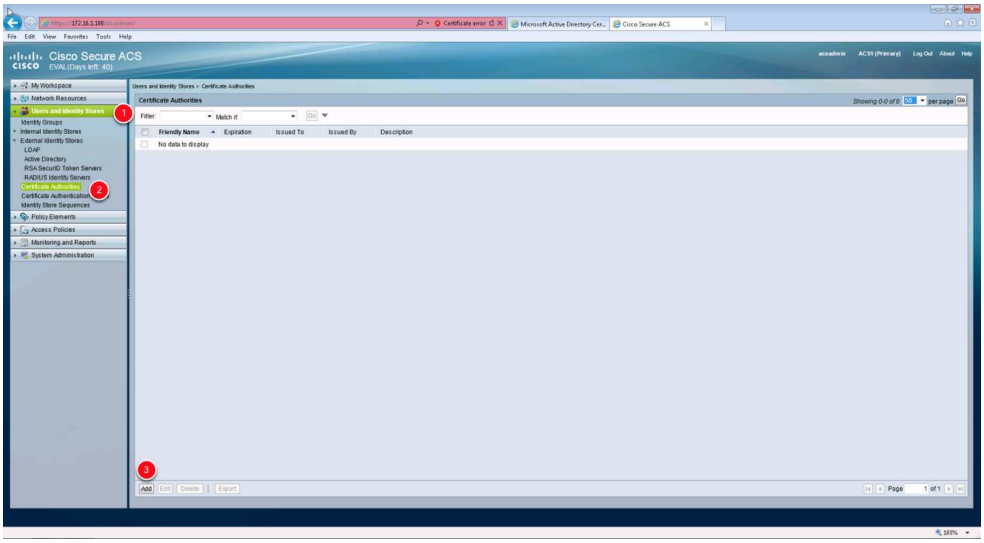


Select the **Base 64** format, download the CA certificate, and save the file locally on the PC:

https://t.me/learningnets

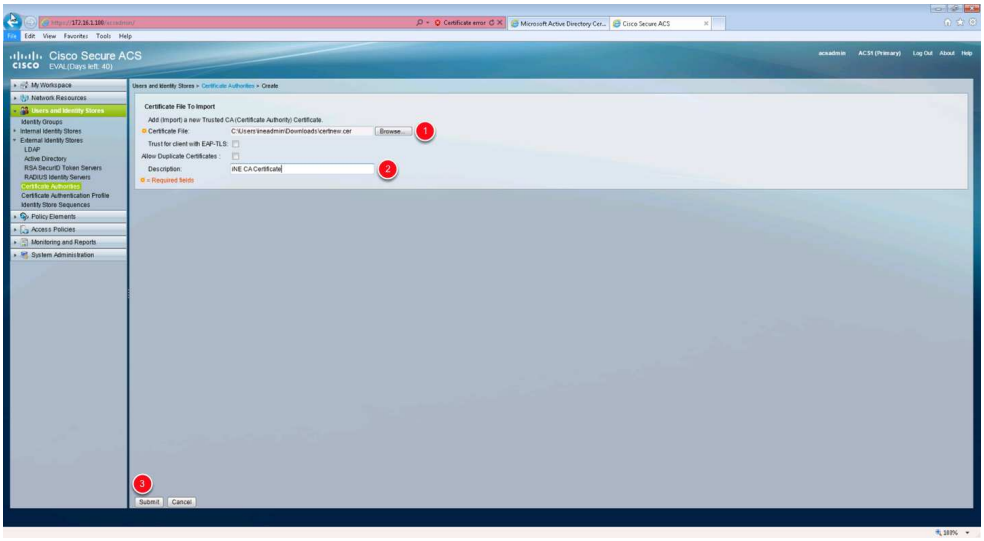


Browse to the **Certificate Authorities** section of the ACS:

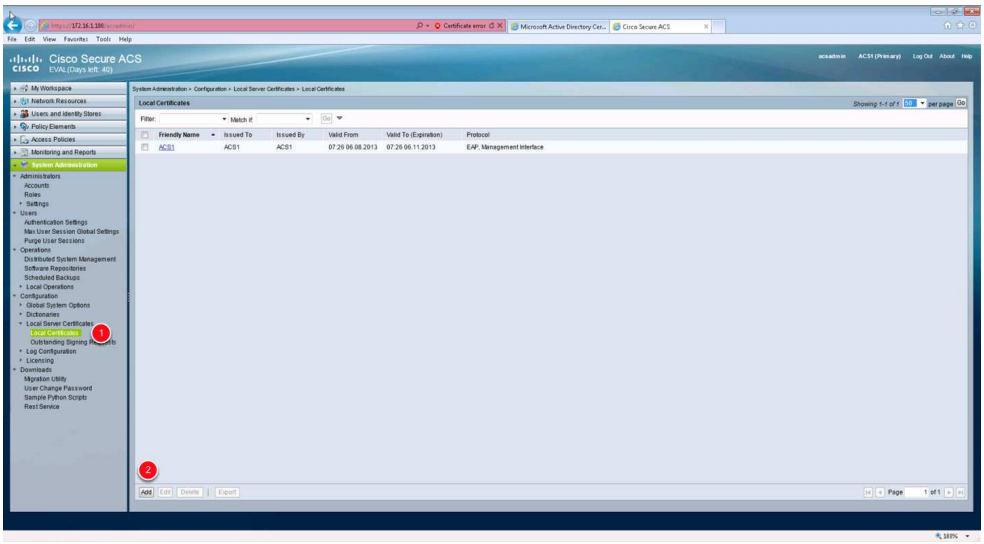


Import the CA certificate from the saved file; to trust client certificates issued by this CA (for example, for EAP-TLS handshake), select the **Trust for client with EAP-TLS** check box:

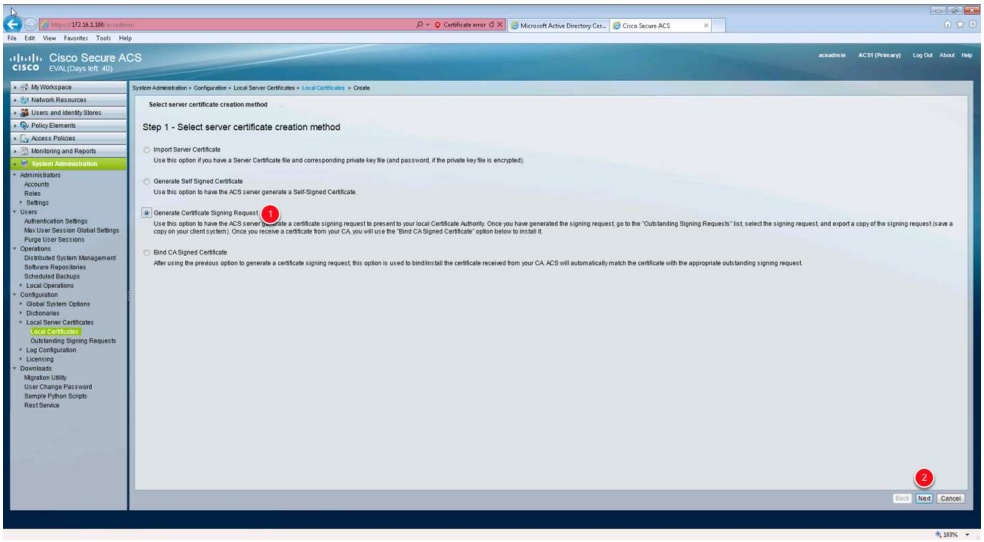
https://t.me/learningnets



Browse to the **Local Certificates** section of the ACS:

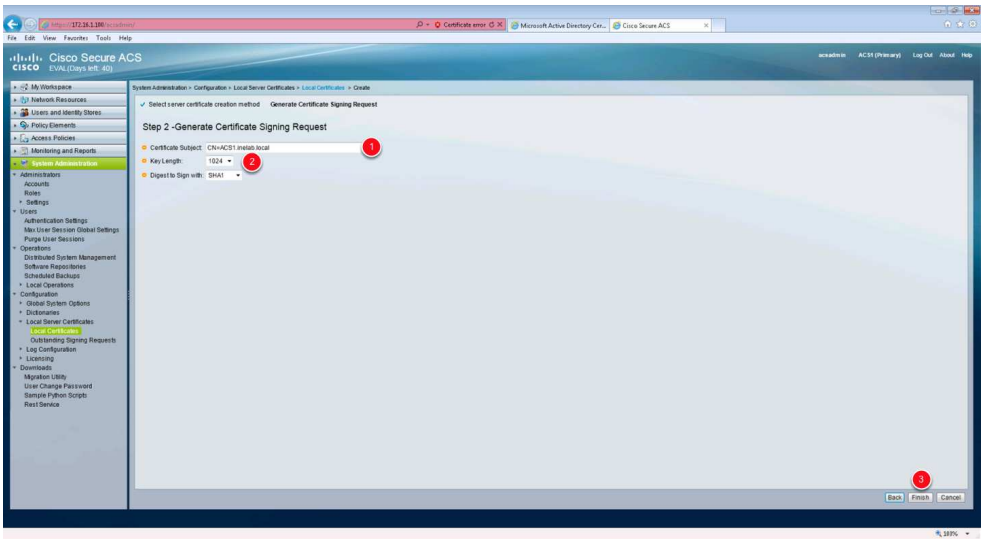


Select to generate a new CSR (Certificate Signing Request):

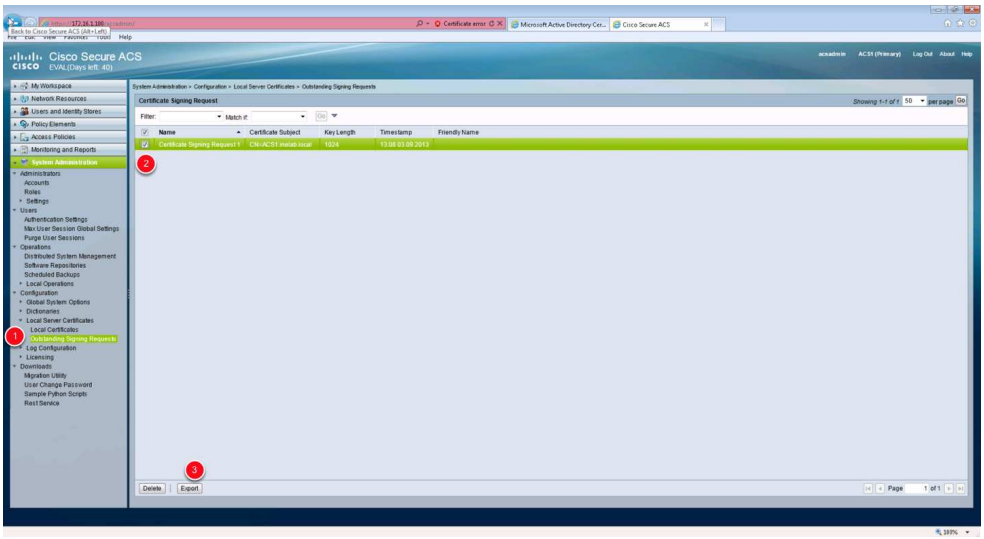


https://t.me/learningnets

Configure the CSR input data as required by the task:

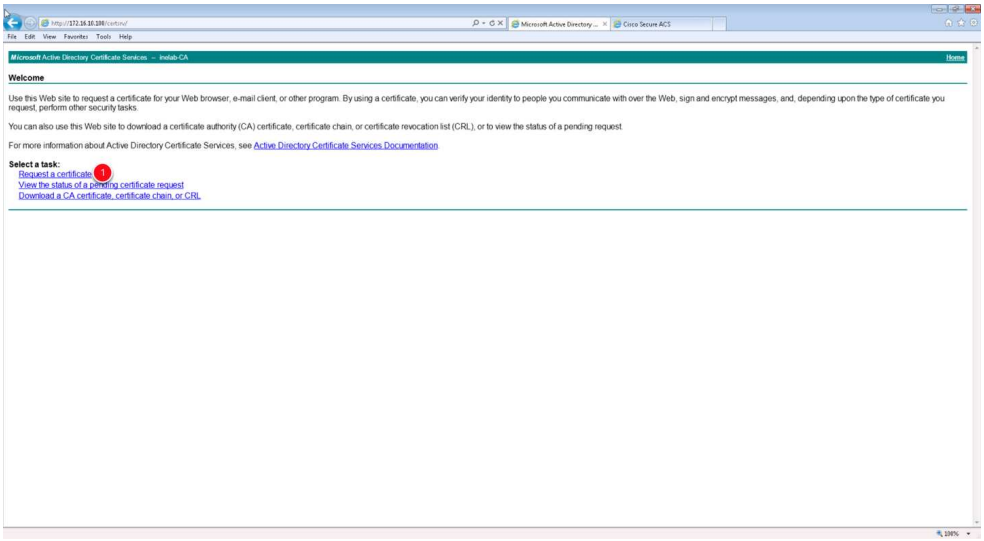


Export the CSR Request to use it asking a certificate from the CA; save the CSR data to a Wordpad file:

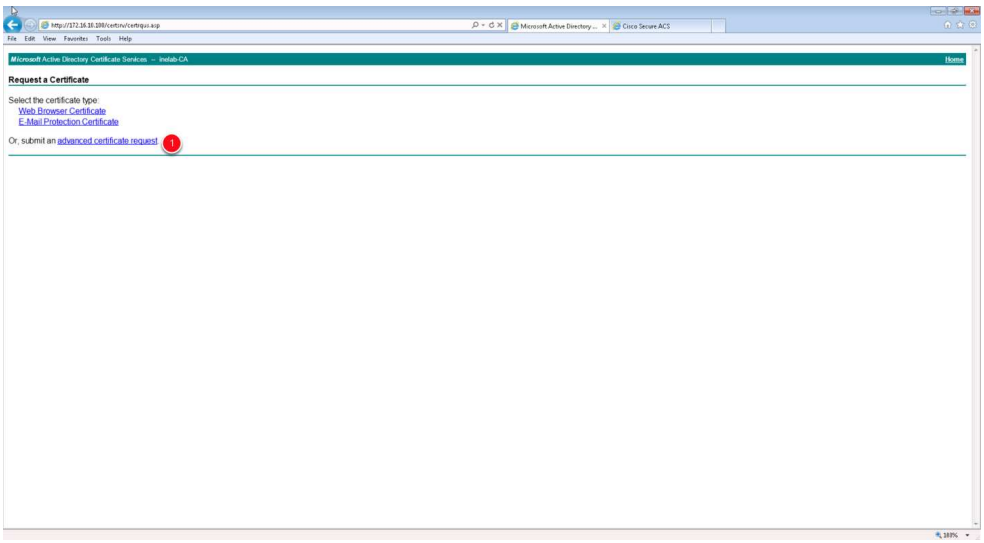


Browse the CA server again to request a certificate:

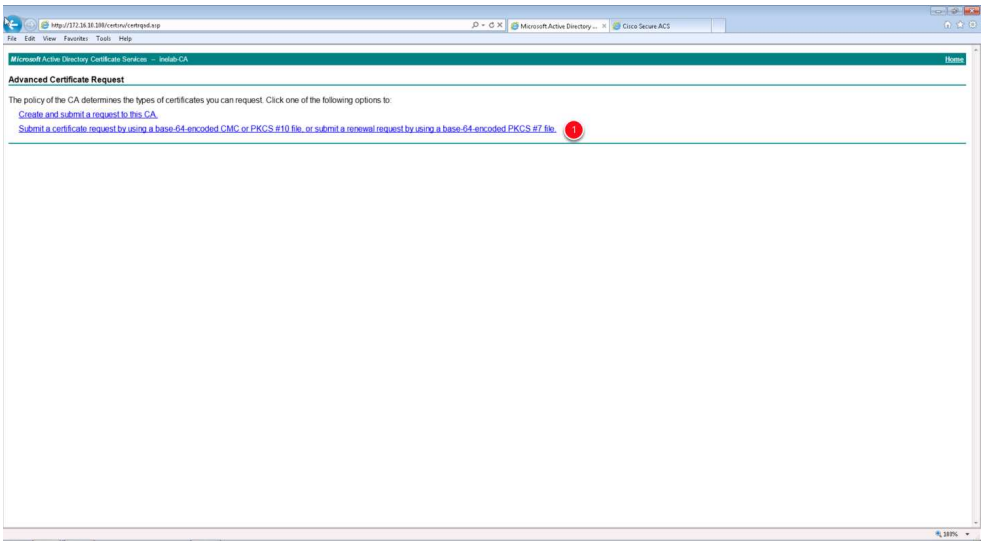
https://t.me/learningnets



Select the **Advanced Certificate Request** option to make use of the CSR:

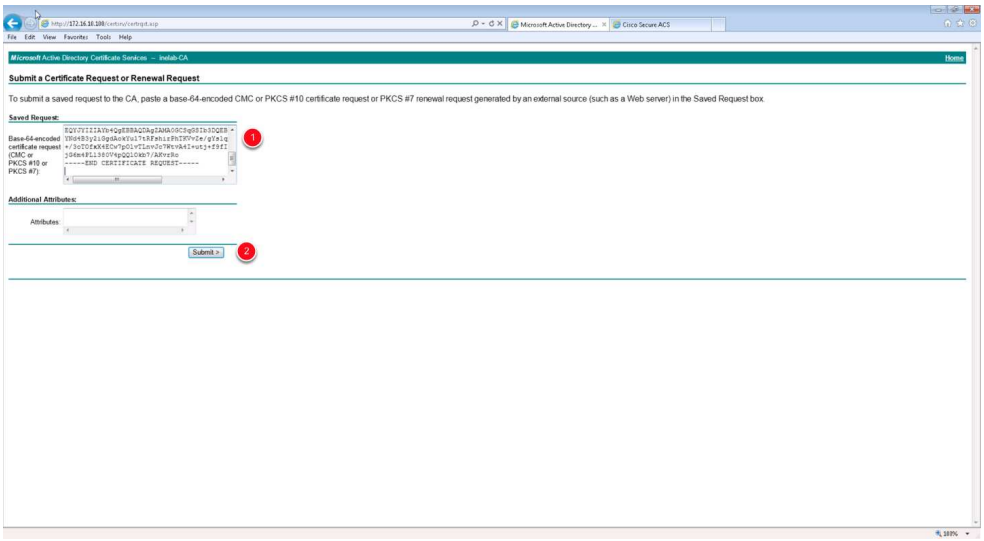


From next screen, use the second option:

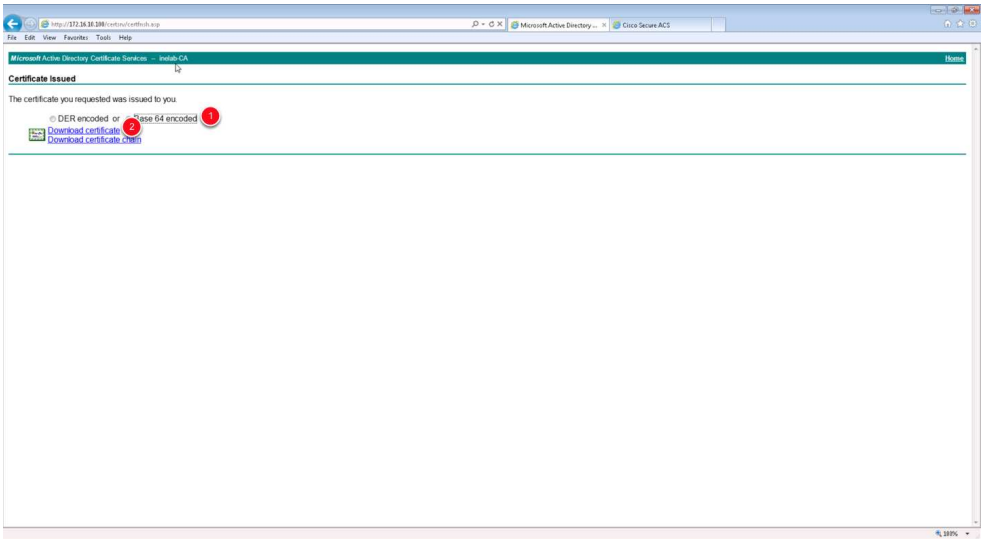


https://t.me/learningnets

Copy and paste the CSR data into the **Saved Request** window:

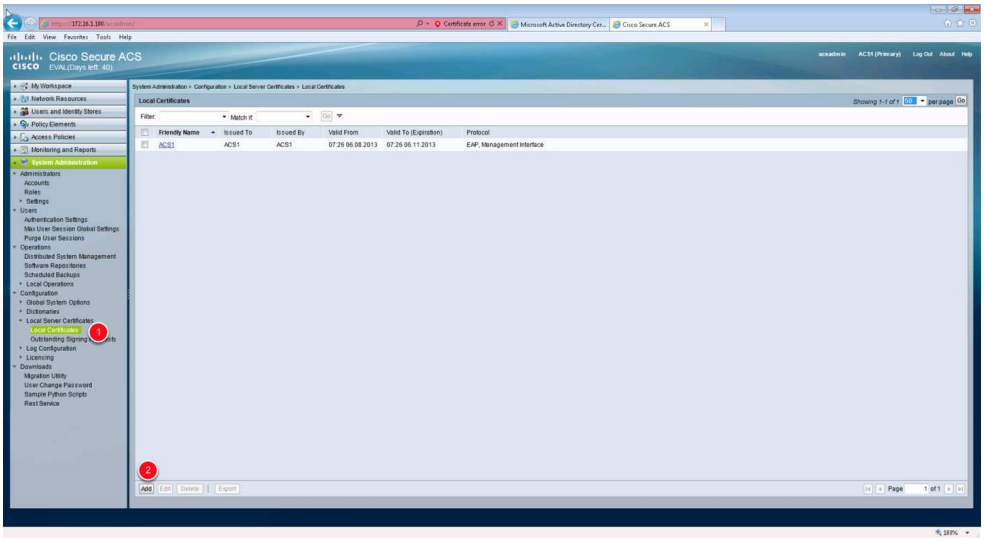


Based on the settings from the CA server, the certificate was automatically granted; browse again to the CA server to download the certificate, and save the file locally on the PC:

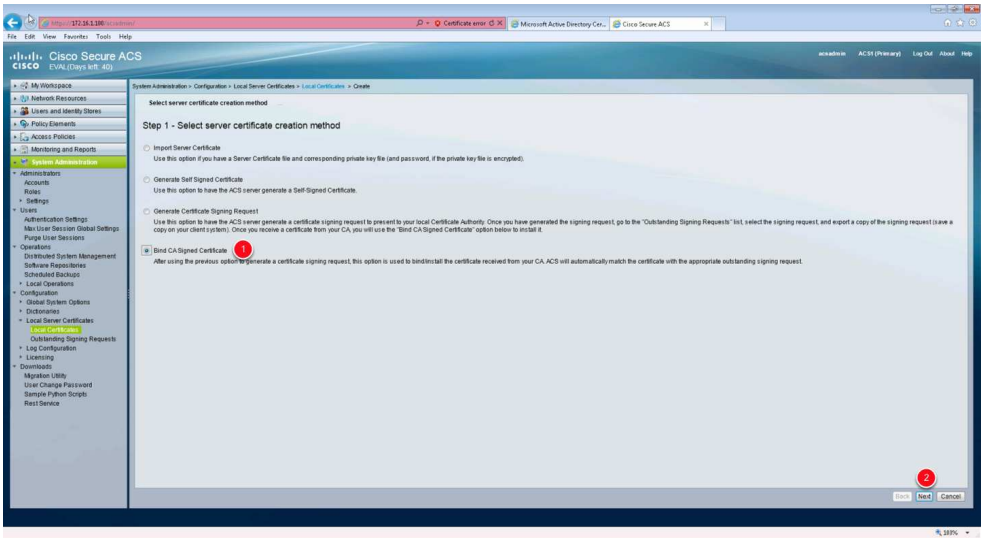


On ACS, browse to the **Local Certificates** section to install the certificate:

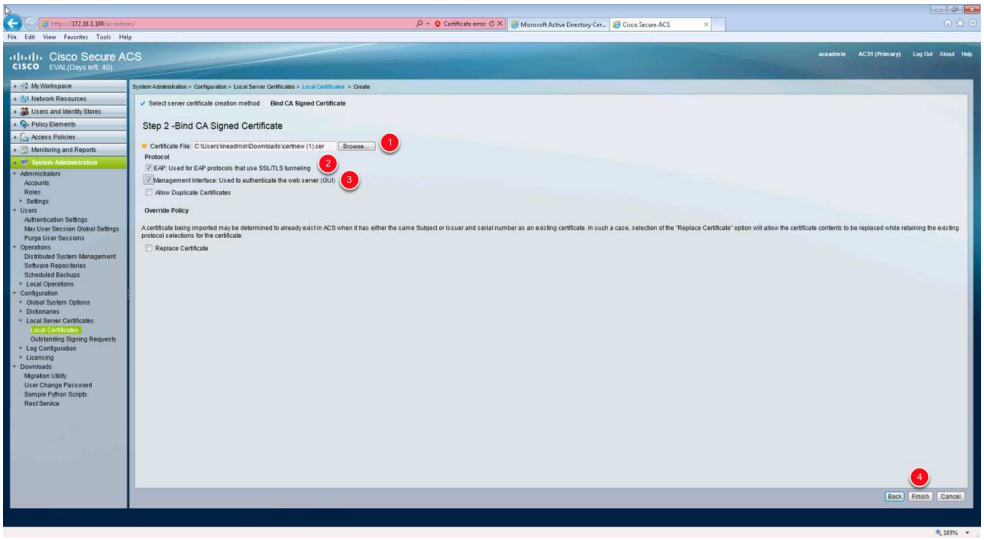
https://t.me/learningnets



Choose the last option to bind the certificate with the generated CSR:



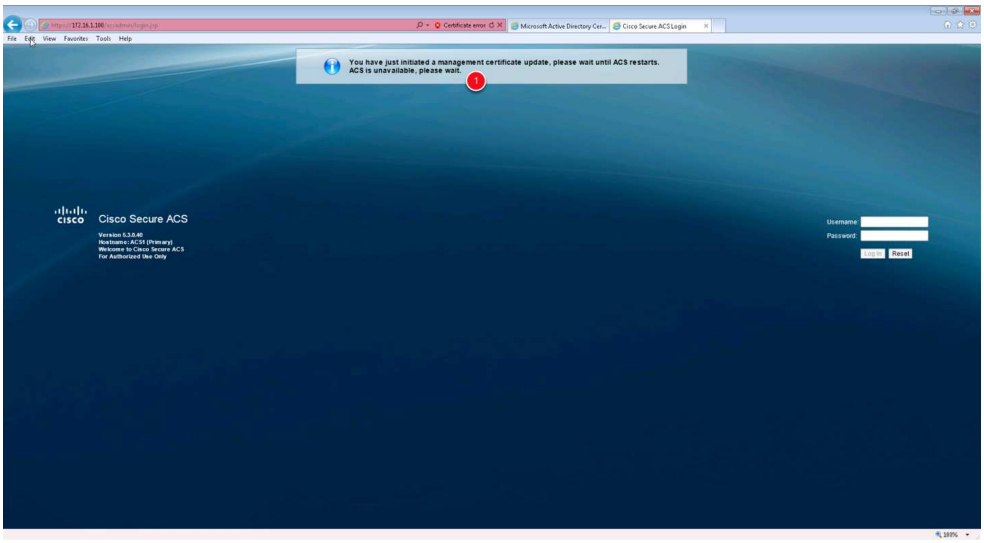
Import the certificate from the saved file and use it for both management and EAP:



To install the new certificate, which is also used for management, ACS will restart

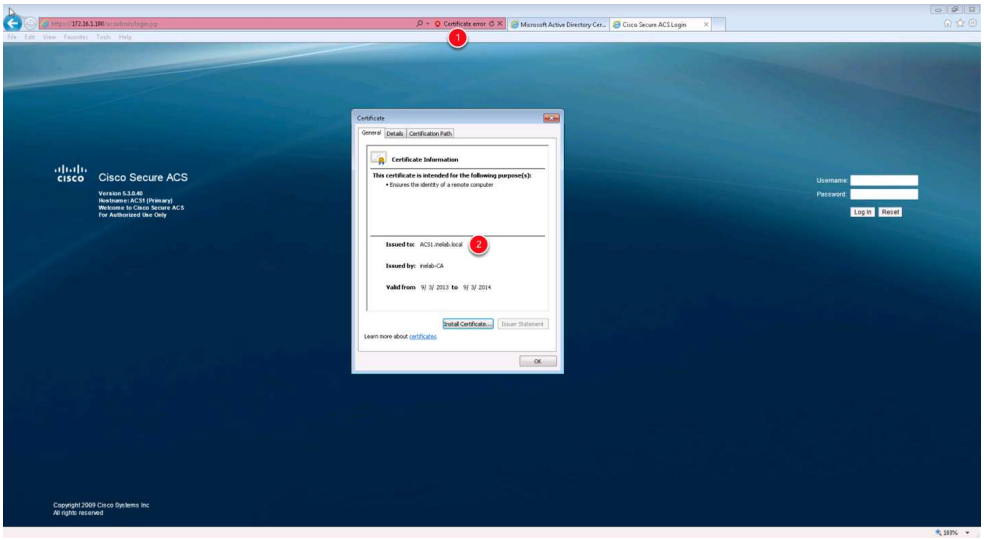
https://t.me/learningnets

and you will be logged out:



Verification

After ACS is restarted, browse it and verify the certificate used for the management session:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

PEAP (EAP-TLS) Using ISE

Task

- Configure **TEST-PC-A** for 802.1x authentication using PEAP; both the supplicant and authentication server should use certificates issued by the CA from **TEST-SRV-A.**
- Ensure that the IP Phone is successfully authorized for VOICE domain.
- Before authentication, all traffic should be allowed; do not use any access-lists for this.
- **TEST-PC-A** should be forced to reauthenticate unconditionally every 8 minutes and 20 seconds; the timer should be received from the RADIUS server.
- If the switch loses connectivity with ISE, when this is recovered any unauthenticated ports should be forced to authenticate.
- Do not add any Network Devices into ISE.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

By default, when a port/device is authenticated and authorized by the RADIUS server, it will not be asked to authenticate again, regardless of whether is its active. The switch can force the reauthentication process if enabled via the port-level command `authentication periodic` . The interval for reauthentication is configured with the port-level command `authentication timer reauthenticate` , and can be locally

https://t.me/learningnets

significant or received from the RADIUS server as RADIUS Attribute 27 (Session-Timeout).

Before a device is authenticated and authorized by the RADIUS server, it is granted to pass traffic as follows:

- With or without a pre-authentication access-list being applied, only EAPoL is allowed.
- If the `authentication open` command is configured without a pre-authentication access-list being applied, all traffic is allowed.
- If the `authentication open` command is configured with a pre-authentication access-list being applied, traffic is allowed based on the access-list.

To verify whether the switch has RADIUS (authentication and accounting) connectivity with ISE or ACS, the switch can be configured to periodically send a dummy **Access-Request** using a specific username. Note that the authentication process does not need to be successful; an **Access-Reject** is good enough for the switch to conclude that RADIUS server is reachable. The dummy account and timer settings are configured in-line with the RADIUS server host and key setup.

The `radius-server vsa send authentication` was applied here just to make use of the default ISE authorization policy for IP Phones, which also sends a dACL with `permit ip any any` . For the same reason, the `ip device tracking` is configured so that the dACL can actually be applied successfully.

Configuration

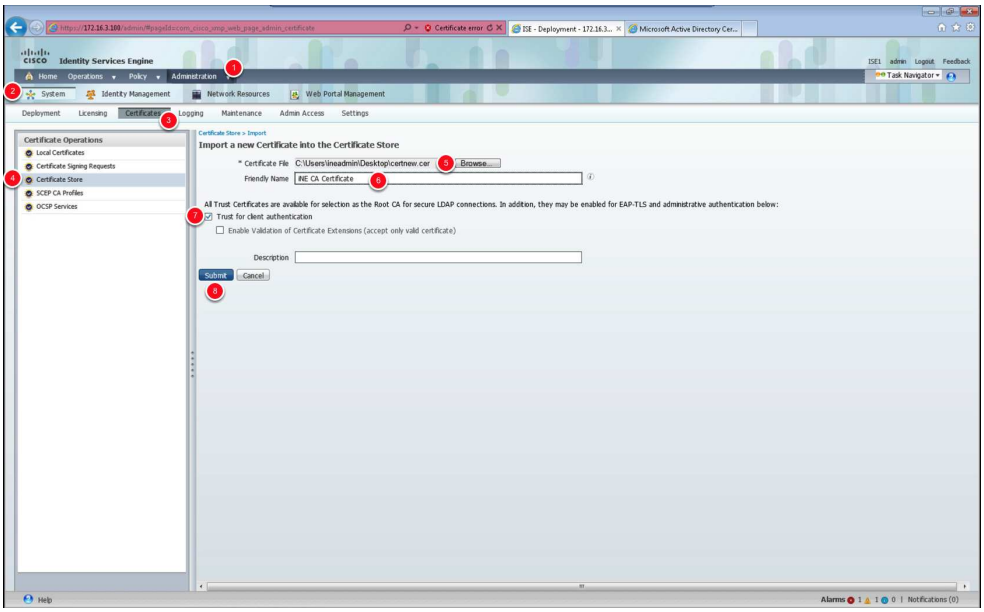
SW1:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
dot1x system-auth-control
radius-server vsa send authentication
ip device tracking
!
interface FastEthernet1/0/5
 authentication open
 authentication periodic
 authentication port-control auto
 authentication host-mode multi-domain
 authentication timer reauthenticate server
 authentication event server alive action reinitialize
 dot1x pae authenticator
```

https://t.me/learningnets

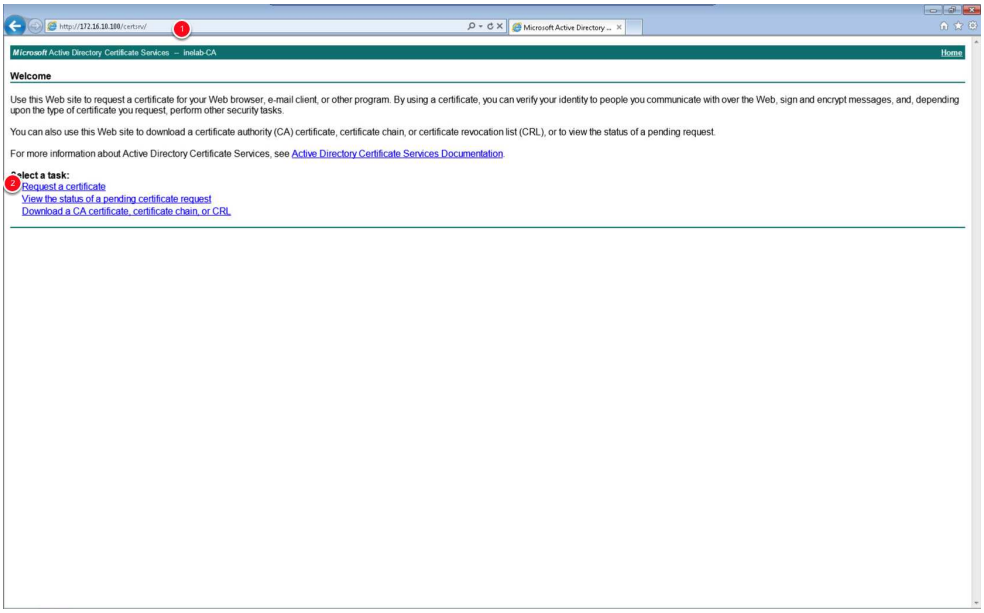
```
 mab
!
aaa group server radius ISE
 server 172.16.3.100
 ip radius source-interface Loopback0
!
radius-server host 172.16.3.100 test username radiustest ignore-acct-port idle-time 1 key CISCO
```

First, request and install a certificate from the CA server for ISE. See the Installing ISE Certificates lab. When installing the certificate, you can specify to use it only for EAP, and leave the self-signed certificate to be used for management via HTTPS. Also make sure to mark the CA certificate as trusted for EAP-TLS, so that client certificates issued by this CA are trusted by ISE and the trust relationship required for EAP-TLS can be established:

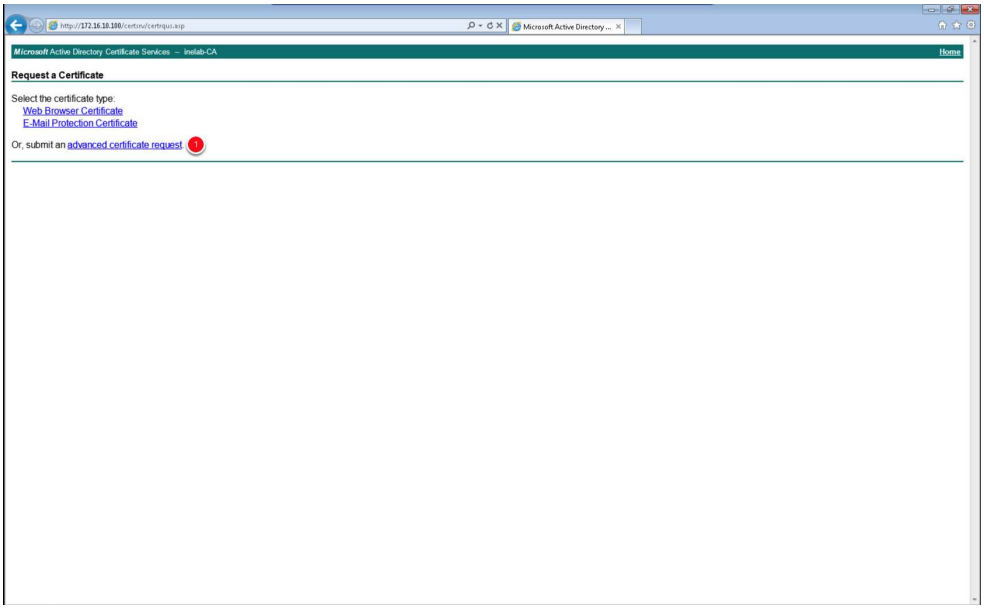


Request a client certificate for **TEST-PC-A** from the CA server; begin by browsing to the CA URL of **http://172.16.10.100/certsrv**:

https://t.me/learningnets

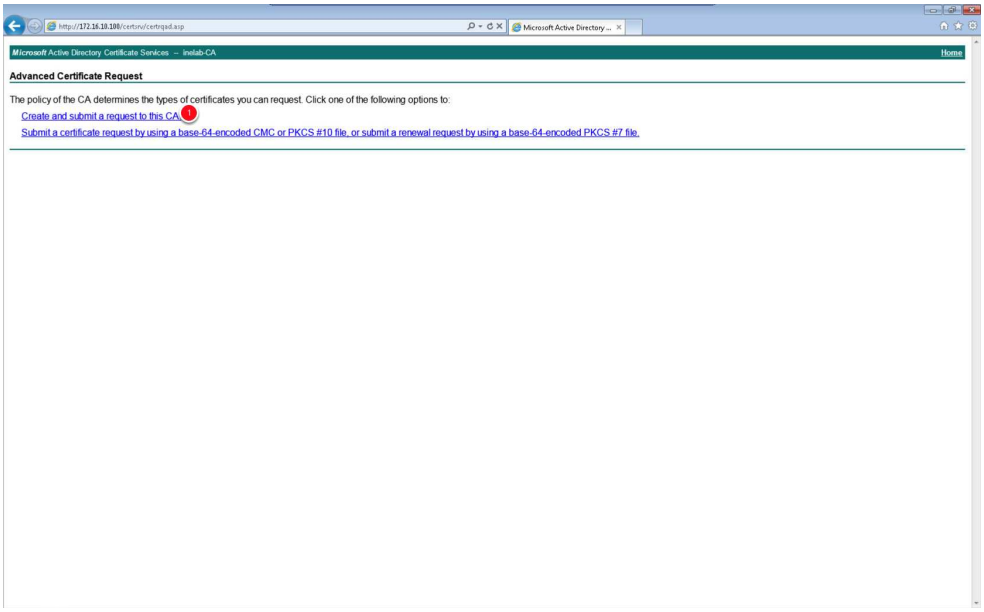


Tto request a client certificate usable for 802.1x, select the **Advanced Service Request** option; this may differ if you are using your own CA infrastructure:

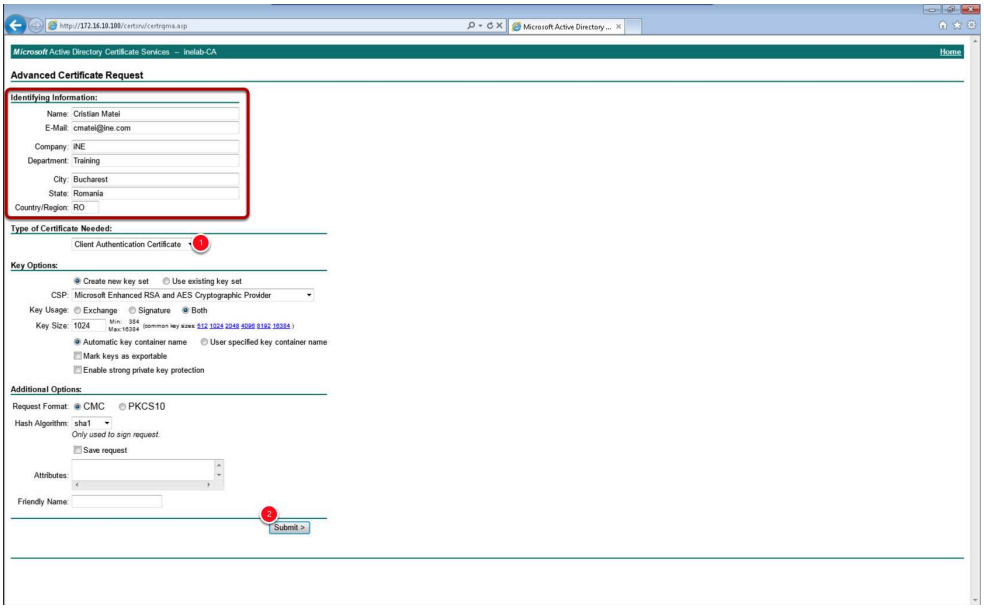


We must create the CSR data and submit it as well, so select the first option. If presented with a warning, select **Yes**. When enrolling the ISE into CA, the second option was selected because the CSR data was generated by ISE itself:

https://t.me/learningnets

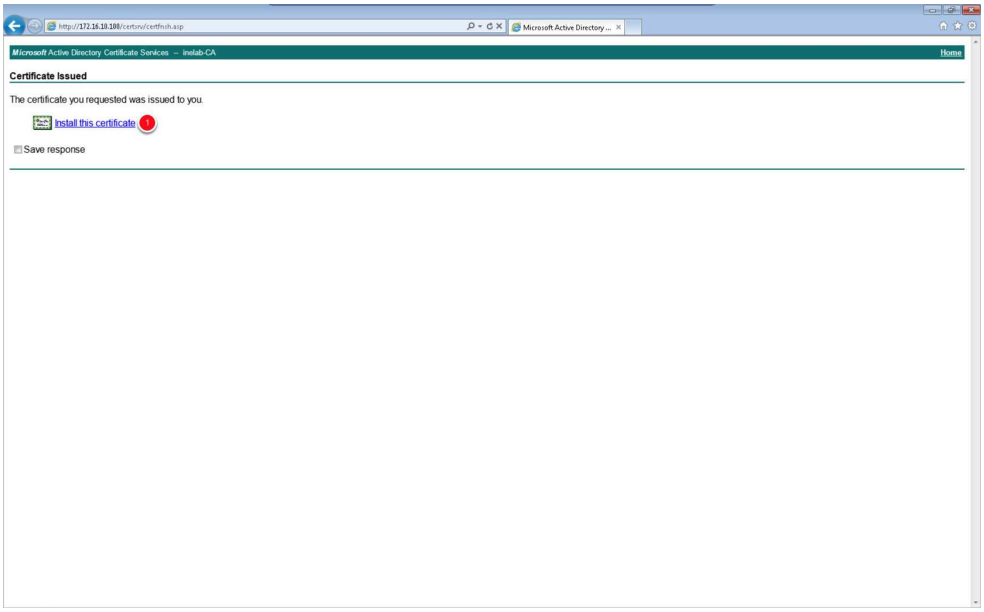


Complete the CSR data with any values because no restrictions are imposed by the task, and click **Submit**:

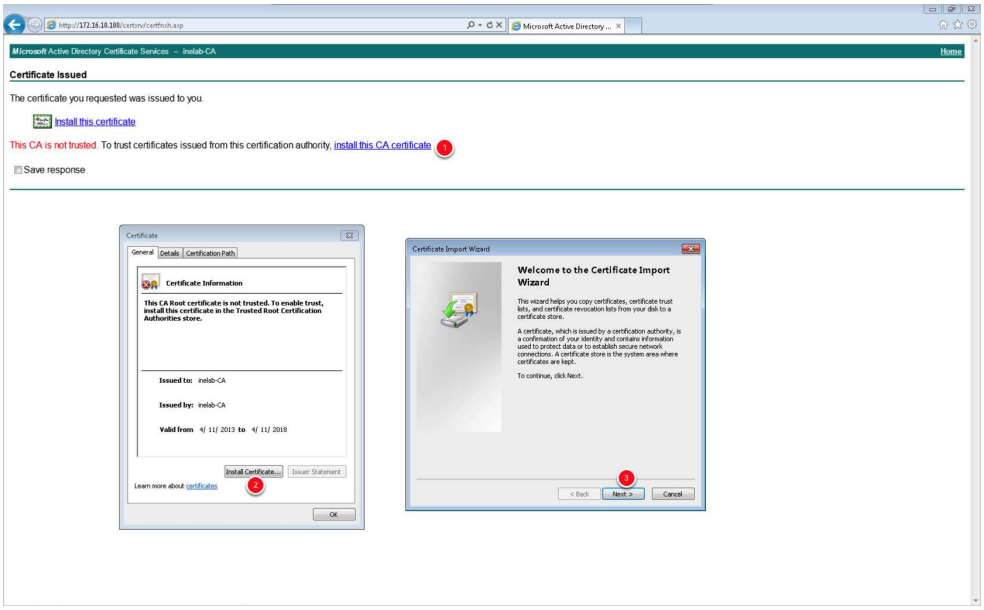


Because the CA is configured to automatically grant certificates, after clicking **Submit** you will be presented with the certificate; click to install it in the local User Repository:

https://t.me/learningnets

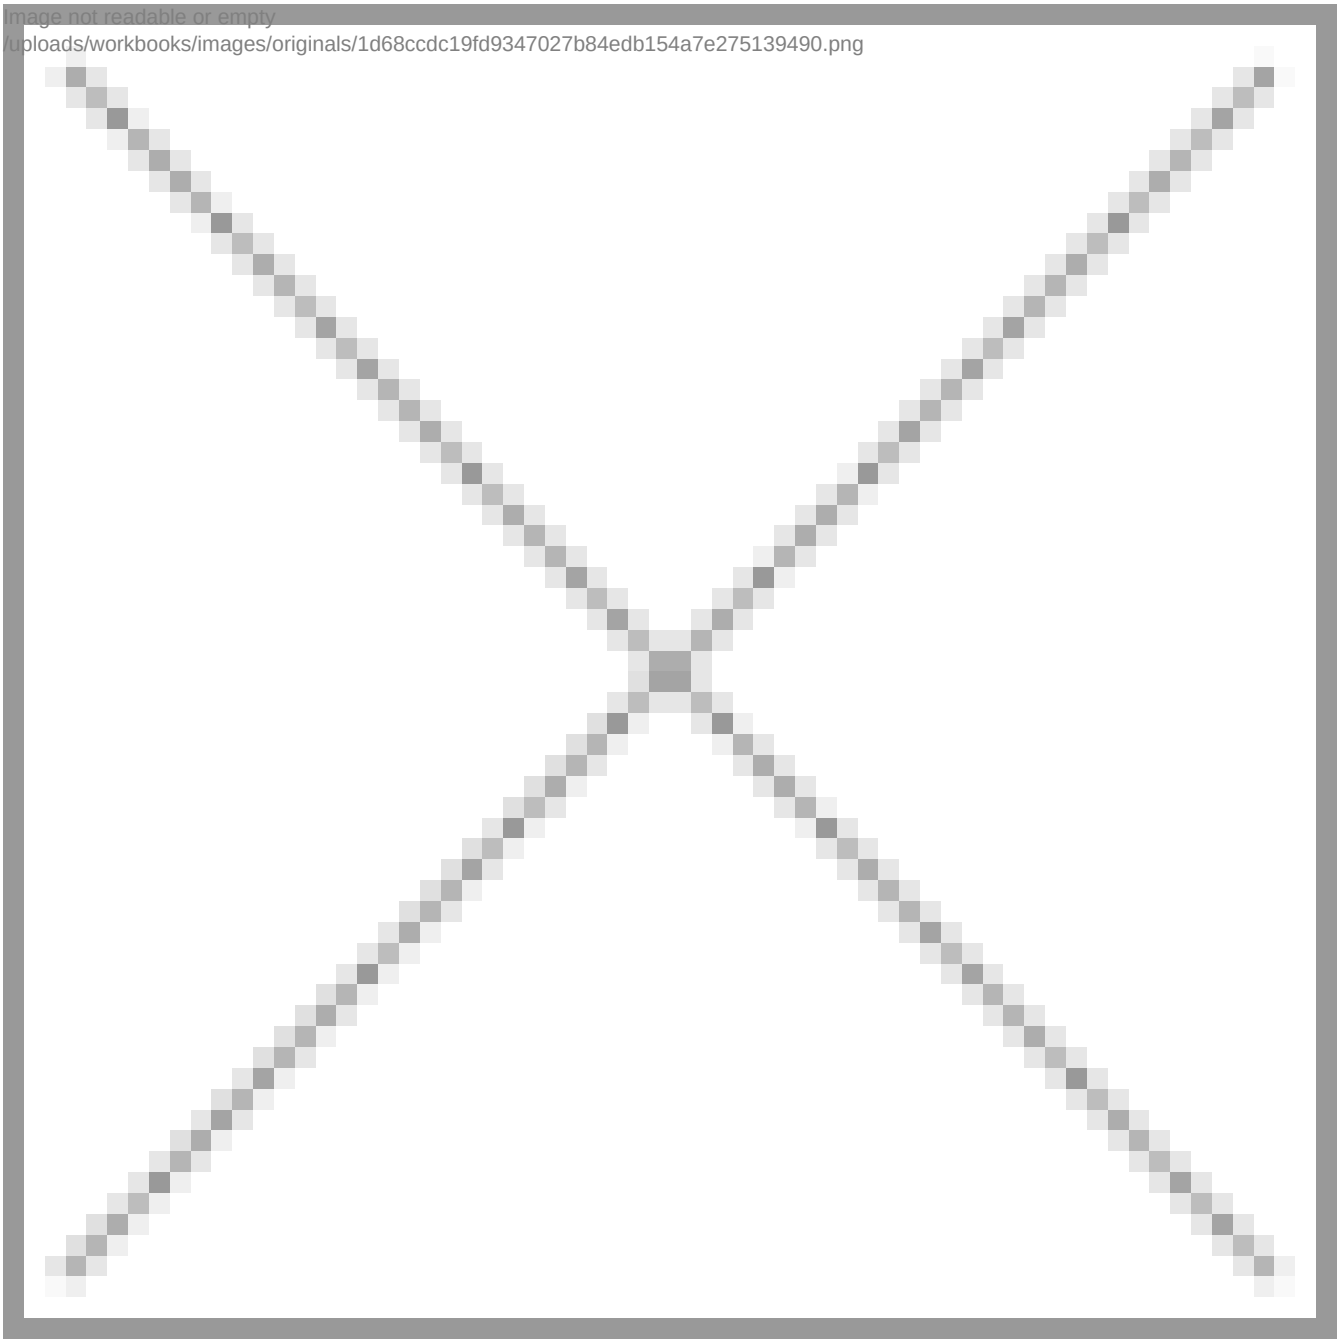


Because the CA is an internal private one, the CA certificate must be added to the **Trusted Root Certification Authorities** before the issued certificate can be installed, so you will be presented with a screen asking for this:



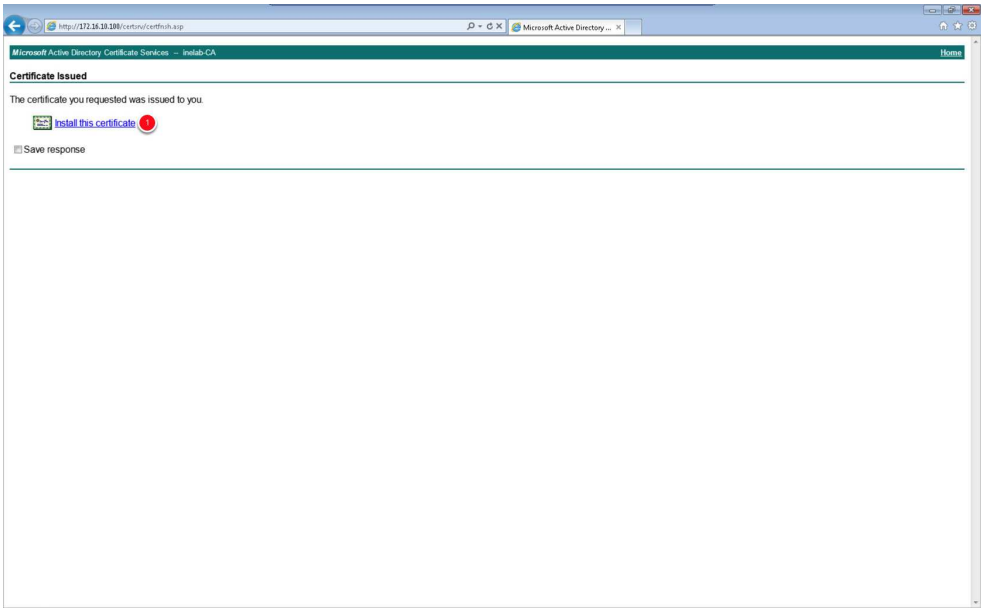
You must manually choose to install the CA certificate into the **Trusted Root Certification Authorities**. Then click **Next** until the end of the wizard, and wait for a successful import message:

https://t.me/learningnets

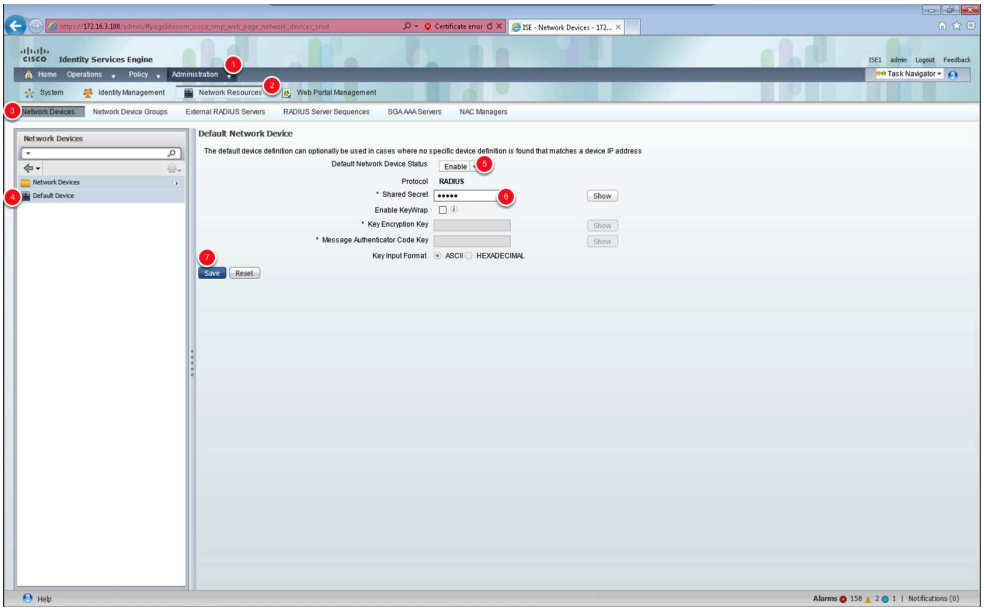


After the CA certificate is installed, click again to install the client certificate:

https://t.me/learningnets

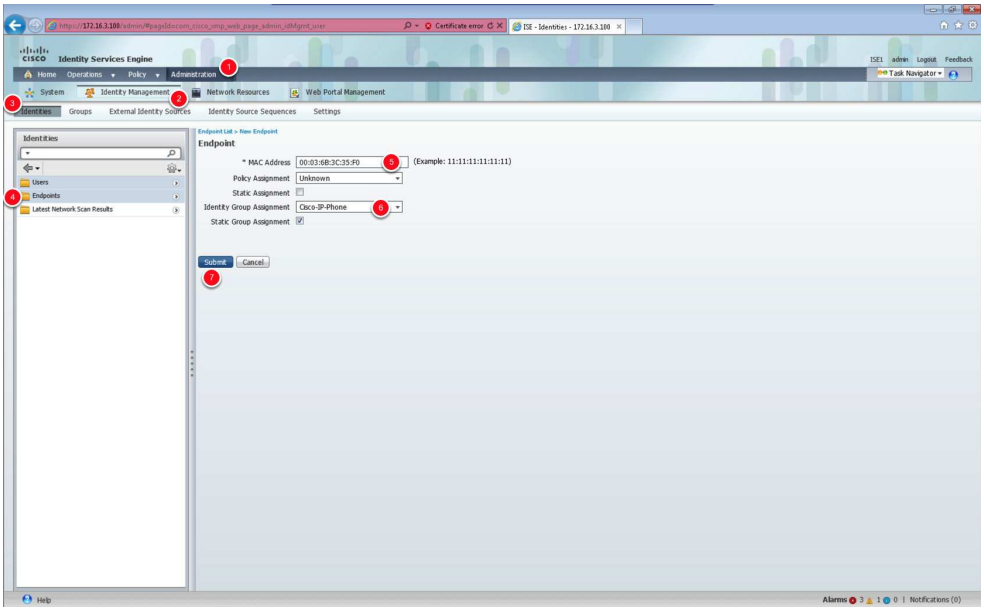


Because no NADs are allowed to be added, configure RADIUS settings for the **Default Device**:

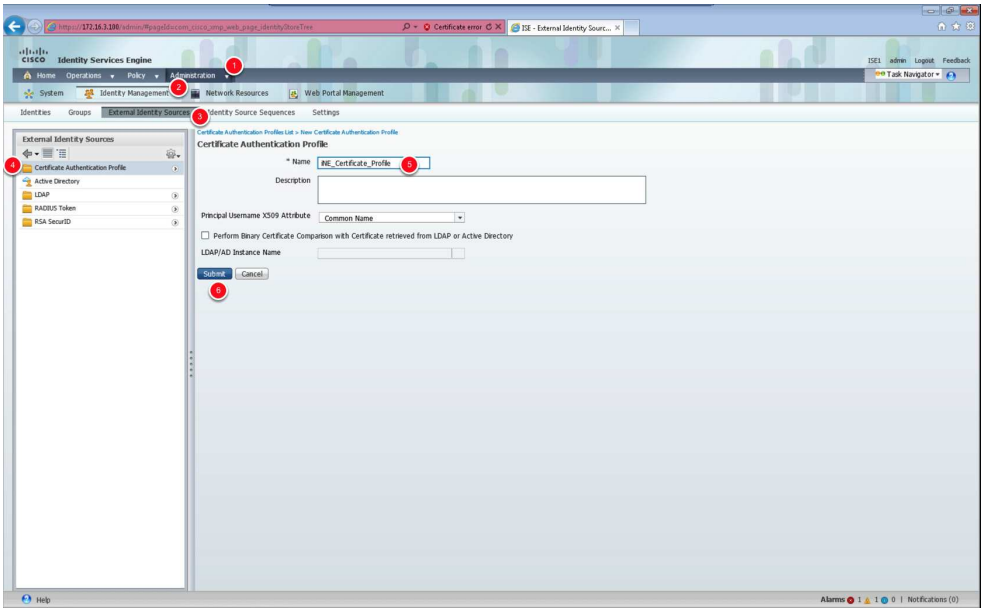


Add the MAC address of the IP Phone in the **Cisco-IP-Phone** Identity Group to be used by MAB and match the default authorization rule for IP Phones. The task does not restrict use of ISE authorization policies, so we can make use of the default one. This rule automatically authorizes the voice VLAN:

https://t.me/learningnets

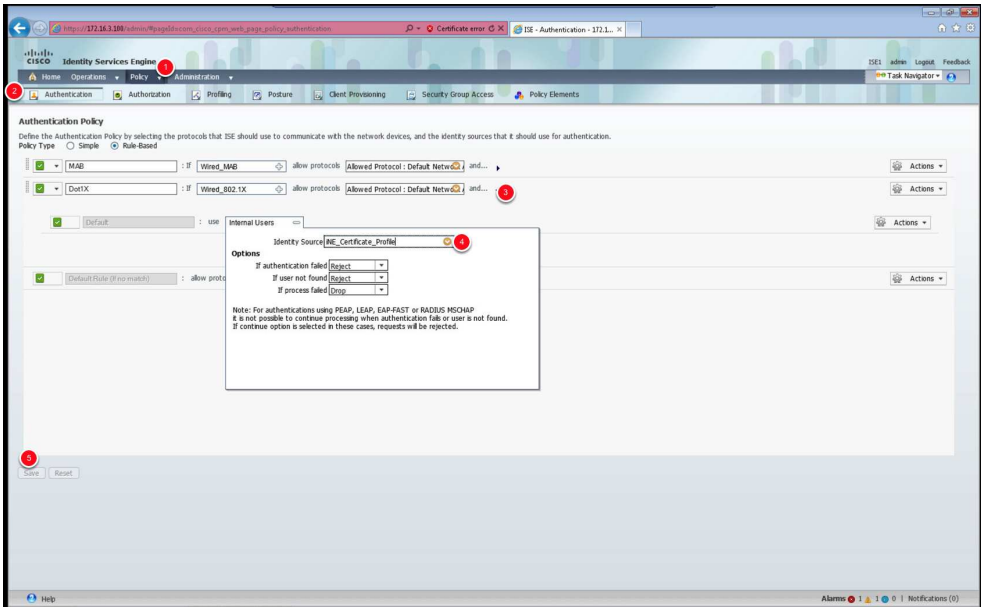


Configure a **Certificate Authentication Profile** to be used in the authentication policy for **TEST-PC-A**:

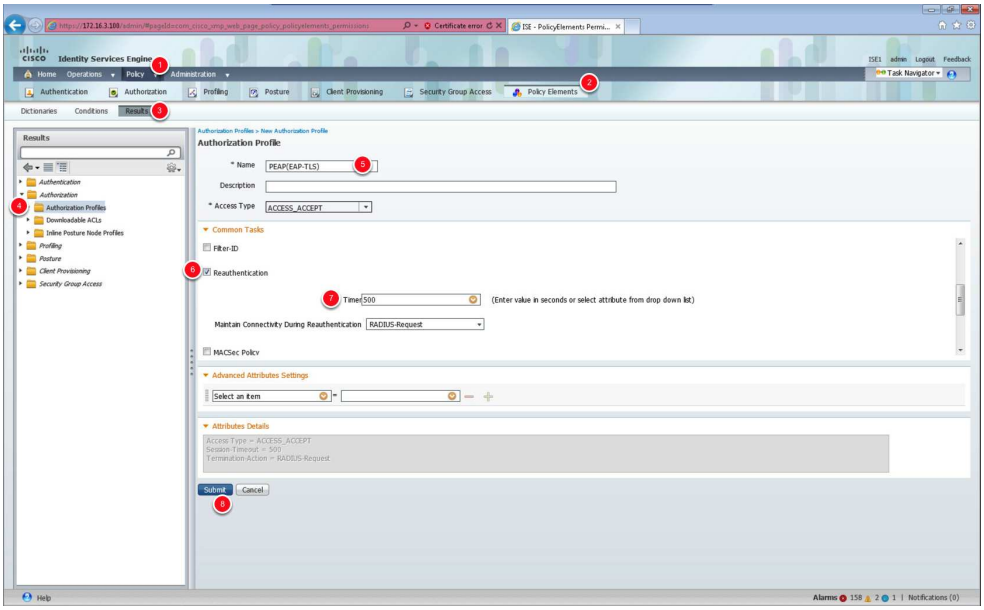


Modify the default 802.1x authentication policy to use the created certificate-based authentication profile as the Identity source:

https://t.me/learningnets

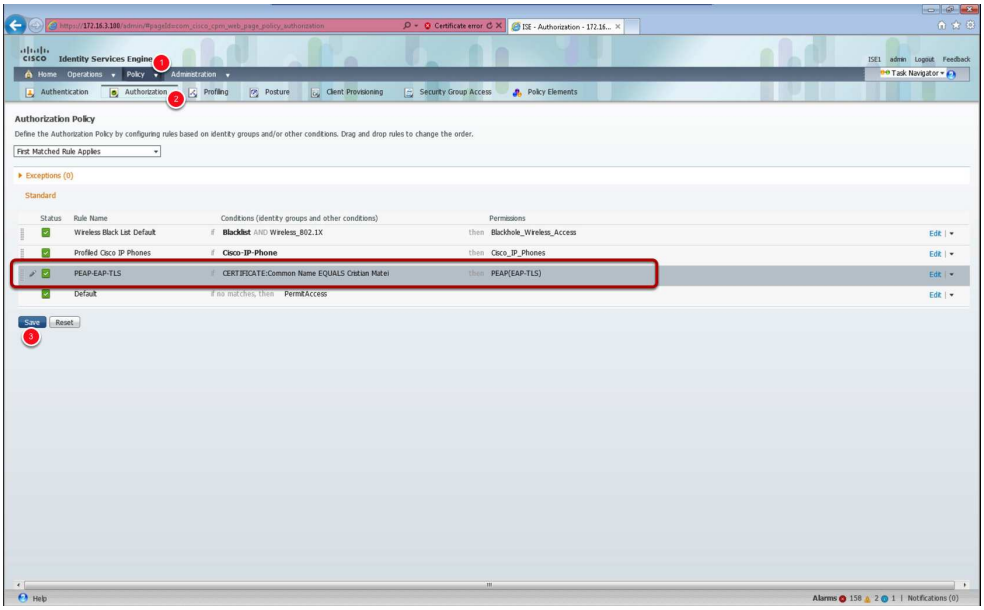


Configure the authorization profile for **TEST-PC-A**, set the RADIUS attribute 27 (Session-Timeout) to a value of 500 seconds:



Configure the authorization policy for **TEST-PC-A**, for example match on the certificate CN:

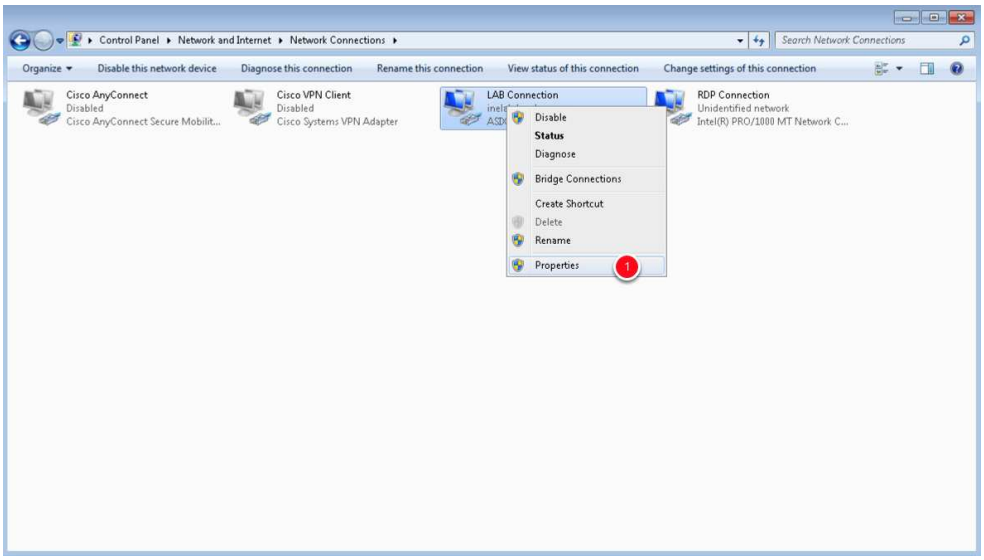
https://t.me/learningnets



If both Windows native supplicant and AnyConnect supplicant are active, AnyConnect has precedence.

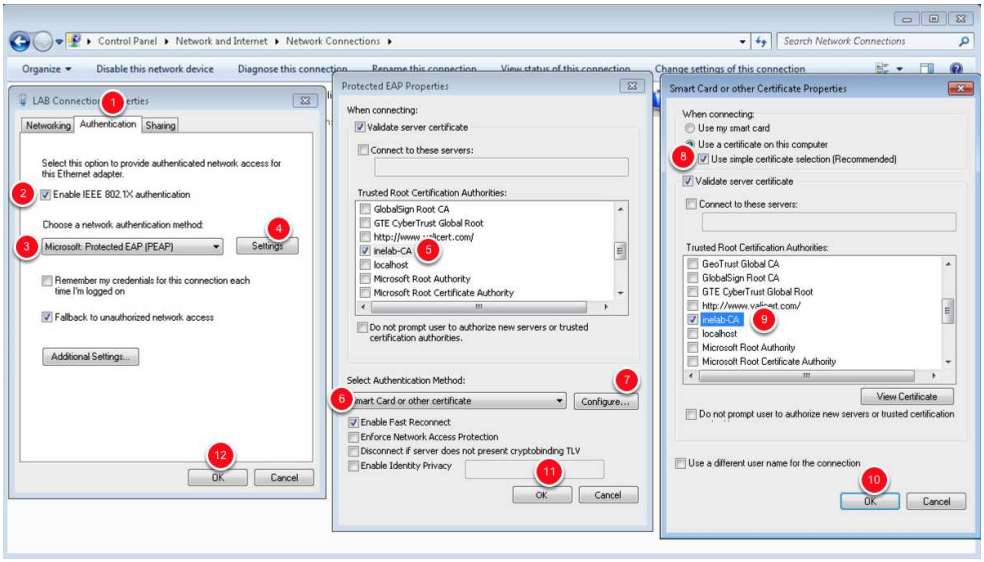
Option 1 Windows Native Supplicant

Configure the native Windows supplicant for the **LAB** Ethernet interface:

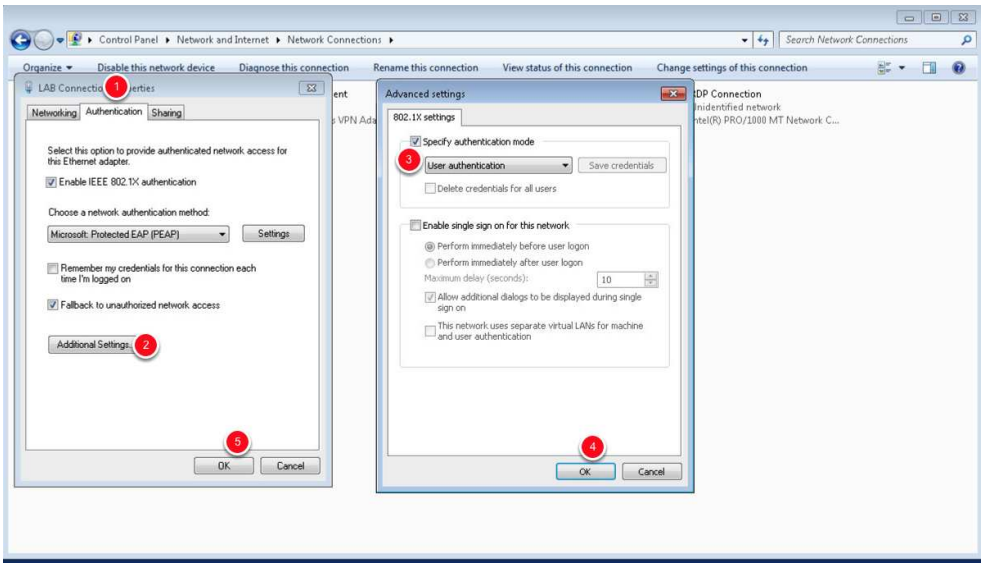


Enable 802.1x for PEAP(EAP-TLS), if the **Authentication** tab is not available, ensure the **Wired AutoConfig** service is started; in order not to be prompted trusting the ISE certificate if using the **Validate Server Certificate** option, specify which Root CA should be trusted for the ISE issued certificate:

https://t.me/learningnets



Ensure user authentication is enabled for PEAP(EAP-TLS) so that a user certificate is selected, not a computer certificate which in this case it does not even exist:



Option 2 AnyConnect Supplicant

First, enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConect 802.1x supplicant service. To access the Windows services, type **services.msc** from command prompt or run. Open AnyConnect and navigate to the **Advanced** tab in order to configure a 802.1x profile for EAP-TLS:

https://t.me/learningnets



Open AnyConnect and apply the configured profile, you'll be asked to trust the ISE certificated issued by the CA on TEST-SRV-A.

Verification

Note that the authentication process will be longer, due to the default order of doing 802.1x first (which the IP Phone can't do) and MAB second. Verify that both domains are successfully authorized, IP Phone using MAB and PC using 802.1x (note the reauthentication parameters active only for the DATA domain and received from the ISE server):

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  Unknown User-Name:  Cristian Matei
Status:  Authz Success
Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A Session timeout:  500s (server), Remaining: 125s
Timeout action:  Reauthenticate
         Idle timeout:  N/A
    Common Session ID:  8801130900000010004CA9A3
      Acct Session ID:  0x00000013
               Handle:  0xF0000010

Runnable methods list:
```

https://t.me/learningnets

```
        Method   State dot1x     Authc Success
        mab      Not run

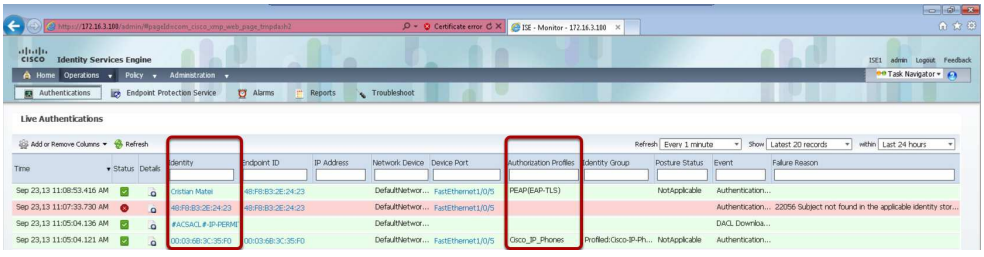

----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  VOICE
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both Authorized By:  Authentication Server
              ACS ACL:  xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  880113090000000E0047C1C6
      Acct Session ID:  0x00000011
               Handle:  0xF700000E

Runnable methods list:
        Method   State
        dot1x    Failed over         mab      Authc Success
```

Upon reauthentication, the following logs appear on the switch console:

```
%AUTHMGR-7-RESULT: Authentication result 'success' from 'dot1x' for client (48f8.b32e.2423) on Interface Fa1/0/5 Aud
%AUTHMGR-5-SUCCESS: Authorization succeeded for client (48f8.b32e.2423) on Interface Fa1/0/5 AuditSessionID 88011309
```

Verify the ISE Authentication Logs:



Verify the list of authorized clients:

https://t.me/learningnets

```
SW1#show authentication sessions


 Interface  MAC Address     Method   Domain   Status         Session ID
 Fa1/0/5    48f8.b32e.2423  dot1x    DATA     Authz Success  8801130900000010004CA9A3
 Fa1/0/5    0003.6b3c.35f0  mab      VOICE    Authz Success  880113090000000E0047C1C6
```

Verify the **authentication open** access-list as well as the dACL received from ISE for VOICE domain:

```
SW1#show ip access-lists
Extended IP access list Auth-Default-ACL-OPEN 10 permit ip any any (867 matches)
Extended IP access list xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2 (per-user)
    10 permit ip any any
!
!SW1#show ip access-lists interface fastethernet1/0/5
permit ip host 136.1.47.51 any
```

Shutdown the switchport interface (SW1 FastEthernet1/0/7) leading to ISE in order for the switch to declare the RADIUS server dead, a message similar to the following should be logged:

```
    %RADIUS-4-RADIUS_DEAD: RADIUS server 172.16.3.100:1645,1646 is not responding.
```

Enable back the switchport interface (SW1 FastEthernet1/0/7) leading to ISE in order for the switch to declare the RADIUS server alive, a message similar to the following should be logged:

```
    %RADIUS-4-RADIUS_ALIVE: RADIUS server 172.16.3.100:1645,1646 is being marked alive.
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

EAP-FAST (EAP-MSCHAPv2) Using ACS

Task

- Configure **TEST-PC-A** for 802.1x authentication using EAP-FAST; make use of the internal ACS user database.
- Ensure that your configuration does not require any server-side certificate used for EAP protocols.
- Configure the switch interface so that when a host is authenticated, all other hosts do not require authentication and are automatically allowed to pass traffic.
- Do not add any Network Devices into ACS.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

EAP-FAST was developed by Cisco as a replacement for LEAP and is very similar in its design to PEAP, which was developed by Microsoft, Cisco, and RSA. Both protocols establish a protected tunnel, and inside the tunnel the actual authentication takes place in a secure manner. To establish the protected tunnel, PEAP relies on server-side certificates, whereas EAP-FAST relies on PAC (Protected Access Credentials) files. PAC files are per-user or per-computer, depending on which entity is being authenticated. These PAC files, which essentially are child-keys generated by ACS from a master key known only by itself, must be provisioned on the clients (802.1x supplicants), and there are two methods

https://t.me/learningnets

to achieve this:

- Static/manual, out-of-band provisioning before the EAP-FAST process starts
- Dynamic/automatic, in-band provisioning during the EAP-FAST process

Because of the PAC provisioning process, EAP-FAST is known for having three phases:

- **Phase Zero** is responsible for in-band PAC provisioning and is optional if PAC files are deployed out-of-band.
- **Phase One** is responsible for establishing the protected tunnel using the PAC files.
- **Phase Two** is the actual authentication phase, where a user or computer can be authenticated.

When in-band PAC files are provisioned, the authentication server and supplicant can be configured to allow this process in an unauthenticated (**Anonymous In-Band PAC Provisioning**) or authenticated (**Authenticated In-Band PAC Provisioning**) manner. The authenticated option requires a server-side certificate issued by a CA, which the 802.1x supplicant needs to trust. Although technically the self-signed ACS certificate can be used, in real deployments a private/public CA-issued certificate is used.

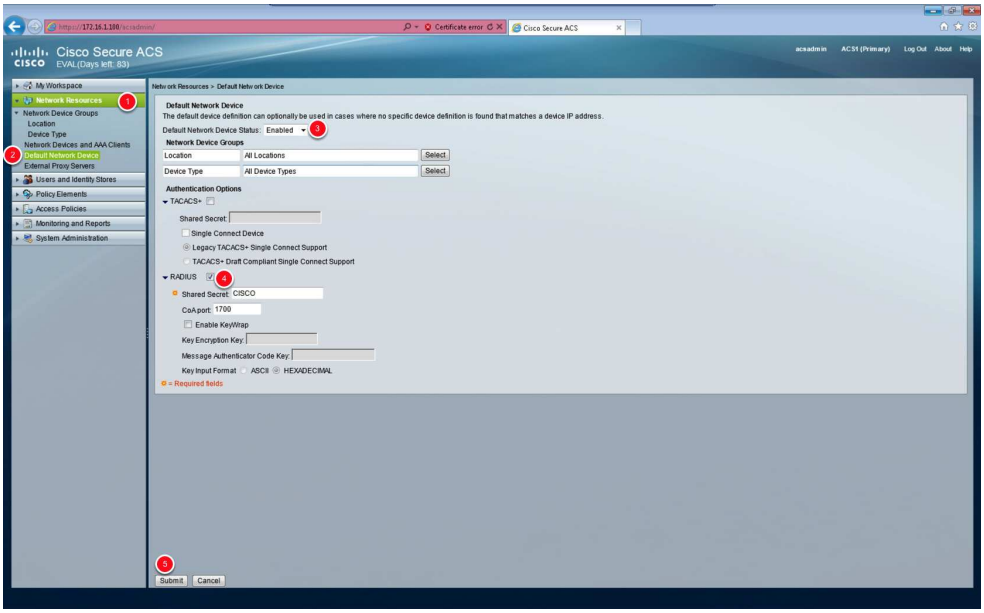
Configuration

SW1:

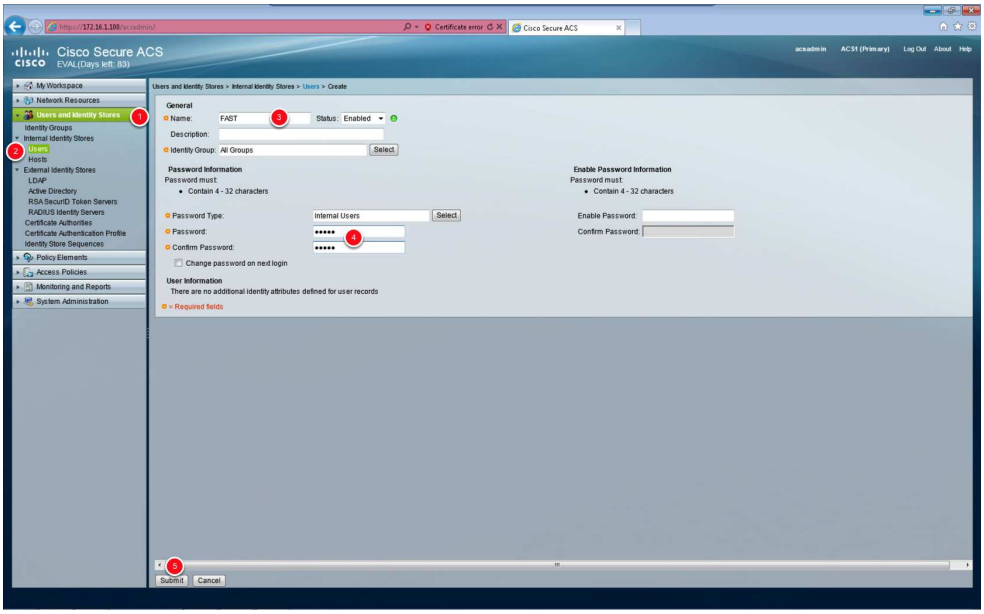
```
aaa new-model
aaa authentication dot1x default group ACS
!
dot1x system-auth-control
!
interface FastEthernet1/0/5
 authentication host-mode multi-host
 authentication port-control auto
 dot1x pae authenticator
!
aaa group server radius ACS
 server-private 172.16.1.100 key CISCO
 ip radius source-interface Loopback0
```

https://t.me/learningnets

Because no NAD's are allowed to be added to ACS, configure the **Default Network Device**:

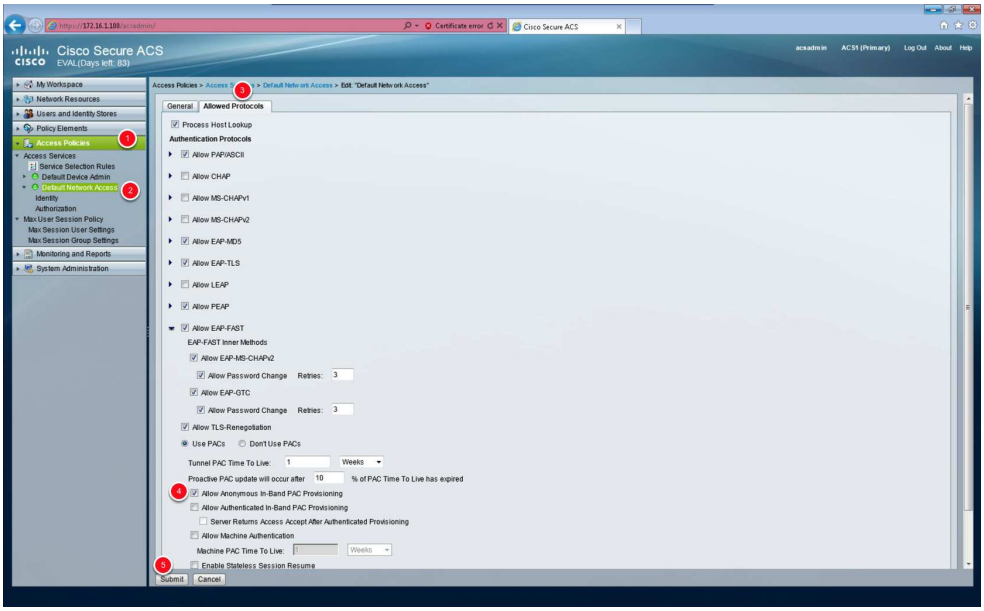


Create a user in the internal ACS store, to be used for EAP-FAST inner method authentication of EAP-MSCHAPv2:



The default authentication and authorization policies will allow EAP-FAST session to be authenticated and authorized, but the default EAP-FAST configuration must be changed to allow only **Anonymous In-Band PAC Provisioning**:

https://t.me/learningnets



Because the native Windows supplicant does not support EAP-FAST, AnyConnect must be used. First, enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConnect 802.1x supplicant service. To access the Windows services, type **services.msc** from the command prompt or Run. Open AnyConnect and navigate to the **Advanced** tab to configure an 802.1x profile for EAP-FAST:



https://t.me/learningnets

Open AnyConnect and apply the configured profile; if configuration is correct on the ACS side you will not be asked to trust the ACS self-signed certificate, because no server-side certificates are used in **Anonymous In-Band PAC Provisioning** mode. When asked by AnyConnect GUI, enter the username/password credentials configured on ACS. These credentials can be preconfigured if the EAP-FAST profile is created using the **Network Access Manager Profile Editor**.

Verification

Verify that **TEST-PC-A** was authenticated:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  Unknown User-Name:  FAST
Status:  Authz Success
Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure Oper host mode:  multi-host
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  8801130900000003501B0C6C5
      Acct Session ID:  0x0000003A
               Handle:  0xC2000035

Runnable methods list:
       Method   State dot1x    Authc Success
```

Verify the authenticated clients and also that the IP Phone was authorized:

```
SW1#show authentication sessions


Interface  MAC Address     Method   Domain   Status         Session ID
Fa1/0/5    48f8.b32e.2423  dot1x    DATA     Authz Success  8801130900000003501B0C6C5
!
!SW1#show mac address-table interface fastEthernet1/0/5
          Mac Address Table
-------------------------------------------

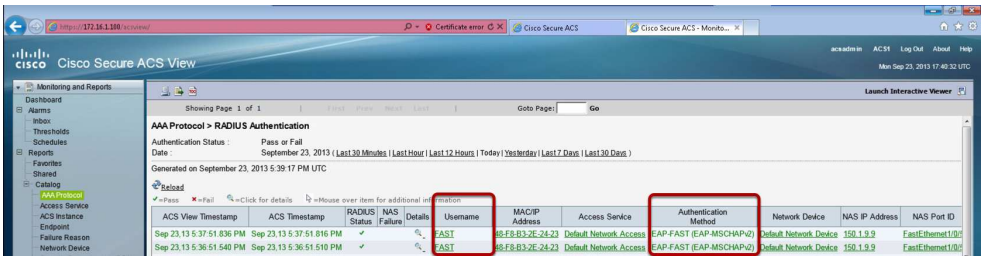
Vlan    Mac Address       Type        Ports
```

https://t.me/learningnets

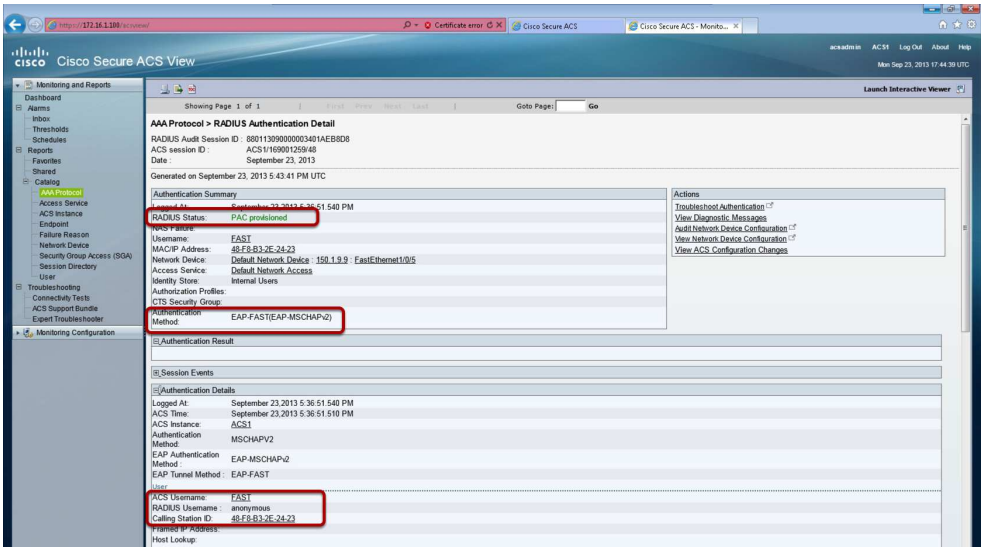
```
----    -----------       --------    ----- 27    48f8.b32e.2423    STATIC      Fa1/0/5
47    0003.6b3c.35f0    STATIC      Fa1/0/5

Total Mac Addresses for this criterion: 2
```

Verify the RADIUS Authentication Logs:

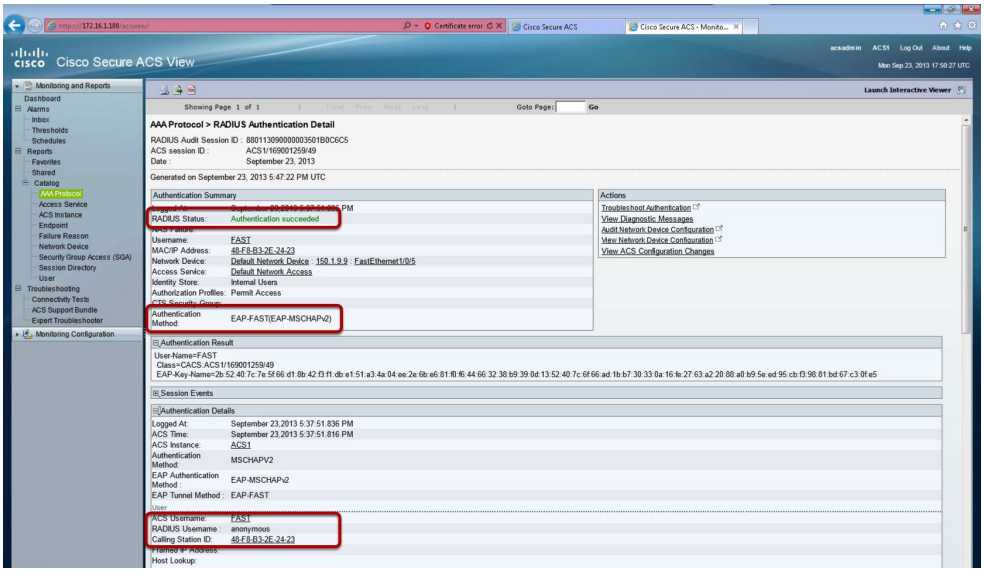


From the above, note that there are two successful authentications. First (bottom-up) is for EAP-FAST Phase zero, and it's related to PAC provisioning, which used EAP-MSCHAPv2. Click **Details** to verify this, and note that for PAC provisioning the user identity was protected and appears as **anonymous**, and there is no **Authentication Result** as expected:



With the provisioned PAC files, the secure tunnel was established in EAP-FAST Phase one, and within EAP-FAST Phase2 the actual authentication process took place using the inner method of EAP-MSCHAPv2. Click **Details** to verify this, and note that for authentication phase the user identity was protected and appears as **anonymous**:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

EAP-FAST (EAP-MSCHAPv2) Using ISE

Task

- Configure **TEST-PC-A** for 802.1x authentication using EAP-FAST; make use of the internal ISE user database.
- Ensure that the supplicant validates the authentication server presented credentials.
- Configure the switch interface so that when a host is authenticated, all other hosts do not require authentication and are automatically allowed to pass traffic.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

EAP-FAST was developed by Cisco as a replacement for LEAP and is very similar in its design to PEAP, which was developed by Microsoft, Cisco, and RSA. Both protocols establish a protected tunnel, and inside the tunnel the actual authentication takes place in a secure manner. To establish the protected tunnel, PEAP relies on server-side certificates, whereas EAP-FAST relies on PAC (Protected Access Credentials) files. PAC files are per-user or per-computer, depending on which entity is being authenticated. These PAC files, which essentially are child-keys generated by ACS from a master key known only by itself, must be provisioned on the clients (802.1x supplicants), and there are two methods to achieve this:

https://t.me/learningnets

- Static/manual, out-of-band provisioning before the EAP-FAST process starts
- Dynamic/automatic, in-band provisioning during the EAP-FAST process

Because of the PAC provisioning process, EAP-FAST is known for having three phases:

- **Phase Zero** is responsible for in-band PAC provisioning and is optional if PAC files are deployed out-of-band.
- **Phase One** is responsible for establishing the protected tunnel using the PAC files.
- **Phase Two** is the actual authentication phase, where a user or computer can be authenticated.

When in-band PAC files are provisioned, the authentication server and supplicant can be configured to allow this process in an unauthenticated (**Anonymous In-Band PAC Provisioning**) or authenticated (**Authenticated In-Band PAC Provisioning**) manner. The authenticated option requires a server-side certificate issued by a CA, which the 802.1x supplicant needs to trust. Although technically the self-signed ISE certificate can be used, in real deployments a private/public CA-issued certificate is used.

Configuration

SW1:

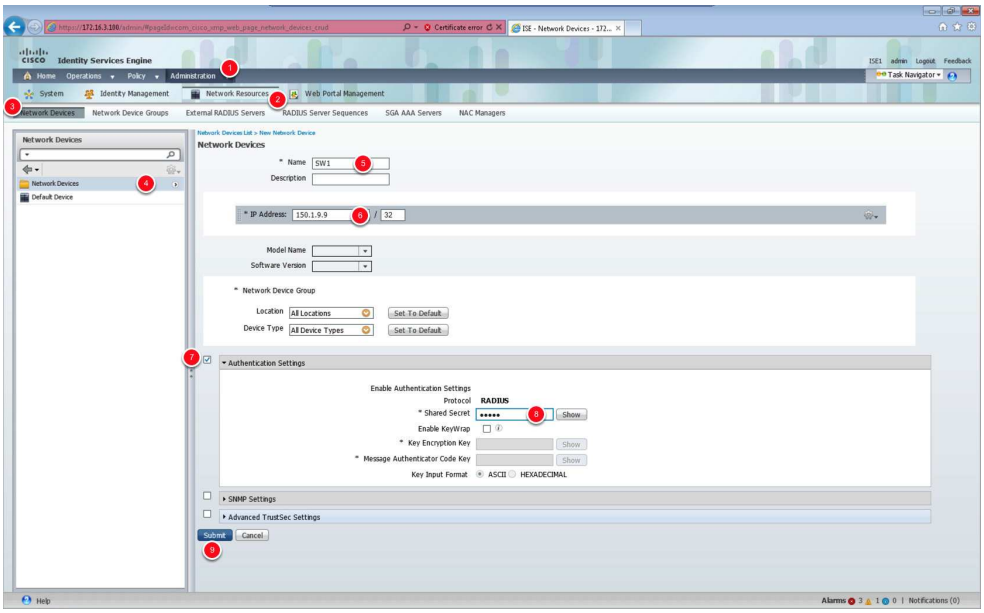
```
aaa new-model
aaa authentication dot1x default group ISE
!
dot1x system-auth-control
!
interface FastEthernet1/0/5
 authentication host-mode multi-host
 authentication port-control auto
 dot1x pae authenticator
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Loopback0
```

First, request and install a certificate from the CA server for ISE. See the Installing ISE Certificates lab . When installing the certificate, you can specify to use it only for EAP, and leave the self-signed certificate to be used for management via HTTPS. There is no need to mark the CA certificate as trusted for EAP-TLS, so that client

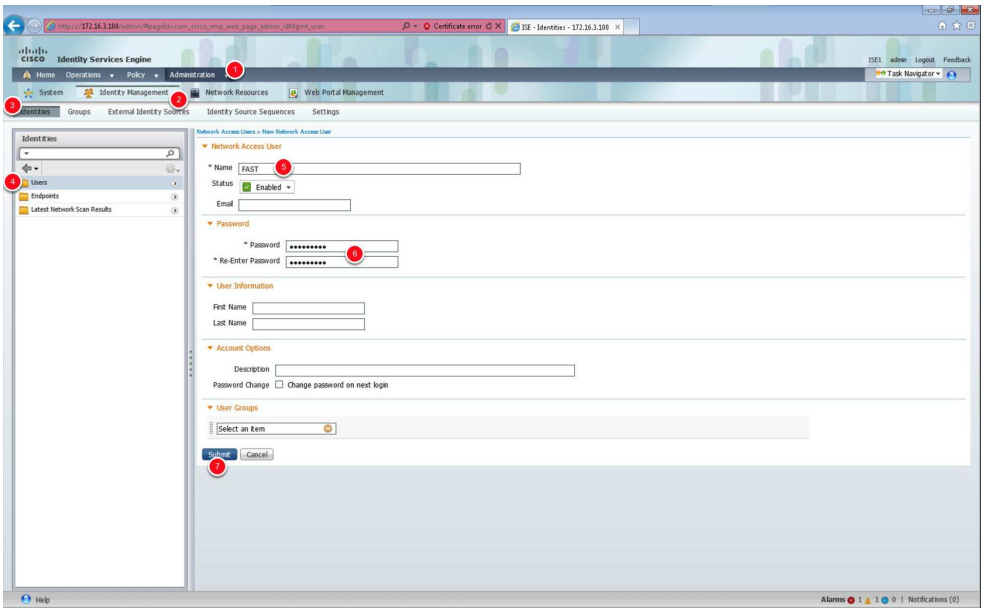
https://t.me/learningnets

certificates issued by this CA are trusted by ISE, because with EAP-FAST there is no requirement for client-side certificates.

Add SW1 as a NAD device into ISE:

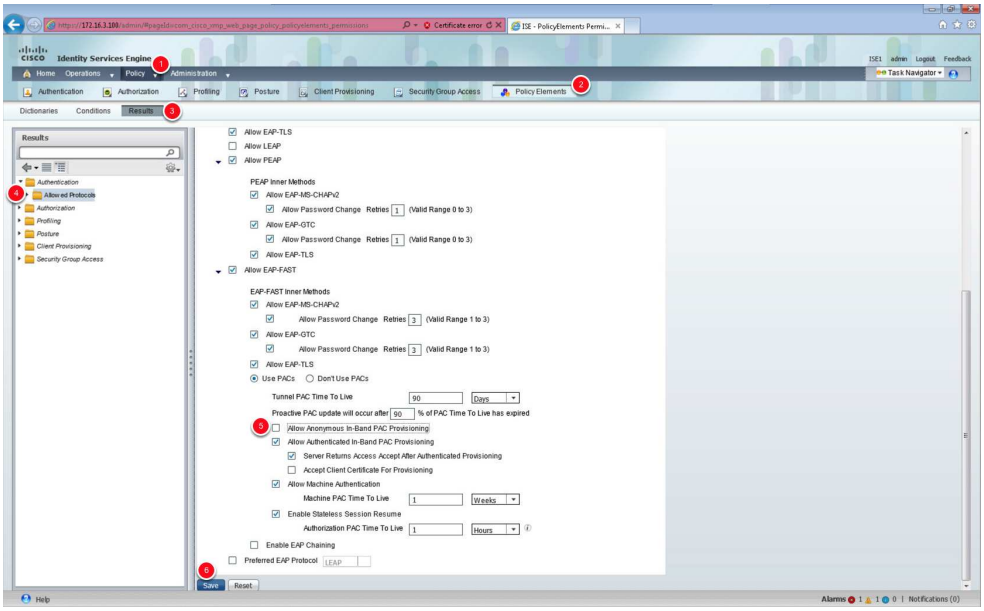


Add a user to the Internal Users identity store to be used for EAP-FAST(EAP-MSCHAPv2):



The default authentication and authorization policies will allow EAP-FAST session to be authenticated and authorized, but the default EAP-FAST configuration must be changed to allow only authenticated PAC file provisioning:

https://t.me/learningnets



Because the native Windows supplicant does not support EAP-FAST, AnyConnect must be used. First, enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConnect 802.1x supplicant service. To access the Windows services, type **services.msc** from the command prompt or Run. Open AnyConnect and navigate to the **Advanced** tab to configure an 802.1x profile for EAP-FAST:



Open AnyConnect and apply the configured profile; if configuration is correct on the ISE side, you will not be asked to trust the ISE self-signed certificate, because no server-side certificates are used in **Anonymous In-Band PAC Provisioning** mode. When asked by AnyConnect GUI, enter the username/password credentials configured on ISE. You will also be prompted to validate the ISE-presented certificate. These credentials, as well as the CA to be trusted, can be preconfigured

https://t.me/learningnets

if the EAP-FAST profile is created using the **Network Access Manager Profile Editor**.

Verification

Verify that **TEST-PC-A** was authenticated:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  Unknown User-Name:  FAST
Status:  Authz Success
Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-host
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  880113090000003A021534AC
      Acct Session ID:  0x0000003F
               Handle:  0x7600003A

Runnable methods list:
       Method   State dot1x    Authc Success
```

Verify the authenticated clients and also that the IP Phone was authorized:

https://t.me/learningnets

```
SW1#show authentication sessions

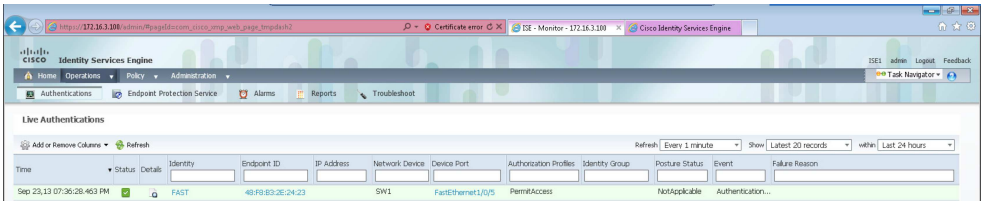

Interface  MAC Address     Method   Domain   Status         Session ID
Fa1/0/5    48f8.b32e.2423  dot1x    DATA     Authz Success  8801130900000003A021534AC
!
!SW1#show mac address-table interface fastEthernet1/0/5

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----  27    48f8.b32e.2423    STATIC      Fa1/0/5
 47    0003.6b3c.35f0    STATIC      Fa1/0/5

Total Mac Addresses for this criterion: 2
```

Verify the ISE Authentication Logs; note that ISE shows only one line compared with ACS, although if you click **Details**, all EAP-FAST Phases will be visible:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

NEAT Using ACS

Task

- Configure SW2 to authenticate its trunk port connection to SW1 using ACS.
- When the trunk is authenticated, the trunk native VLAN should be VLAN 1.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

NEAT (Network Edge Authentication Topology) further enhances the 802.1x deployments by also allowing authentication and authorization of downstream switch-to-switch ports. The downstream switch being authenticated will play the role of 802.1x supplicant, and the upstream switch will perform the regular 802.1x authenticator role. The upstream switch will have its downstream port configured as access (although trunk can be used as well), and when 802.1x authentication is successful, it will automatically convert the port as trunk; in the case of a link down event, the port is reverted back in access mode. To authorize the port as trunk, the authenticator switch must receive as part of the RADIUS authorization process the Cisco AV Pair with a value of **device-traffic-class=switch**. For the authenticator switch to process the received VSA information, the `aaa authorization network default` command must be configured.

Normally on an 802.1x enabled port, for all MAC addresses seen on the port the switch will ask for authentication before allowing it access (except when the port

https://t.me/learningnets

runs in **multi-user** mode and only one MAC needs to be authenticated). With NEAT, the host-mode is not relevant and the upstream switch assumes that the downstream switch already authenticated its access ports, so there is no need to authenticate it as well. For this reason, when the downstream switch (the supplicant switch) is authorized, it needs to signal to the upstream switch (the authenticator switch) its learned MAC addresses (or authenticated MAC addresses if it uses 802.1x on its access ports), so that the upstream switch allows traffic from these without asking for authentication. The protocol used for switch-to-switch communication to synchronize these MAC addresses is CISP (Client Information Signaling Protocol).

NEAT does not support redundant links between the authenticator and the supplicant.

If the authenticator switch has BPDU Guard enabled on its downstream port toward the supplicant switch, it is likely that the port will enter the err-disables state before actually being authenticated. To fix this, use the global command `dot1x supplicant controlled transient` on the supplicant switch, which basically blocks the port before being authenticated, so no BPDU's are being sent. Authorization received from the RADIUS server will remove the BPDU Guard configuration from the port.

Configuration

SW1:

```
cisp enable
!
dot1x credentials NEAT
 username NEAT-USER
 password NEAT-PASS
!
interface FastEthernet1/0/20
 dot1x pae supplicant
 dot1x credentials NEAT
```

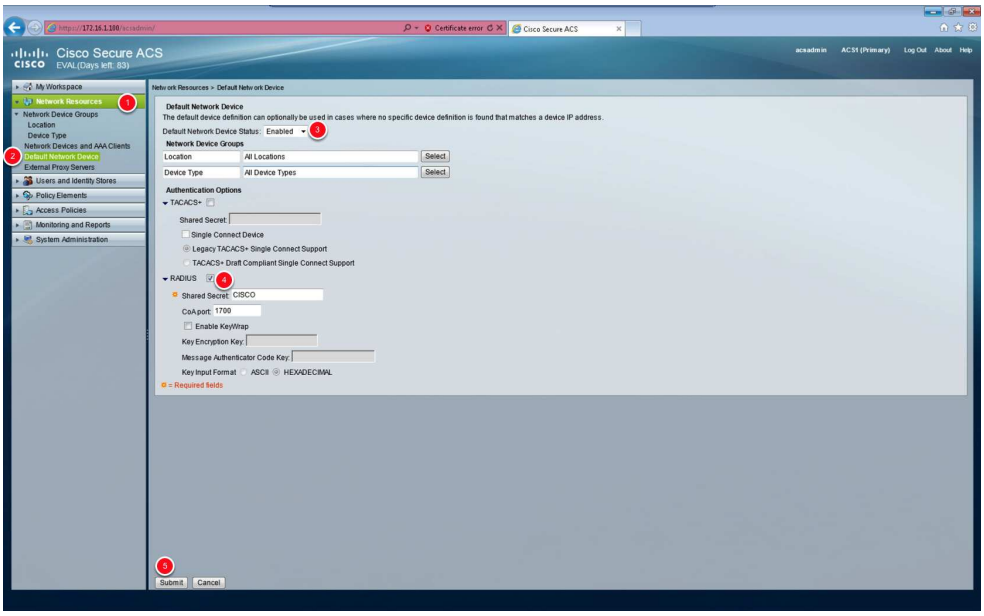
SW2:

```
aaa new-model
aaa authentication dot1x default group ACS
aaa authorization network default group ACS
!
cisp enable
dot1x system-auth-control
!
default interface FastEthernet1/0/20
!
```

https://t.me/learningnets

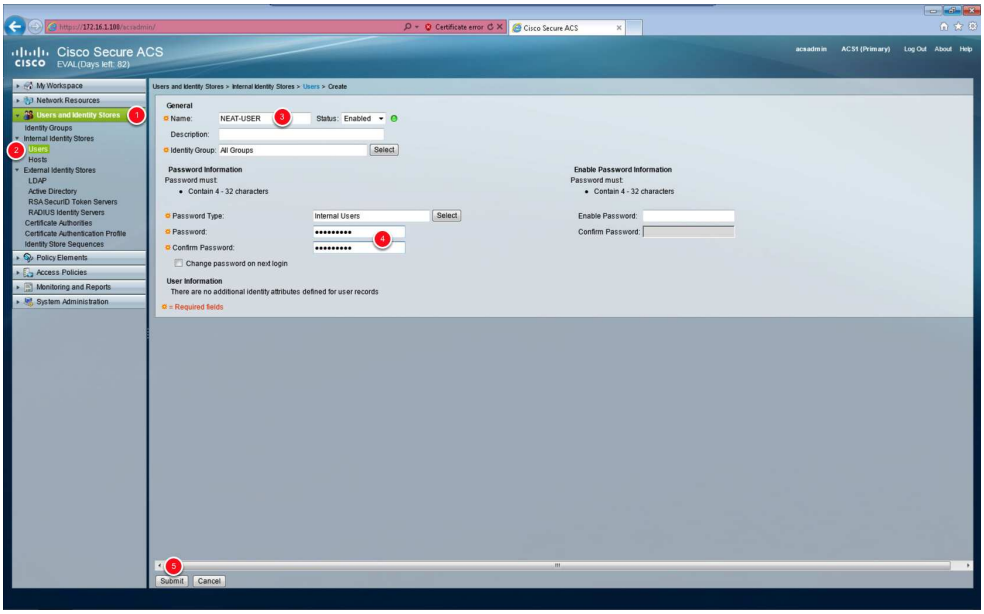
```
interface FastEthernet1/0/20
 switchport mode access
 authentication port-control auto
 dot1x pae authenticator
!
interface Vlan22
 ip address 172.16.1.10 255.255.255.0
 no shutdown
!
aaa group server radius ACS
 server-private 172.16.1.100 key CISCO
 ip radius source-interface Vlan22
```

Configure the **Default Network Device**; alternatively, configure SW2 specifically:

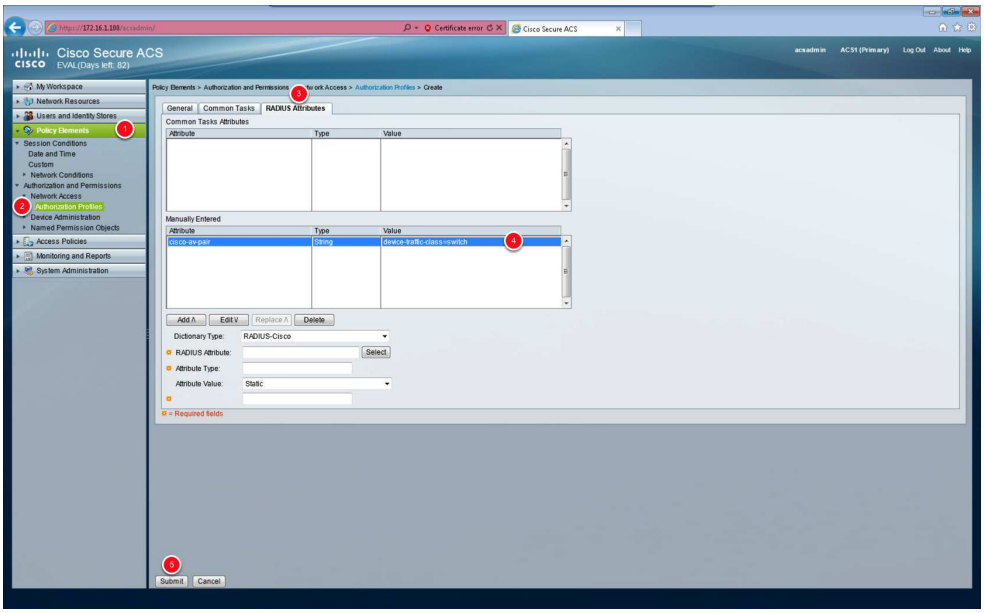


Create a user in the internal ACS store to be used by SW1, which is the 802.1x supplicant; this must match the username/password configured on SW1:

https://t.me/learningnets

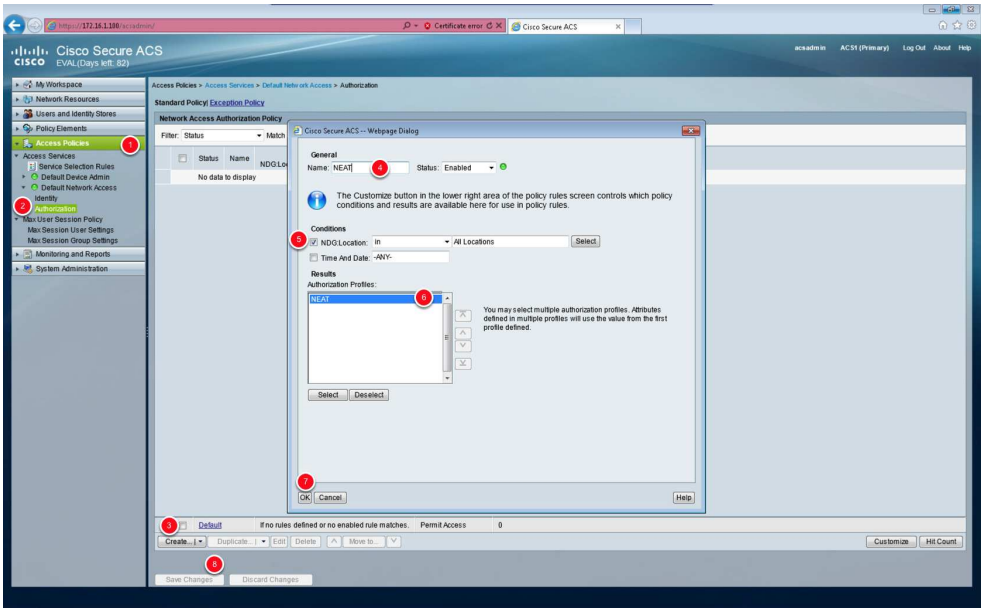


Configure the authorization profile that will authorize the trunk port by using the Cisco AV Pair with a value of **device-traffic-class=switch**:



Configure the RADIUS authorization policy to apply the previously configured profile:

https://t.me/learningnets



Verification

On successful authorization, the authenticating switch, which is SW2 in this case, will log the following commands as being applied at the interface (trunk native VLAN will be the access VLAN used before authentication):

```
Applying command... 'no switchport access vlan 1' at Fa1/0/20
Applying command... 'no switchport nonegotiate' at Fa1/0/20
Applying command... 'switchport trunk encapsulation dot1q' at Fa1/0/20
Applying command... 'switchport mode trunk' at Fa1/0/20
Applying command... 'switchport trunk native vlan 1' at Fa1/0/20
Applying command... 'spanning-tree portfast trunk' at Fa1/0/20
```

Verify that SW1 was authenticated:

```
SW2#show authentication sessions interface fastEthernet1/0/20
            Interface:  FastEthernet1/0/20
          MAC Address:  001b.d58b.2196
           IP Address:  Unknown User-Name:  NEAT-USER
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  single-host
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000000B24DAFCF8
```

https://t.me/learningnets

```
      Acct Session ID:  0x0000000C
               Handle:  0xF900000B

Runnable methods list:
       Method   State dot1x     Authc Success
```

Verify that the port is actually functioning in trunk mode:

```
SW2#show interfaces fastEthernet1/0/20 switchport | i runk
Administrative Mode: trunk
Operational Mode: trunkAdministrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1qNegotiation of Trunking: On
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: none
Trunking VLANs Enabled: ALL
!
!SW2#show spanning-tree vlan 22 interface fastEthernet1/0/20 detail
 Port 22 (FastEthernet1/0/20) of VLAN0022 is designated forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.22.
   Designated root has priority 32790, address 001b.5408.f380
   Designated bridge has priority 32790, address 001b.5408.f380
   Designated port id is 128.22, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
The port is in the portfast mode by portfast trunk configuration

   Link type is point-to-point by default
   BPDU: sent 559, received 0
```

Verify the CISP roles:

```
SW1#show cisp summary

CISP is running on the following interface(s):
---------------------------------------------- Fa1/0/20 (supplicant)

!
!SW2#show cisp summary
```

https://t.me/learningnets

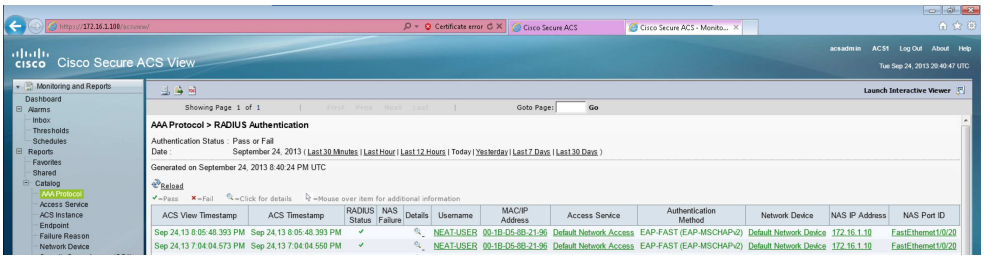
```
CISP is running on the following interface(s):
---------------------------------------------- Fa1/0/20 (authenticator)
```

Verify the authenticated MAC address learned through CISP:

```
SW2#show cisp clients


Authenticator Client Table:
---------------------------
  MAC Address     VLAN    Interface
  ---------------------------------
  d48c.b5da.d548    19    Fa1/0/20
  a493.4caa.f505    17    Fa1/0/20
  001e.f779.4770    37    Fa1/0/20
  001b.d58b.21c1    17    Fa1/0/20
  001b.d58b.21c2    19    Fa1/0/20
  001b.d58b.21c3    27    Fa1/0/20
  001b.d58b.21c4    29    Fa1/0/20
  d867.d9e0.bbc0    29    Fa1/0/20
  001b.d58b.21c5    47    Fa1/0/20
  0018.199f.161f    24    Fa1/0/20
  0018.199f.161d    26    Fa1/0/20
  0003.6b3c.35f0    27    Fa1/0/20
  0003.6b3c.35f0    47    Fa1/0/20
  000c.2937.4dc5    24    Fa1/0/20
  48f8.b32e.2423    27    Fa1/0/20
```

Verify the RADIUS Authentication Logs:



Note that when shutting down the port, the commands authorized through 802.1x are automatically removed and the port is being configured back in access mode; the switch will log the following messages:

```
Applying command... 'no switchport nonegotiate' at Fa1/0/20
Applying command... 'switchport mode access' at Fa1/0/20
```

https://t.me/learningnets

```
Applying command... 'no switchport trunk encapsulation dot1q' at Fa1/0/20
Applying command... 'no switchport trunk native vlan 1' at Fa1/0/20
Applying command... 'no spanning-tree portfast trunk' at Fa1/0/20
Applying command... 'switchport access vlan 1' at Fa1/0/20
Applying command... 'spanning-tree portfast' at Fa1/0/20
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

NEAT Using ISE

Task

- Configure SW1 to authenticate its trunk port connection to SW2 using ISE.
- When the trunk is authenticated, trunk native VLAN should be VLAN 1.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

NEAT (Network Edge Authentication Topology) further enhances the 802.1x deployments by also allowing authentication and authorization of downstream switch-to-switch ports. The downstream switch being authenticated will play the role of 802.1x supplicant, and the upstream switch will perform the regular 802.1x authenticator role. The upstream switch will have its downstream port configured as access (although trunk can be used as well), and when 802.1x authentication is successful, it will automatically convert the port as trunk; in the case of a link down event, the port is reverted back in access mode. To authorize the port as trunk, the authenticator switch must receive as part of the RADIUS authorization process the Cisco AV Pair with a value of **device-traffic-class=switch**. For the authenticator switch to process the received VSA information, the `aaa authorization network default` command must be configured.

Normally on an 802.1x enabled port, for all MAC addresses seen on the port the switch will ask for authentication before allowing it access (except when the port

https://t.me/learningnets

runs in **multi-user** mode and only one MAC needs to be authenticated). With NEAT, the host-mode is not relevant and the upstream switch assumes that the downstream switch already authenticated its access ports, so there is no need to authenticate it as well. For this reason, when the downstream switch (the supplicant switch) is authorized, it needs to signal to the upstream switch (the authenticator switch) its learned MAC addresses (or authenticated MAC addresses if it uses 802.1x on its access ports), so that the upstream switch allows traffic from these without asking for authentication. The protocol used for switch-to-switch communication to synchronize these MAC addresses is CISP (Client Information Signaling Protocol).

NEAT does not support redundant links between the authenticator and the supplicant.

If the authenticator switch has BPDU Guard enabled on its downstream port toward the supplicant switch, it is likely that the port will enter the err-disables state before actually being authenticated. To fix this, use the global command `dot1x supplicant controlled transient` on the supplicant switch, which basically blocks the port before being authenticated, so no BPDU's are being sent. Authorization received from the RADIUS server will remove the BPDU Guard configuration from the port.

Configuration

SW1:

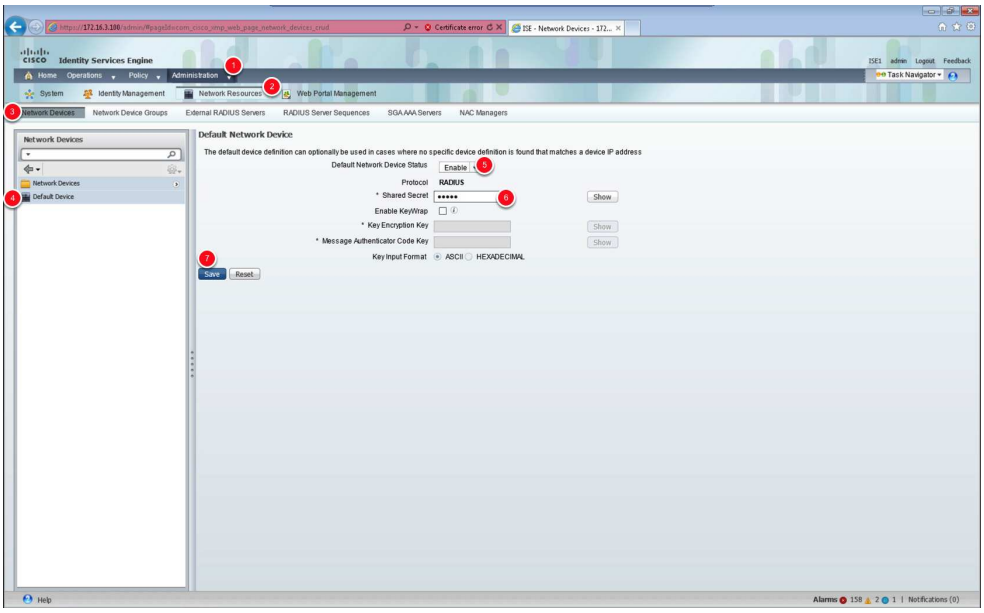
```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
cisp enable
dot1x system-auth-control
!
default interface FastEthernet1/0/20
!
interface FastEthernet1/0/20
 switchport mode access
 authentication port-control auto
 dot1x pae authenticator
!
interface Vlan24
 ip address 172.16.3.9 255.255.255.0
 no shutdown
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Vlan24
```

https://t.me/learningnets

SW2:

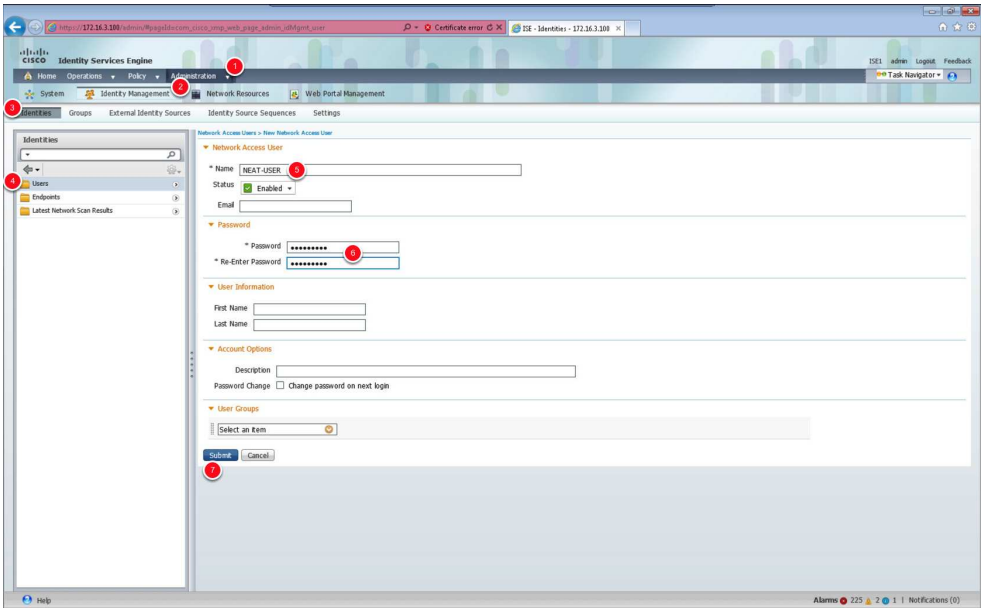
```
cisp enable
!
dot1x credentials NEAT
 username NEAT-USER
 password 12Cisco34
!
interface FastEthernet1/0/20
 dot1x pae supplicant
 dot1x credentials NEAT
```

Configure RADIUS settings for the **Default Device**; alternatively, configure SW1 specifically:

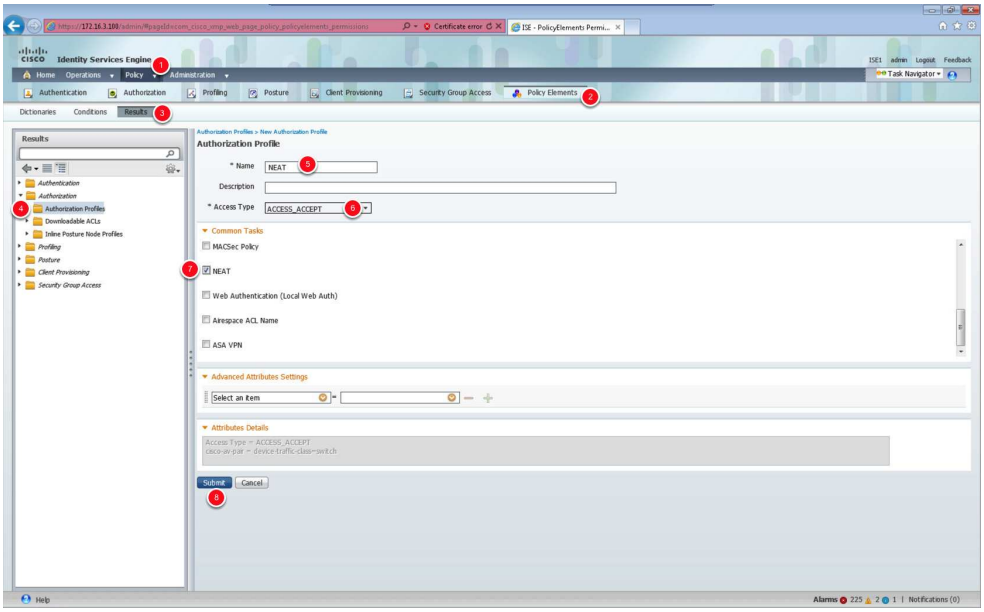


Create a user in the internal Internal Users Identity Store to be used by SW2, which is the 802.1x supplicant; this must match the username/password configured on SW2:

https://t.me/learningnets

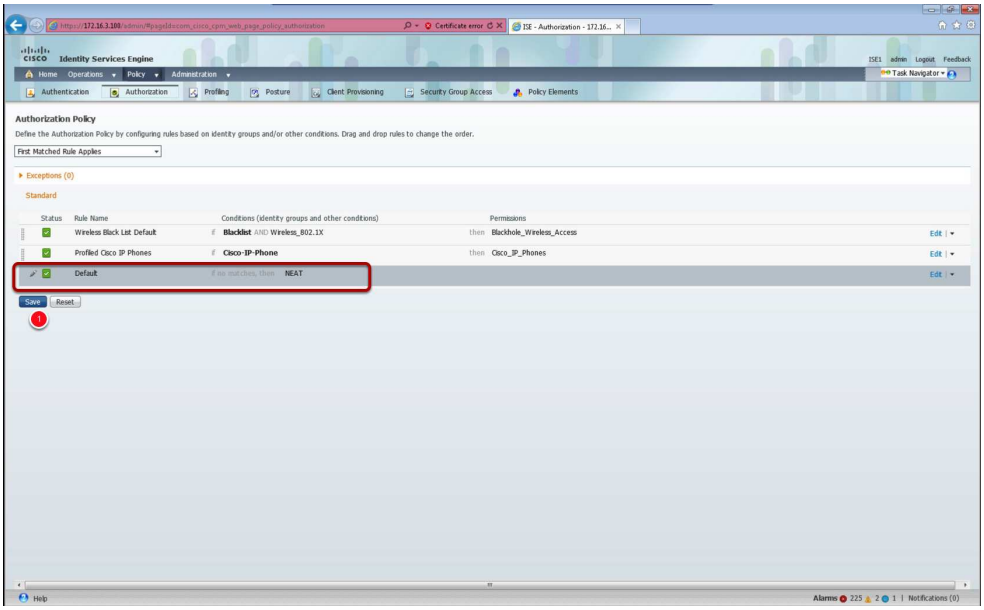


Configure the authorization profile that will authorize the trunk port by using the Cisco AV Pair with a value of **device-traffic-class=switch**; on ISE, the AV pair value is pre-configured using the **NEAT** option:



Modify the default authorization policy to apply the previously configured profile:

https://t.me/learningnets



Verification

On successful authorization the authenticating switch, which is SW2 in this case, will log the following commands as being applied at the interface (trunk native VLAN will be the access VLAN used before authentication):

```
Applying command... 'no switchport access vlan 1' at Fa1/0/20
Applying command... 'no switchport nonegotiate' at Fa1/0/20
Applying command... 'switchport trunk encapsulation dot1q' at Fa1/0/20
Applying command... 'switchport mode trunk' at Fa1/0/20
Applying command... 'switchport trunk native vlan 1' at Fa1/0/20
Applying command... 'spanning-tree portfast trunk' at Fa1/0/20
```

Verify that SW2 was authenticated:

```
SW1#show authentication sessions interface fastEthernet1/0/20
            Interface:  FastEthernet1/0/20
          MAC Address:  001b.5408.f396
           IP Address:  Unknown User-Name:  NEAT-USER
Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  single-host
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000003001F4DAD
```

https://t.me/learningnets

```
      Acct Session ID:  0x00000004
               Handle:  0x2C000003

Runnable methods list:
       Method   State dot1x     Authc Success
```

Verify that the port is actually functioning in trunk mode:

```
SW1#show interfaces fastEthernet1/0/20 switchport | i runk
Administrative Mode: trunk
Operational Mode: trunkAdministrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1qNegotiation of Trunking: On
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: none
Trunking VLANs Enabled: ALL
!
!SW1#show spanning-tree vlan 22 interface fastEthernet1/0/20 detail
 Port 22 (FastEthernet1/0/20) of VLAN0022 is designated forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.22.
   Designated root has priority 16406, address 001b.d58b.2180
   Designated bridge has priority 16406, address 001b.d58b.2180
   Designated port id is 128.22, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
The port is in the portfast mode by portfast trunk configuration

   Link type is point-to-point by default
   BPDU: sent 21107, received 0
```

Verify the CISP roles:

```
SW1#show cisp summary

CISP is running on the following interface(s):
---------------------------------------------- Fa1/0/20 (authenticator)
!
!SW2#show cisp summary

CISP is running on the following interface(s):
```

https://t.me/learningnets

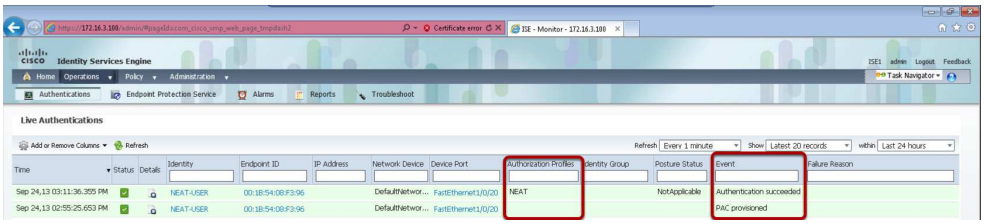
```
---------------------------------------------- Fa1/0/20 (supplicant)
```

Verify the authenticated MAC address learned through CISP:

```
SW1#show cisp clients


Authenticator Client Table:
---------------------------
  MAC Address     VLAN    Interface
  ---------------------------------
  d48c.b5da.d549    26    Fa1/0/20
  a493.4caa.f508    37    Fa1/0/20
  d48c.b5da.d549     1    Fa1/0/20
  0018.199f.1620    22    Fa1/0/20
  001b.5408.f3c1    19    Fa1/0/20
```

Verify the Authentication Logs and notice EAP-FAST Phases zero and two:



Note that when shutting down the port, the commands authorized through 802.1x are automatically removed and the port is being configured back in access mode; the switch will log the following messages:

```
Applying command... 'no switchport nonegotiate' at Fa1/0/20
Applying command... 'switchport mode access' at Fa1/0/20
Applying command... 'no switchport trunk encapsulation dot1q' at Fa1/0/20
Applying command... 'no switchport trunk native vlan 1' at Fa1/0/20
Applying command... 'no spanning-tree portfast trunk' at Fa1/0/20
Applying command... 'switchport access vlan 1' at Fa1/0/20
Applying command... 'spanning-tree portfast' at Fa1/0/20
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ISE Profiling

Task

- Configure SW1 and ISE for device profiling using DHCP and RADIUS.
- Configure ISE with the following two authorization rules:
  - All devices profiled as **Cisco-IP-Phone** should be allowed all IP traffic and VOICE vlan access.
  - All other devices using MAB should be allowed only for ICMP and TCP traffic and VOICE vlan access.
- Before authentication, all traffic should be allowed.
- Ensure that **TEST-PC-A** is authorized for the guest VLAN, which is VLAN 27.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

ISE can profile devices and use this information for making more device-relevant authorization policies. For example, it can give Cisco IP Phones one authorization policy and Avaya IP Phones a different authorization policy. Because both devices will probably use MAB, without the ability to profile the devices, it would be more difficult to create differentiated and scalable authorization policies. ISE can profile by the following mechanisms:

- **Netflow** can be used to detect whether the device has been misclassified, based on

https://t.me/learningnets

its traffic (source/destination IP and ports, protocol used).
- **DHCP** uses DHCP attributes for end host operating system fingerprinting.
- **HTTP** can be used to detect the end host operating system based on the user-agent presented in the HTTP header.
- **RADIUS** can use attributes from both authentication and accounting; the most relevant is the MAC address from the **Calling-Station-ID** attribute and IP address from the **Framed-IP-Address** attribute. If the NAD supports the device sensor functionality, further information can be gathered by the NAD through CDP/LLDP/DHCP.
- **NMAP** can detect end host operating systems and type through TCP/UDP port scanning.
- **DNS** requires knowledge of the IP address of the end host and existing PTR records on the DNS server, and works by identifying the FQDN of the end host.
- **SNMP** works with traps received from the switch and queries sent to the switch when a device is authenticated; it uses CDP/LLDP, MAC address, and ARP table data.

In the regular RADIUS 802.1x authentication process, the NAD (switch) is the one initiating it through an **Access-Request** message as soon as the link UP is detected or frames are seen from non-authenticated MAC addresses. If the RADIUS server is configured correctly, it will reply with an **Access-Accept** message and possible authorization attributes. CoA (Change of Authorization) is an addition to RADIUS, defined in RFC 3576, which allows the authentication server to initiate a RADIUS action at any point in time for any previously authenticated devices. CoA uses UDP port 1700 by default, and ACS does not support dynamic authorization; it support CoA, but it needs the CoA to be manually triggered by the ACS administrator. CoA is required when device authorization must be changed as a result of profiling or device posture.

The RADIUS server can instruct the NAD (switch) to take actions on a specific device based on one or more attributes:

- Acct-Session-Id (IETF attribute 44)
- Audit-Session-Id (Cisco VSA)
- Calling-Station-Id (IETF attribute 31, which contains the host MAC address)
- Plain IP Address (IETF attribute 8)

Of the multiple possible actions the RADIUS server can instruct the switch to do, only two are relevant for 802.1x, with implementation through Cisco AV Pair attributes:

- **Bounce the port** works only for when the host-mode is **single-host** and uses a

https://t.me/learningnets

value of **subscriber:command=bounce-host-port**.

- **Reauthenticate the device** works for any host-mode and uses any of the following values: **subscriber:command=reauthenticate**, **subscriber:reauthenticate-type=last**, **subscriber:reauthenticate-type=rerun**

In our case, the IP Phone will first match the **MAB** rule and receive an authorization from ISE. As soon as ISE will profile the device as a Cisco IP Phone, it will signal the switch to perform a reauthentication and now the IP Phone will receive a different authorization from ISE.

The command `radius-server attribute 6 on-for-login-auth` configures the NAD to send the RADIUS **Service-Type** attribute in its **Access-Request** messages toward the RADIUS server, so that the RADIUS server can, for example, differentiate between a MAB and 802.1x request. This command is optional, because the switch sends the **Service-Type** attribute by default.

The command `radius-server attribute 8 include-in-access-req` configures the NAD to send the supplicant's IP address in the **Access-Request** RADIUS message as a value in the **Framed-IP-Address** attribute. This attribute can be used in RADIUS authorization policies, or it can be used by profiling services that need to know the IP address of the supplicant such as DNS, Netflow, and NMAP. The IP address of the supplicant can also be sent in RADIUS accounting packets, if the `radius-server vsa send accounting` command is configured. The IP address of the supplicant sent in RADIUS Accounting packets can only be learned if the IP address is assigned via DHCP, or if `ip device tracking` is configured on the switch.

The command `radius-server attribute 25 access-request include` is optional as well and adds the RADIUS class attribute in the **Access-Request** message.

Configuration

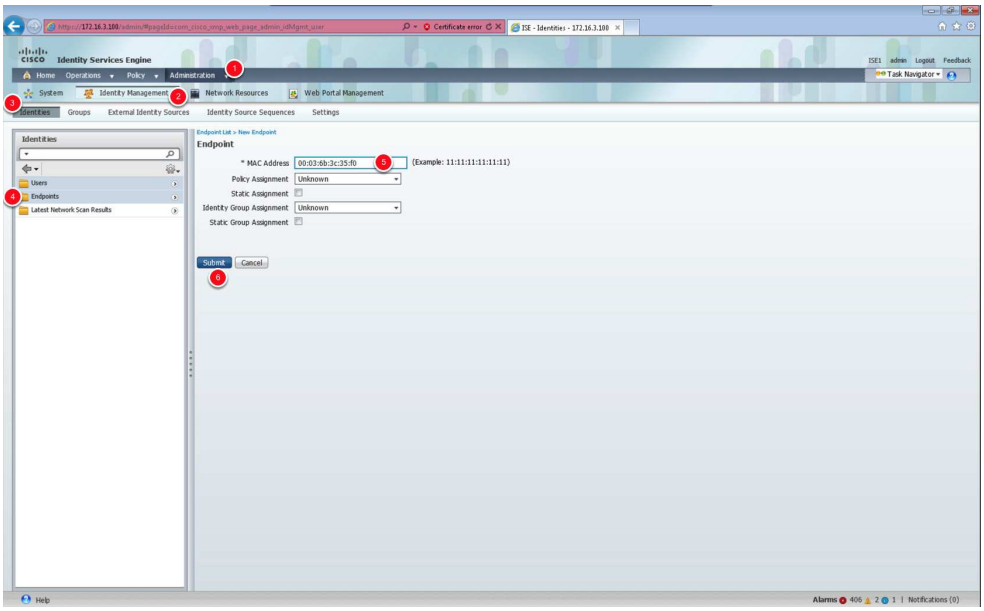
SW1:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
ip device tracking
dot1x system-auth-control
radius-server vsa send authentication
!
interface FastEthernet1/0/5
 authentication port-control auto
 authentication host-mode multi-domain
```

https://t.me/learningnets

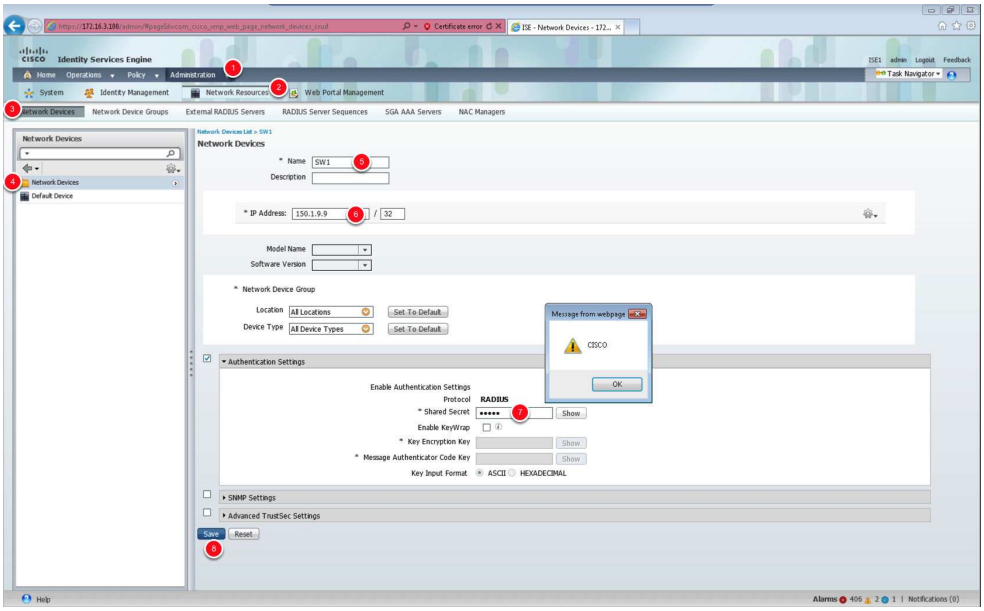
```
 authentication event no-response action authorize vlan 27
 authentication open
 dot1x pae authenticator
 mab
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Loopback0
!
aaa server radius dynamic-author
 client 172.16.3.100 server-key CISCO
!
interface Vlan47
 ip helper-address 172.16.3.100
```

First add the MAC address of the IP Phone in the ISE Endpoint store; do not statically profile it:

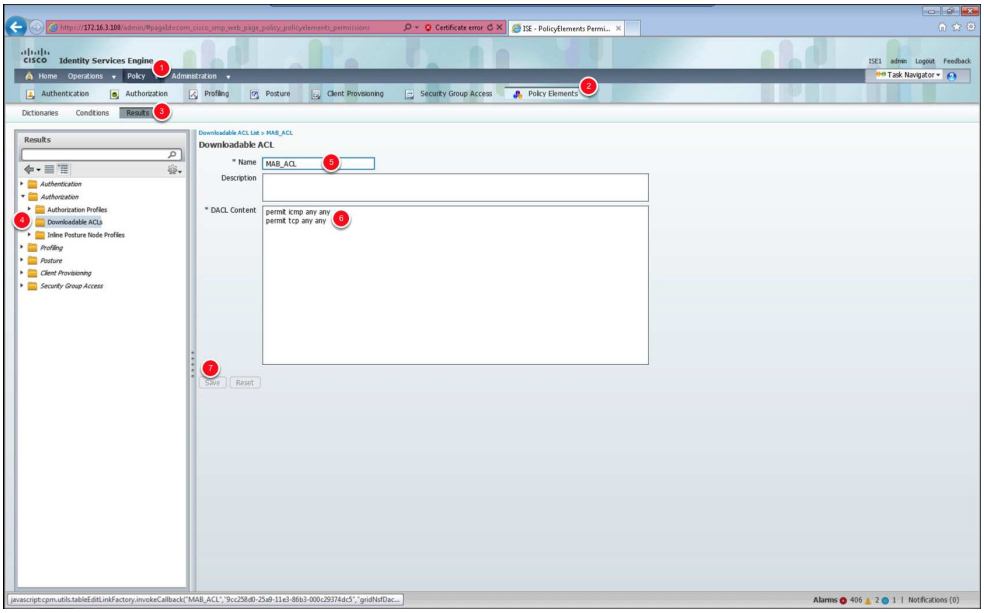


Add SW1 as a NAD into ISE, and note there are no specific settings for CoA:

https://t.me/learningnets

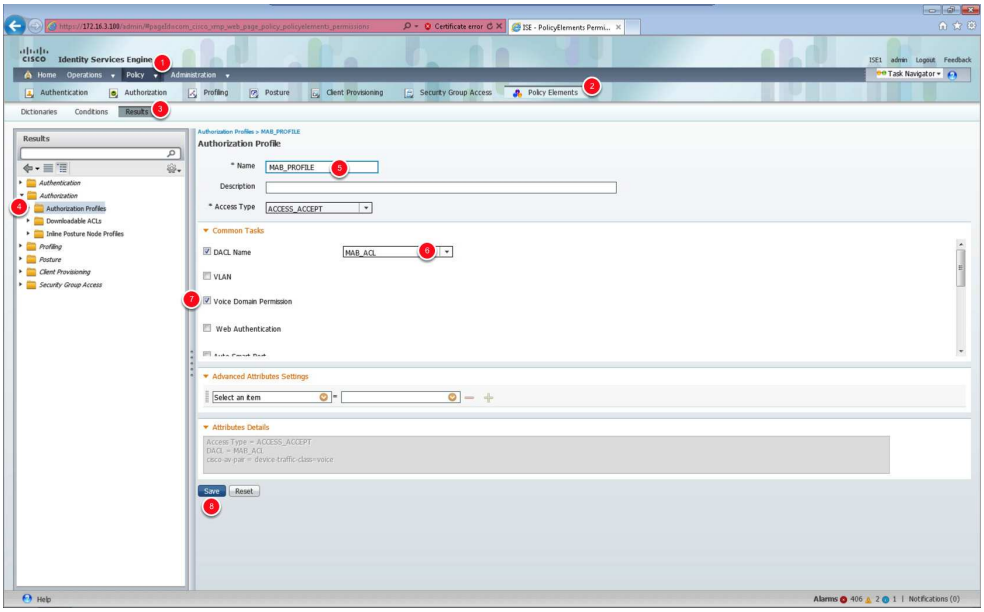


Configure a dACL to be used for the non-Cisco IP Phone devices authorization profile:

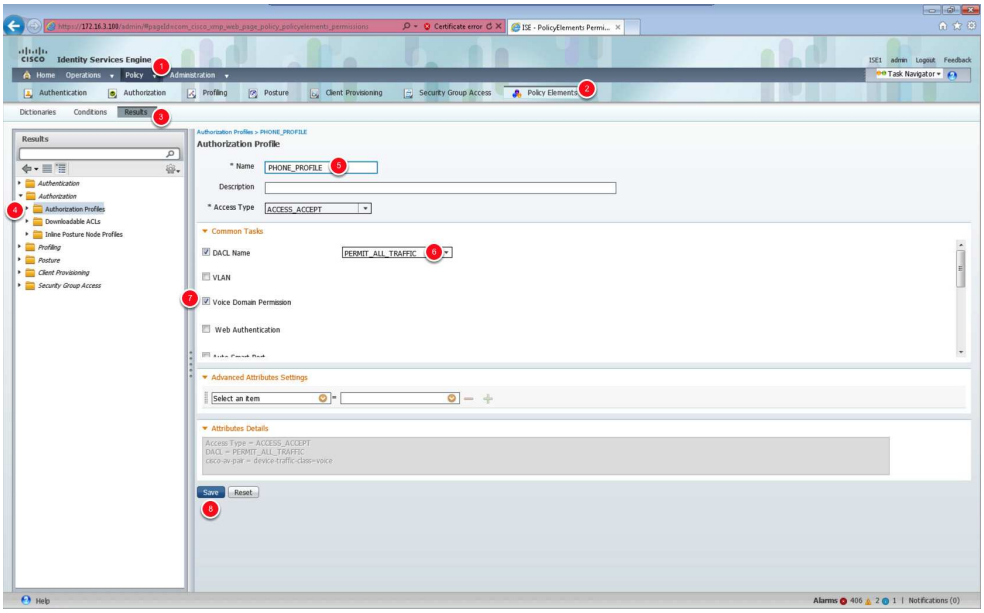


Configure the MAB authorization profile:

https://t.me/learningnets

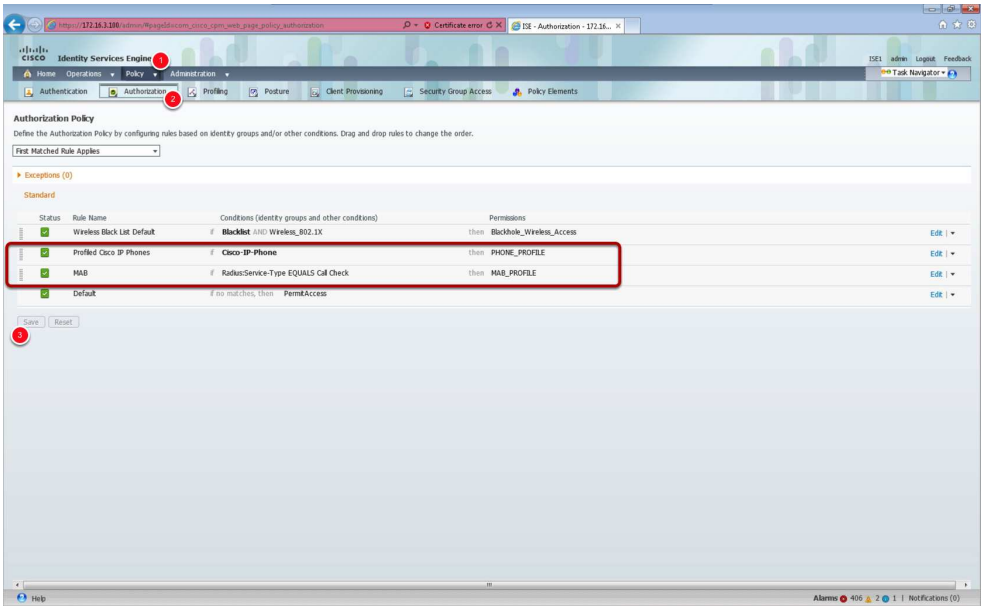


Configure the IP Phone authorization profile:

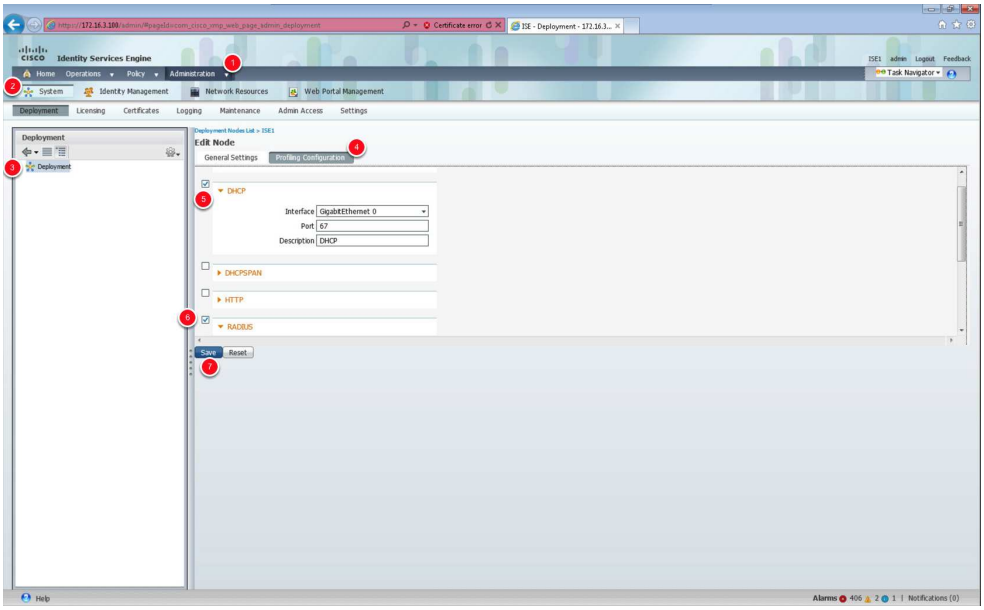


Configure the authorization policies for both MAB and IP Phone, applying the previously created profiles; because authorization policies are processed top down, just like an ACL, make sure that the MAB policy is after the IP Phone policy:

https://t.me/learningnets

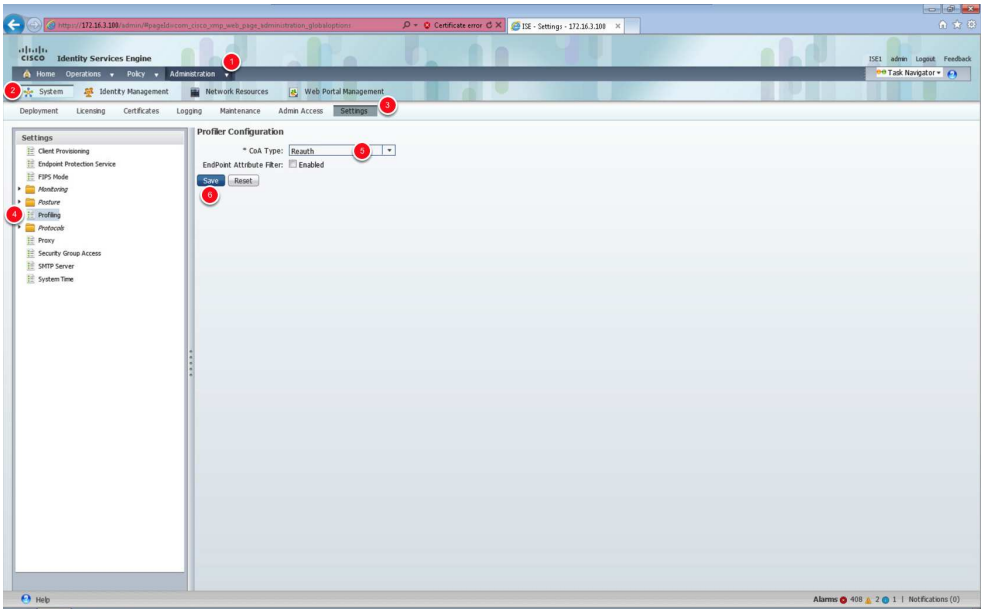


Activate DHCP and RADIUS profiling:



Configure the CoA action as reauthentication; by default, no action is taken:

https://t.me/learningnets



Verification

Bring the port up before activating DHCP and RADIUS Profiling, and verify the received authorization policy; also note that the TEST-PC-A was authorized in the guest VLAN because its 802.1x supplicant is not configured:

```
SW1#show authentication sessions interface fastethernet1/0/5

            Interface:  FastEthernet1/0/5
          MAC Address:  Unknown
           IP Address:  Unknown User-Name:  UNRESPONSIVE
Status:  Authz Success
Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-host
     Oper control dir:  both Authorized By:  Guest Vlan
Vlan Policy:  27
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000002C035BA083
      Acct Session ID:  0x0000002D
               Handle:  0x5100002C

Runnable methods list:
       Method   State
       mab      Failed over

----------------------------------------
```

https://t.me/learningnets

```
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  136.1.47.51 User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  VOICE
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both Authorized By:  Authentication Server
ACS ACL:  xACSACLx-IP-MAB_ACL-52427f10
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000002D035C025B
      Acct Session ID:  0x0000002E
               Handle:  0xB300002D

Runnable methods list:
       Method   State mab        Authc Success
```

Verify the access-list received from ISE:

```
SW1#show ip access-lists xACSACLx-IP-MAB_ACL-52427f10
Extended IP access list xACSACLx-IP-MAB_ACL-52427f10 (per-user) 10 permit icmp any any
20 permit tcp any any
!
!SW1#show ip access-lists interface fastEthernet1/0/5
permit icmp host 136.1.47.51 any
permit tcp host 136.1.47.51 any
```

Activate the RADIUS and DHCP profiling on ISE, which should trigger a reauthentication, and a new authorization policy should be received from ISE:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  Unknown
           IP Address:  136.1.27.51
            User-Name:  UNRESPONSIVE
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-host
     Oper control dir:  both
```

https://t.me/learningnets

```
        Authorized By:  Guest Vlan
          Vlan Policy:  27
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000002E03838897
      Acct Session ID:  0x0000002F
              Handle:  0xDD00002E

Runnable methods list:
       Method   State
       mab      Failed over

---------------------------------------

            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  136.1.47.51 User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  VOICE
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-domain
     Oper control dir:  both
        Authorized By:  Authentication Server ACS ACL:  xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000002F0383EA5F
      Acct Session ID:  0x00000030
              Handle:  0x5300002F

Runnable methods list:
       Method   State mab       Authc Success
```

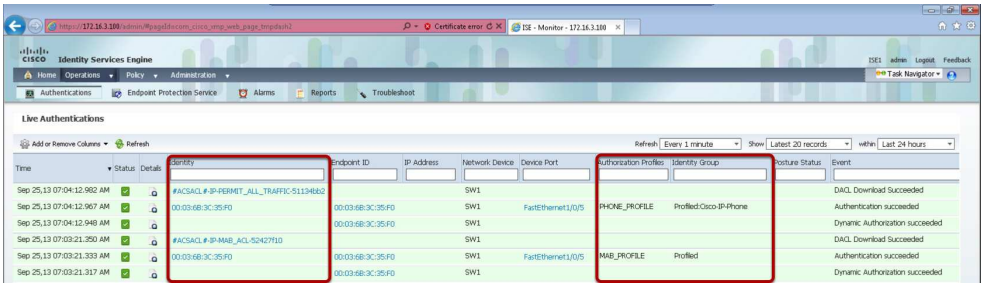
Verify the access-list received from ISE:

```
SW1#show ip access-lists xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2
Extended IP access list xACSACLx-IP-PERMIT_ALL_TRAFFIC-51134bb2 (per-user) 10 permit ip any any
!
!SW1#show ip access-lists interface fastEthernet1/0/5
permit ip host 136.1.47.51 any
```

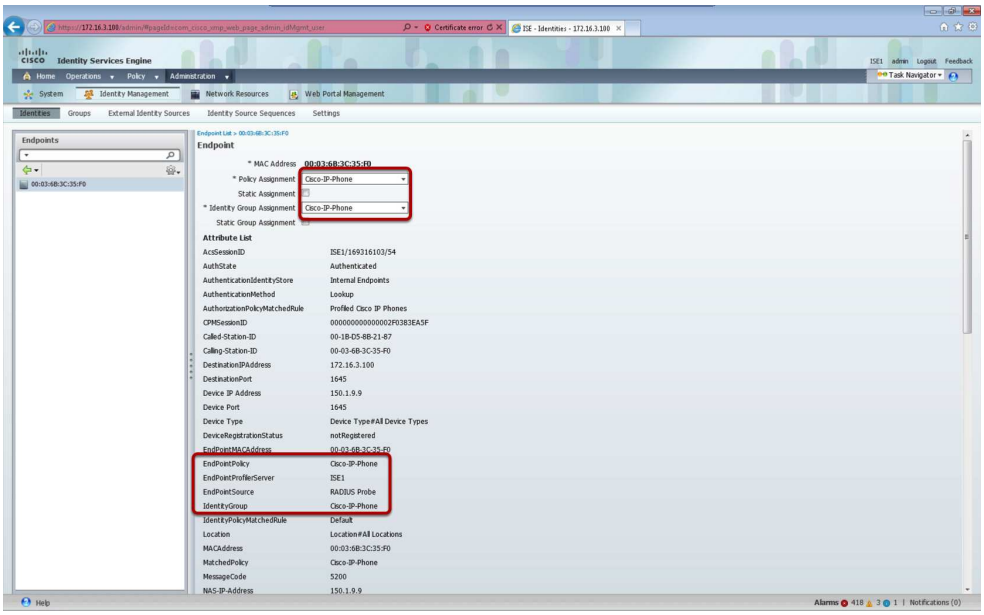
Verify the authentication logs from ISE by looking bottom-up; notice that first the **MAB**

https://t.me/learningnets

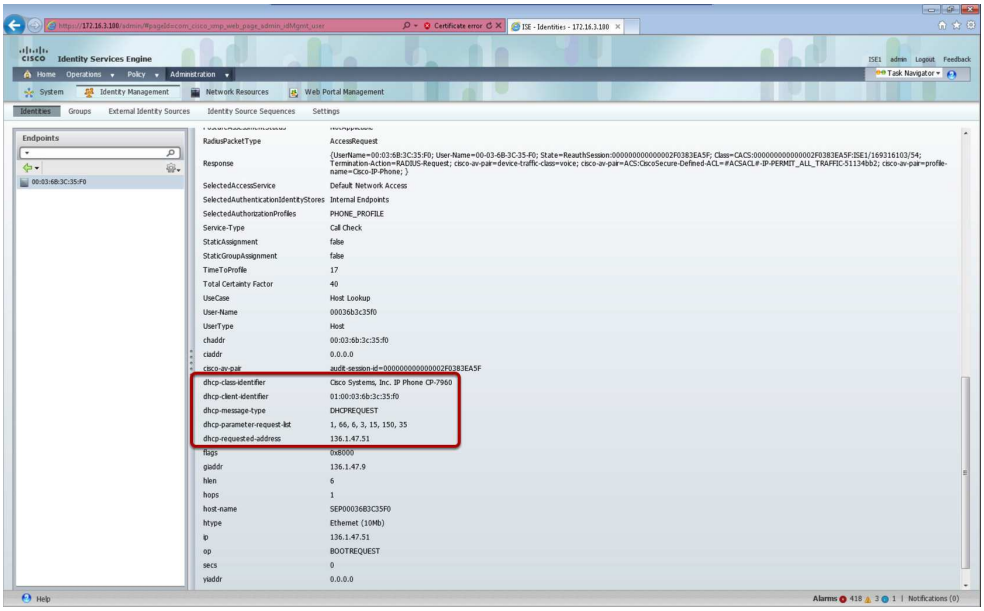
authorization profile was matched, the profiling happened automatically, and the authorization was changed to **PHONE**:



Verify that the IP Phone was correctly profiled, and check the RADIUS data:



Verify that the IP Phone was correctly profiled, and check the DHCP data:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ISE Central Web Authentication

Task

- Configure SW1 and ISE to authorize **TEST-PC-A** using Central Web Authentication.
- When authorized, **TEST-PC-A** should be allowed all ICMP, TCP, and UDP traffic.
- IP Phone should not be authorized to join the VOICE VLAN.
- Do not allow users to self-provision their devices into ISE database.
- Ensure that SW1's HTTP/HTTPS service are not available for managing it.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The **Central Web Authentication** makes use of the ISE Web Portal and works as follows:

- Because we want to use CWA, we cannot authenticate the client because either it is not 802.1x compliant, or MAB can't work because its MAC address is not known to ISE.
- For this reason, we must configure the matching authentication policy on ISE with an action of **Continue** if the user/MAC is unknown, so that ISE can allow the device to authenticate via CWA even though MAC is unknown.
- We configure an authorization policy on ISE specifying that if the user is unknown, to

https://t.me/learningnets

send authorization to the switch to force the client to use Central Web Authentication (the key is correct configuration of dACL and Redirect ACL, both being mandatory unless the port already has a pre-authentication access-list applied on the port, in which case only Redirect ACL is required to be configured on ISE).
- When the user successfully authenticates in the Guest Portal by CWA, ISE will attach an invisible flag for this session-id, remembering that it was a guest authenticated user through web portal.
- ISE will automatically trigger a CoA event and signal the switch to reauthenticate the device based on the audit-session-id (note that the CoA request is acknowledged by the switch, so a response is sent to ISE before reauthentication).
- When the switch triggers the reauthentication, ISE will match the same authentication policy, but now it needs to match on another authorization policy using a specific condition of **Network Access Use Case Equals GuestFlow**; this condition will actually look for that flag and see if it was set for this session-id.
- Finally, the switch will receive the new authorization policy and apply it for the respective session-id.

CWA is basically a regular 802.1x authentication, but the initial authorization received from the RADIUS server forces the user into web authentication using ISE web portal. ACS does not support this functionality.

The dACL used for the initial authorization must allow all required user traffic, such as DHCP, DNS and access to ISE Guest Portal on the default port of 8443. Without an IP address and DNS name resolution for the URL being redirected, the client will not be able to access ISE Guest Portal. HTTP and HTTPS traffic must also be allowed so that the switch can intercept it based on the rules in the Redirect ACL. The Redirect ACL is sent by ISE to the switch in the format of a Cisco AV Pair, but just the name of the ACL is sent, so the ACL itself must be configured at the switch level. The logic of the Redirect ACL for the switch (on the Wireless LAN Controller is the opposite logic) is that all denied traffic is not intercepted, and all permitted traffic is intercepted; therefore, we want to permit all HTTP and HTTPS traffic only. For CWA to work, the user must manually open the web browser and initiate an HTTP/HTTPS session.

On SW1's FastEthernet1/0/5 there are two hosts, the IP Phone and the TEST-PC-A. The IP Phone learns the VOICE VLAN from the switch using CDP version 2 advertisements. However, when it boots, it will send untagged traffic (before it learns the VOICE VLAN), so the switch will learn its MAC address in the DATA VLAN first. Based on task requirements, the IP Phone is not allowed to join the VOICE VLAN and it will remain in the DATA VLAN, which results in the DATA VLAN having two MAC addresses. Having more than one MAC address in the DATA VLAN for an

https://t.me/learningnets

802.1x enabled port is allowed only in host-mode of **multi-auth** or **multi-host**; otherwise, a violation occurs the port enters the **err-disabled** state by default.

The `ip device tracking` and `radius-server vsa send authentication` commands are used so that the dACLs received from RADIUS server can be enforced on the port successfully. Even without dACL being used, command `ip device tracking` is still required for CWA to work, as the switch needs to first learn the IP address of the client so it knows for which source IP address to perform the redirection.

The HTTP and HTTPS servers must be active on the switch so that redirection can take place, but this would leave these services opened for user access as well (so you can actually access the switch for management via HTTP/HTTPS). Using the commands `ip http active-session-modules none` and `ip http secure-active-session-modules none`, the services are disabled while still allowing the redirection.

Configuration

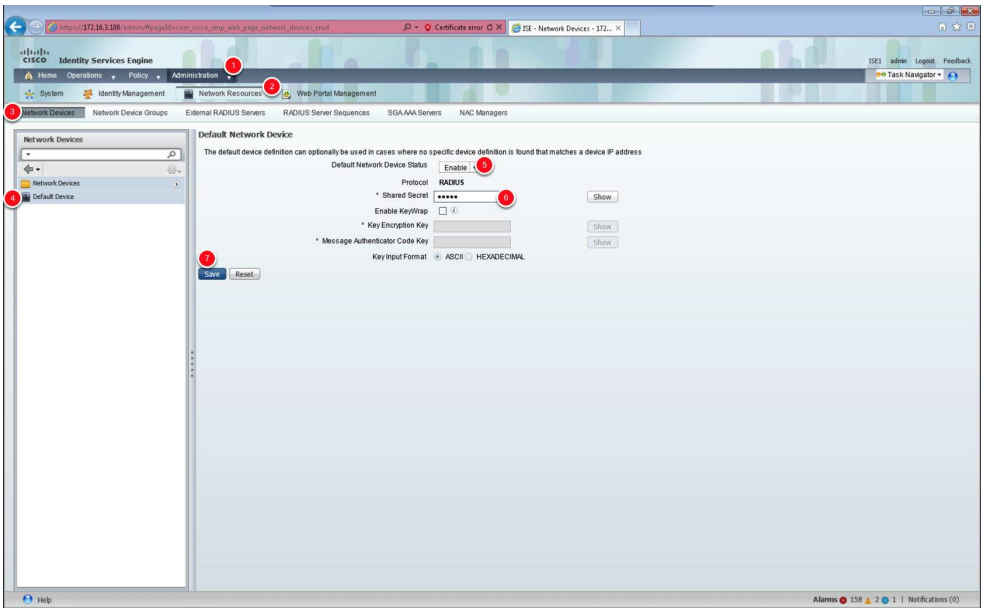
SW1:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
ip device tracking
radius-server vsa send authentication
!
ip access-list extended REDIRECT_ACL
 deny udp any eq bootpc any eq bootps
 deny udp any host 172.16.10.100 eq domain
 deny tcp any host 172.16.3.100 eq 8443
 permit tcp any any eq www
 permit tcp any any eq 443
!
ip http server
ip http secure-server
ip http active-session-modules none
ip http secure-active-session-modules none
!
interface FastEthernet1/0/5
 authentication port-control auto
 authentication host-mode multi-auth
 mab
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
```

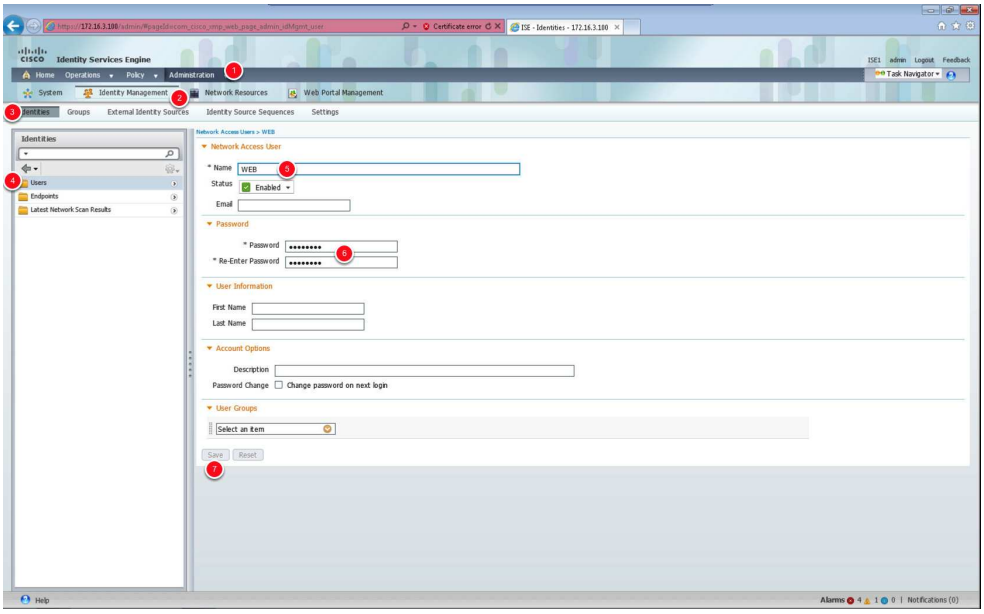
https://t.me/learningnets

```
 ip radius source-interface Loopback0
!
aaa server radius dynamic-author
 client 172.16.3.100 server-key CISCO
```

Configure RADIUS settings for the **Default Device**; alternatively, configure SW1 specifically:

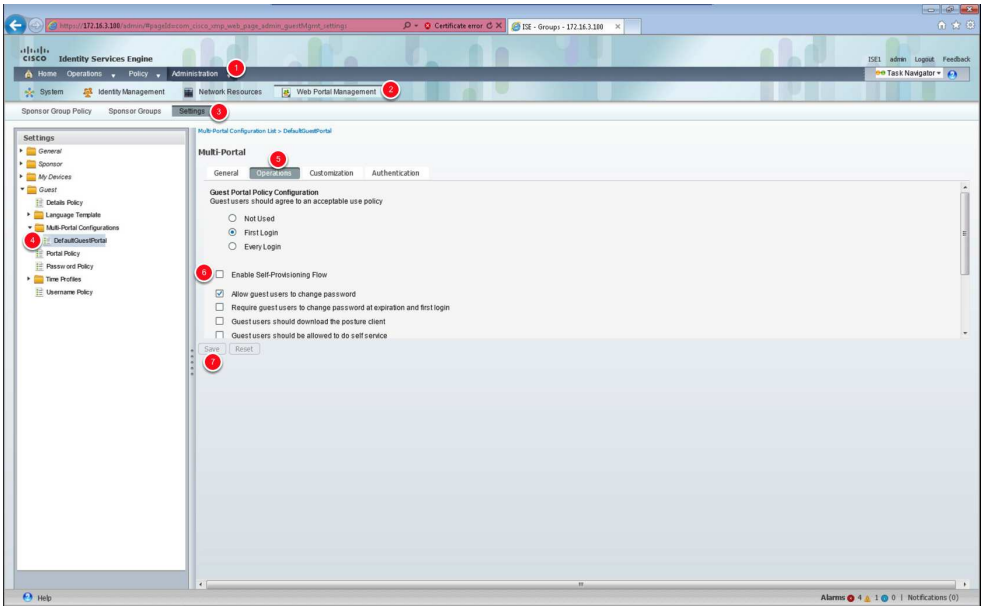


Create a user in the internal Internal Users Identity Store to be used by **TEST-PC-A** when doing the web authentication:

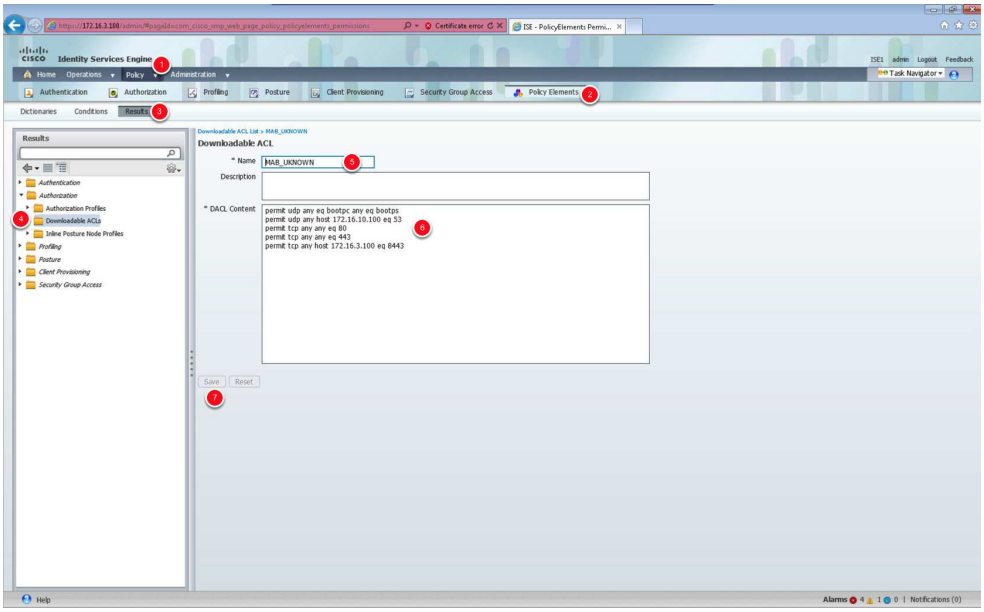


Do not allow users to access the self-provisioning functionality; just allow them to authenticate in the web guest portal:

https://t.me/learningnets

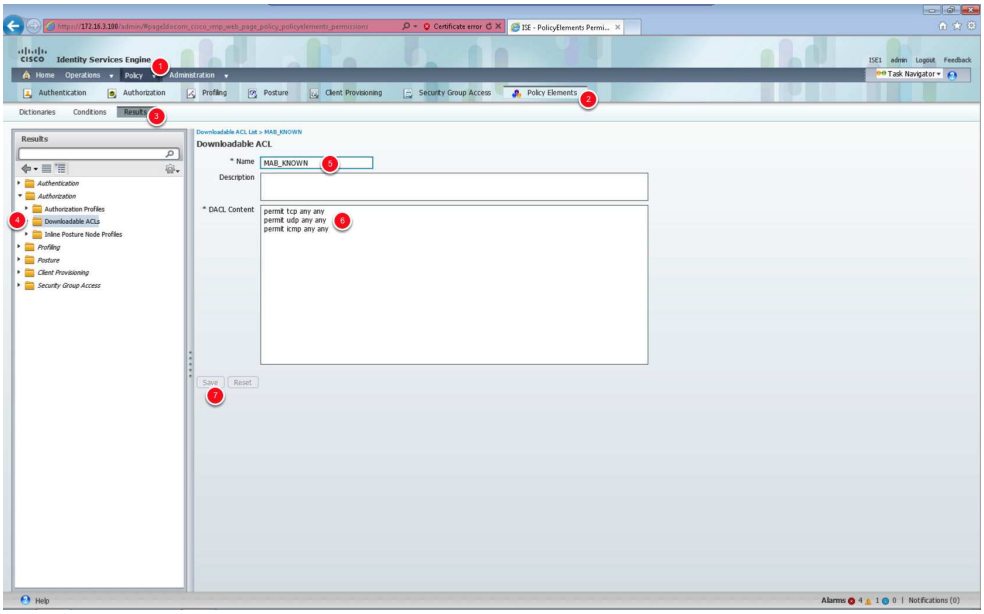


Configure the dACL used to authorize clients before web authentication:

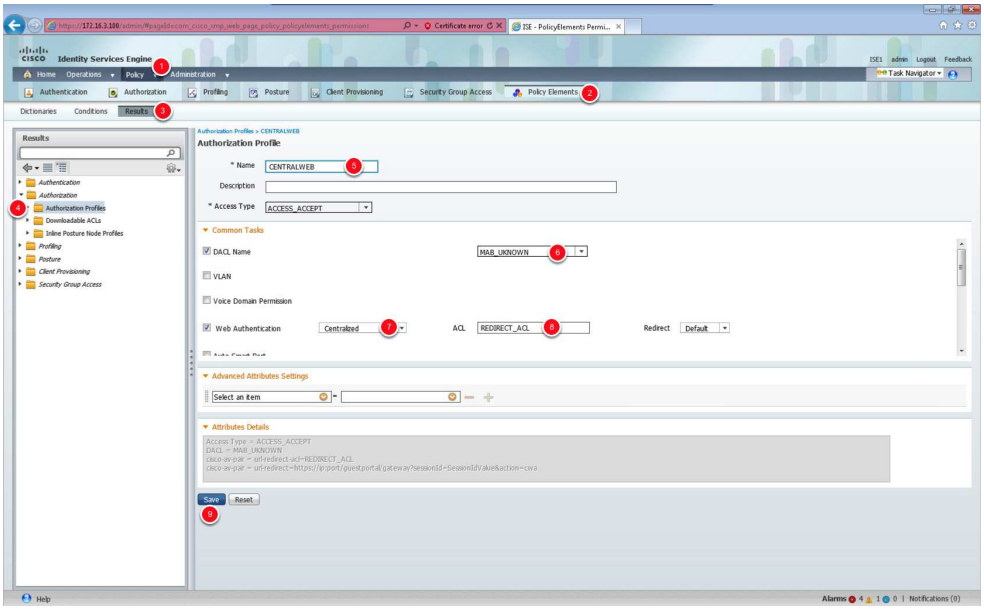


Configure the dACL used to authorize clients after successful web authentication:

https://t.me/learningnets

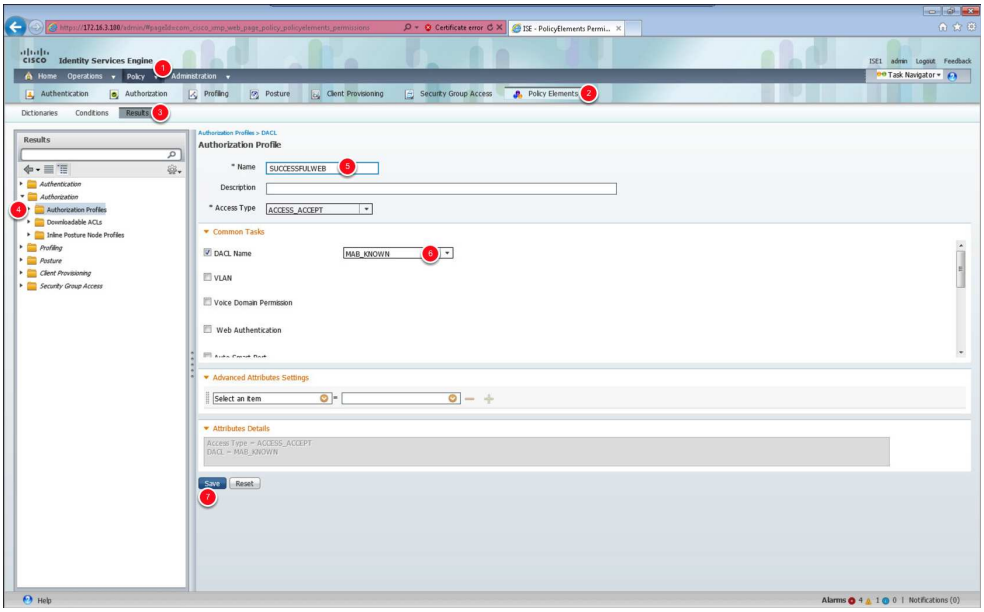


Configure the authorization profile enforcing Central Web Authentication:

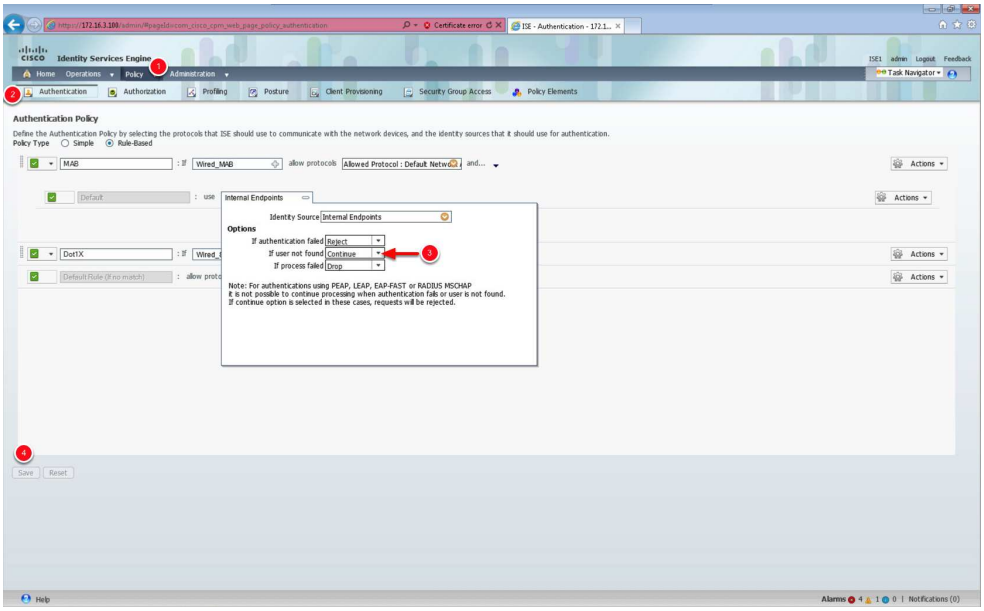


Configure the authorization profile used after successful web authentication:

https://t.me/learningnets

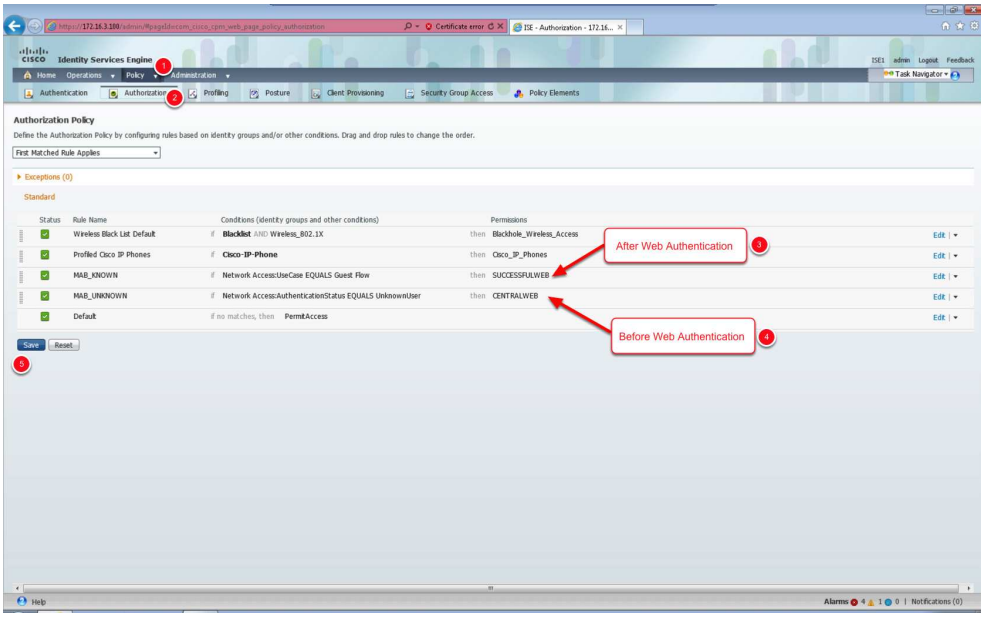


Modify the default MAB authentication policy to use an action of **Continue** if the MAC address is not found in the Internal Endpoint Store:



Configure the two authorization policies, one redirecting unknown users to use Central Web Authentication and one authorizing users who successfully authenticated on the web portal. Note that the authorization policies are matched top-down, just like an ACL, so it's important to place the policy used to redirect users to the web portal at the bottom; otherwise, you will find yourself in a loop, where even after successful authentication in the web portal you will actually match the same policy and be forced to authenticate again.

https://t.me/learningnets



Verification

Verify the initial authorization, where the switch will intercept HTTP/HTTPS traffic and redirect it to ISE at the received redirect-URL:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  136.1.27.51 User-Name:  48-F8-B3-2E-24-23
 Status:  Authz Success
 Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A ACS ACL:  xACSACLx-IP-MAB_UKNOWN-52447e79
 URL Redirect ACL:  REDIRECT_ACL
          U
 RL Redirect:  https://ISE1.inelab.local:8443/guestportal/gateway?sessionId=88011309000000910C1844AE&action=cwa
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  88011309000000910C1844AE
      Acct Session ID:  0x00000099
               Handle:  0x23000091

Runnable methods list:

       Method   State mab       Authc Success

----------------------------------------
```

https://t.me/learningnets

```
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown User-Name:  00-03-6B-3C-35-F0
Status:  Authz Success
Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A ACS ACL:  xACSACLx-IP-MAB_UKNOWN-52447e79
URL Redirect ACL:  REDIRECT_ACL

URL Redirect:  https://ISE1.inelab.local:8443/guestportal/gateway?sessionId=880113090000000920C185DC8&action=cwa
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  880113090000000920C185DC8

      Acct Session ID:  0x0000009A
               Handle:  0x87000092

Runnable methods list:
       Method   State mab        Authc Success
```

Verify the authorization access-list received from ISE:

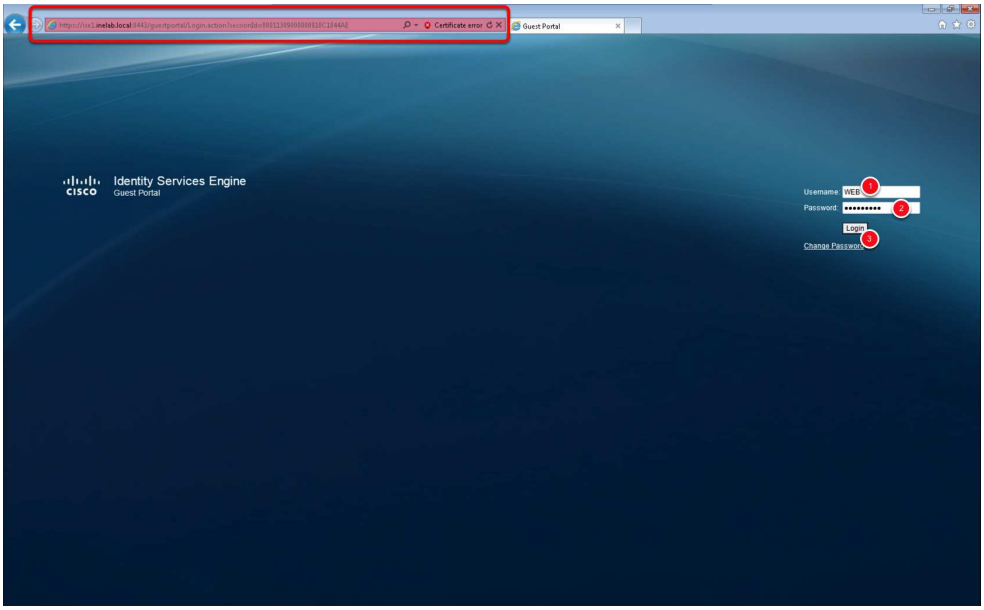
```
SW1#show ip access-lists xACSACLx-IP-MAB_UKNOWN-52447e79
Extended IP access list xACSACLx-IP-MAB_UKNOWN-52447e79 (per-user)
    10 permit udp any eq bootpc any eq bootps
    20 permit udp any host 172.16.10.100 eq domain
    30 permit tcp any any eq www
    40 permit tcp any any eq 443
    50 permit tcp any host 172.16.3.100 eq 8443
!
!SW1#show ip access-lists interface fastEthernet1/0/5

     permit udp host 136.1.27.51 eq bootpc any eq bootps (2 matches)
     permit udp host 136.1.27.51 host 172.16.10.100 eq domain (164 matches)
     permit tcp host 136.1.27.51 any eq www
     permit tcp host 136.1.27.51 any eq 443
     permit tcp host 136.1.27.51 host 172.16.3.100 eq 84431
```

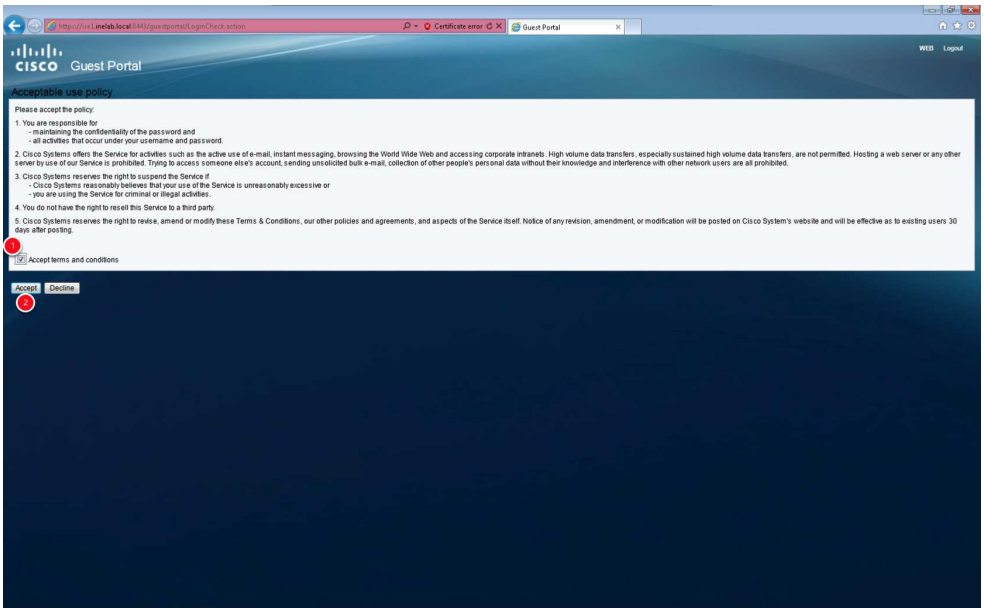
On TEST-PC-A, open Internet Explorer and browse to any HTTP/HTTPS IP address, so that the switch intercepts it and redirects you to the Guest Web Portal

https://t.me/learningnets

from ISE. Note that the URL you've been redirected to is the one received in the authorization policy.



When authenticated with the configured username/password on ISE, you will be presented with the AUP (Acceptable Use Policy). Accept it so that ISE can trigger the CoA process:



Verify the new authorization received dynamically from ISE. Note that it is applicable only for the TEST-PC-A; the IP Phone remained in the captive portal.

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423
           IP Address:  136.1.27.51 User-Name:  WEB
 Status:  Authz Success
```

https://t.me/learningnets

```
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A ACS ACL:  xACSACLx-IP-MAB_KNOWN-52447e90
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  88011309000000910C1844AE
      Acct Session ID:  0x00000099
               Handle:  0x23000091

Runnable methods list:
       Method   State mab        Authc Success



----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown User-Name:  00-03-6B-3C-35-F0
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
           Vlan Group:  N/A
              ACS ACL:  xACSACLx-IP-MAB_UKNOWN-52447e79 URL Redirect ACL:  REDIRECT_ACL

URL Redirect:  https://ISE1.inelab.local:8443/guestportal/gateway?sessionId=88011309000000920C185DC8&action=cwa
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  88011309000000920C185DC8
      Acct Session ID:  0x0000009A
               Handle:  0x87000092


Runnable methods list:
       Method   State mab        Authc Success
```

Verify the authorization access-list received from ISE:

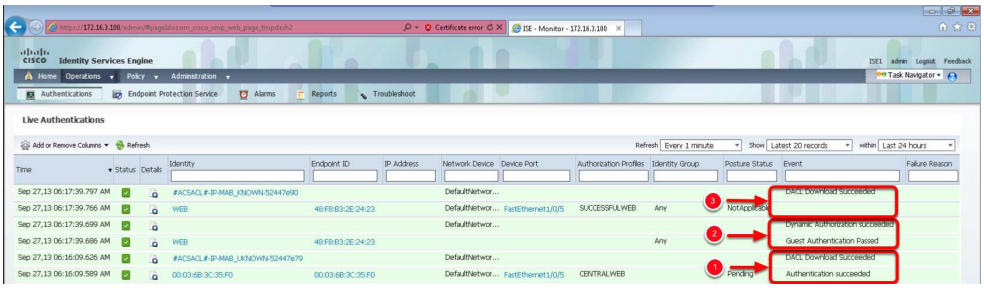
https://t.me/learningnets

```
SW1#show ip access-lists xACSACLx-IP-MAB_KNOWN-52447e90

Extended IP access list xACSACLx-IP-MAB_KNOWN-52447e90 (per-user)
    10 permit tcp any any
    20 permit udp any any
    30 permit icmp any any
!
!SW1#show ip access-lists interface fastEthernet1/0/5

     permit tcp host 136.1.27.51 any
     permit udp host 136.1.27.51 any (52 matches)
     permit icmp host 136.1.27.51 any
```

Verify the Authentication Logs and note the three steps, from the top down: first the Web Authentication succeeded, then the CoA was triggered, and then the new authorization policy was sent to the switch:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

Local Web Authentication Using ACS

Task

- Configure SW1 and ACS to authorize **TEST-PC-A** using Local Web Authentication:
  - **TEST-PC-A** should use Local Web Authentication in case of 802.1x authentication failure.
  - SW1 should try authentication through 802.1x first and webauth second.
  - The authorization policy on ACS should be based on the **Service-Type** RADIUS attribute.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

LWA (local web authentication) makes use of the IOS HTTP cut-through or authentication proxy functionality, which has been discussed in previous labs. Thus, from a configuration perspective it's a combination of 802.1x and authentication proxy commands. From the end host perspective, there is no difference between using LWA and CWA; it still authenticates using a web portal. In most cases, LWA (webauth) will be used as a fallback mechanism if 802.1x and MAB fail, but you can use webauth as the primary or the only method of authentication on a port (so no MAB or 802.1x configuration is required for this last option). When using LWA with RADIUS authentication, the switch will use a value of **Outbound** for the RADIUS **Service-Type** attribute. In general, for the switch to send this attribute in the **Access-Request**

https://t.me/learningnets

RADIUS message, the `radius-server attribute 6 on-for-login-auth` command may be required depending on the IOS code version. As seen also in the IOS authentication proxy lab, RADIUS authorization received from the RADIUS server must use the Cisco AV Pair with a value of **priv-lvl=15** for LWA to work. Because authentication proxy works with both RADIUS and TACACS, this is one case in which 802.1x (with fallback to webauth) works with TACACS.

If guest VLAN (host has no 802.1x supplicant active) and local web authentication are configured for the port, local web authentication takes precedence and guest VLAN is ignored. Authentication-fail VLAN (host has an active 802.1x supplicant but fails authentication) and local web authentication can't be active at the same time on a port, because both are configured as an option of the same `authentication event fail action [authorize|next-method]` command. If critical VLAN authentication (AAA server is dead) and local web authentication are configured for the port, local web authentication will take precedence only if the switch failed over to webauth because it couldn't declare the AAA server as dead in the appropriate time (before the switch failed over to webauth).

Before the switch fails over to webauth and tries to authenticate the port using 802.1x and MAB, traffic is permitted as follows:

- If `authentication open` is configured, all traffic is allowed.
- If a port ACL is configured, traffic is allowed based on the ACE.
- If the switch IOS version supports the default **Auth-Default-ACL** being automatically applied, traffic is allowed based on the ACE; most importantly, DHCP is allowed.
- If no port ACL has been applied and the switch IOS version does not support the default **Auth-Default-ACL** being automatically applied, all traffic is dropped; most importantly, DHCP is not allowed.

When the switch fails over to webauth, traffic is permitted prior to authentication as follows:

- If `authentication open` is configured, all traffic is allowed, but the switch still intercepts HTTP/HTTPS traffic to authenticate the user.
- If a port ACL is configured, traffic is allowed based on the sum of the two ACLs (port ACL and webauth ACL), but the switch still intercepts HTTP/HTTPS traffic to authenticate the user.
- If the switch IOS version supports the default **Auth-Default-ACL** being automatically applied, traffic is allowed based on the sum of the two ACLs (port ACL and webauth ACL), but the switch still intercepts HTTP/HTTPS traffic to authenticate the user.
- If no port ACL has been applied and the switch IOS version does not support the default **Auth-Default-ACL** being automatically applied, traffic is allowed based on

https://t.me/learningnets

the webauth ACL.

Whether DHCP traffic is allowed is of great importance, because with the default timers and default order of authentication where 802.1x is first, the switch will fail 802.1x for an unresponsive host and pass to the next available authentication method after 90 seconds. If the host requires DHCP addressing, the 90-second interval is too big and the host will stop trying to request an IP address via DHCP; it will fail over to assigning itself an IP address from the APIPA range (169.254.0.0/16). This presents a problem for authentication methods such as CWA (central web authentication) and LWA (local web authentication), where the host requires a routable IP address to complete the web authentication process.

So for webauth to function, a PACL may or may not be required, depending on:

- The IOS code version used
- Configured 802.1x timers (using defaults or not)
- IP addressing of the host (static or dynamic through DHCP)
- Whether a dACL is received from the RADIUS server as part of the authorization

The same is true for the webauth fallback profile, where technically the admission rule is mandatory, and the access-list may or may not be required. As long as the authenticating host gets an IP address and can resolve FQDN via DNS query (if the user will browse based on domain names and not IP addresses), the access-list applied in the webauth fallback profile is not really required. In our example it is still used, because most of the times you will use it just to ensure that the webauth process works.

When successfully authenticated, if no ACL is received from RADIUS server, webauth will automatically apply a **permit ip any any** rule for the authenticated host, replacing the source IP address with the IP address of the host identified from the tracking table. The command `ip device tracking` is required to trigger the session state used for local web authentication, while tracking relies on DHCP or ARP packets seen from the host to learn the IP address of the authenticating host.

For more information on why the 802.1x host-mode of **multi-auth** was used, see the ISE Central Web Authentication lab; this is because we don't authenticate/authorize the IP Phone, which will cause a violation to occur and the port to enter an **err-disabled** state.

A dedicated authentication method is configured for the console to keep the default method of no-login for the console, because when the `aaa authentication login default group ACS` command is configured, it applies also for console access.

https://t.me/learningnets

Webath authorization does not support per-user ACL; if an ACL is received from RADIUS as part of the authorization process, it must be a proxy-ACL, Filter-ID ACL, or dACL.

Configuration Option 1

LWA (webauth) is used as the only authentication method for the port (this does not satisfy the task requirements, but is given as an reference example for when only webauth is necessary). No verifications are done for this option because option 2 is more relevant seeing option 2, with session state before and after switch fails over to webauth.

https://t.me/learningnets

SW1:

```
aaa new-model
aaa authentication login default group ACS
aaa authentication login CONSOLE none
aaa authorization auth-proxy default group ACS
!
aaa group server radius ACS
 server-private 172.16.1.100 key CISCO
 ip radius source-interface Loopback0
!
ip device tracking
!
ip access-list extended PREAUTH_WEB
 permit udp any eq bootpc any eq bootps
 permit udp any host 172.16.10.100 eq domain
!
ip http server
ip http secure-server
!
ip admission name WEBAUTH proxy http
fallback profile WEBAUTH_PROFILE
 ip admission WEBAUTH
 ip access-group PREAUTH_WEB in
!
interface FastEthernet1/0/5
 authentication port-control auto
 authentication order webauth
 authentication fallback WEBAUTH_PROFILE
!
line console 0
 login authentication CONSOLE
```

Configuration Option 2

LWA (webauth) is used as a fallback method for 802.1x; no MAB is configured, based on task requirements. Unlike in option 1, 802.1x-specific commands are added.

SW1:

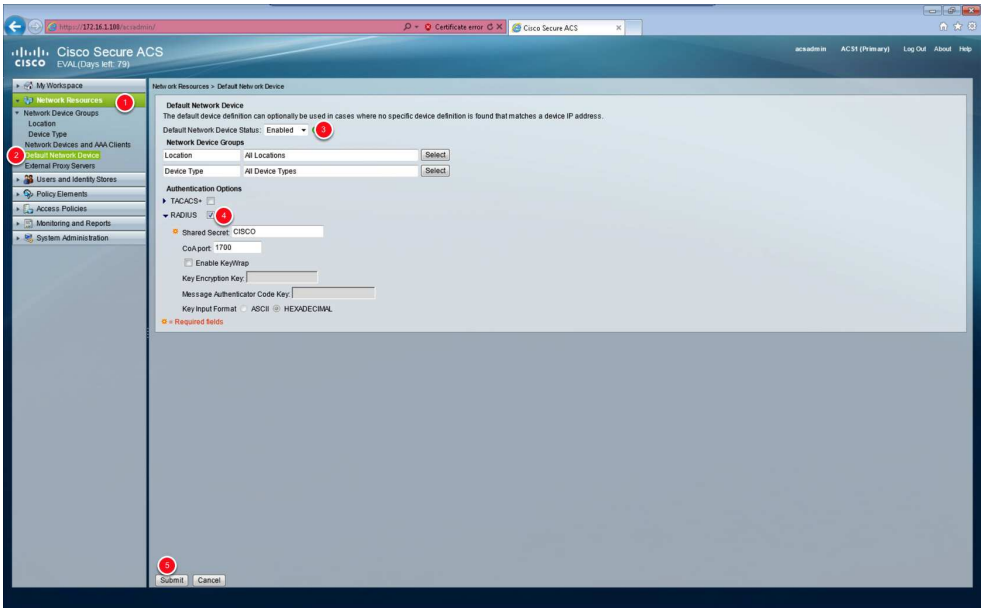
```
aaa new-model
aaa authentication login default group ACS
```

https://t.me/learningnets

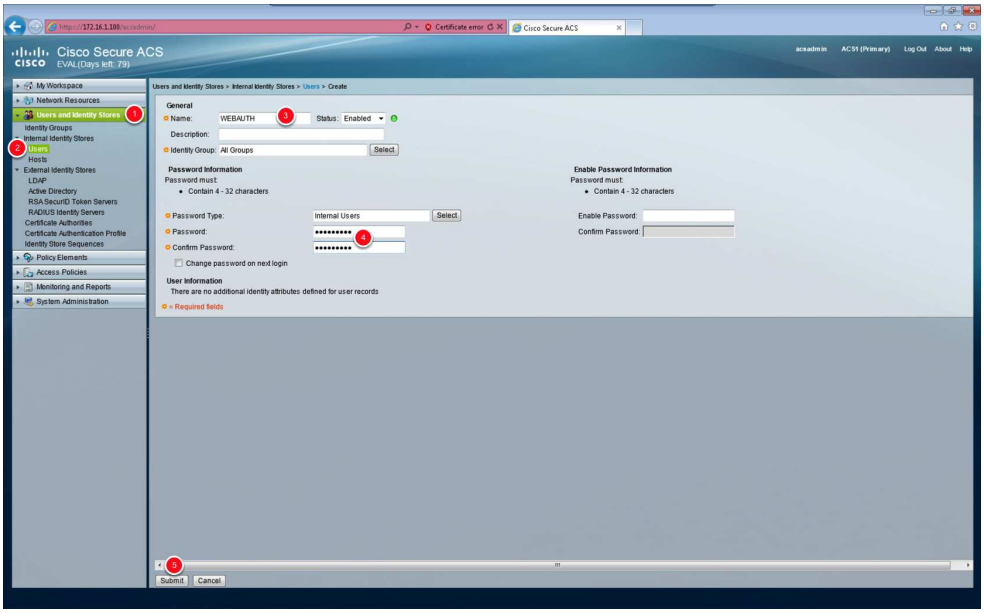
```
aaa authentication login CONSOLE none
aaa authentication dot1x default group ACS
aaa authorization auth-proxy default group ACS
!
aaa group server radius ACS
 server-private 172.16.1.100 key CISCO
 ip radius source-interface Loopback0
!
ip device tracking
dot1x system-auth-control
!
ip access-list extended PREAUTH_WEB
 permit udp any eq bootpc any eq bootps
 permit udp any host 172.16.10.100 eq domain
!
ip http server
ip http secure-server
!
ip admission name WEBAUTH proxy http
fallback profile WEBAUTH_PROFILE
 ip admission WEBAUTH
 ip access-group PREAUTH_WEB in
!
interface FastEthernet1/0/5
 authentication port-control auto
 authentication host-mode multi-auth
 authentication order dot1x webauth
 authentication event fail action next-method
 authentication fallback WEBAUTH_PROFILE
 dot1x pae authenticator
!
line console 0
 login authentication CONSOLE
```

On the ACS server, add SW1 as a NAD using RADIUS only:

https://t.me/learningnets

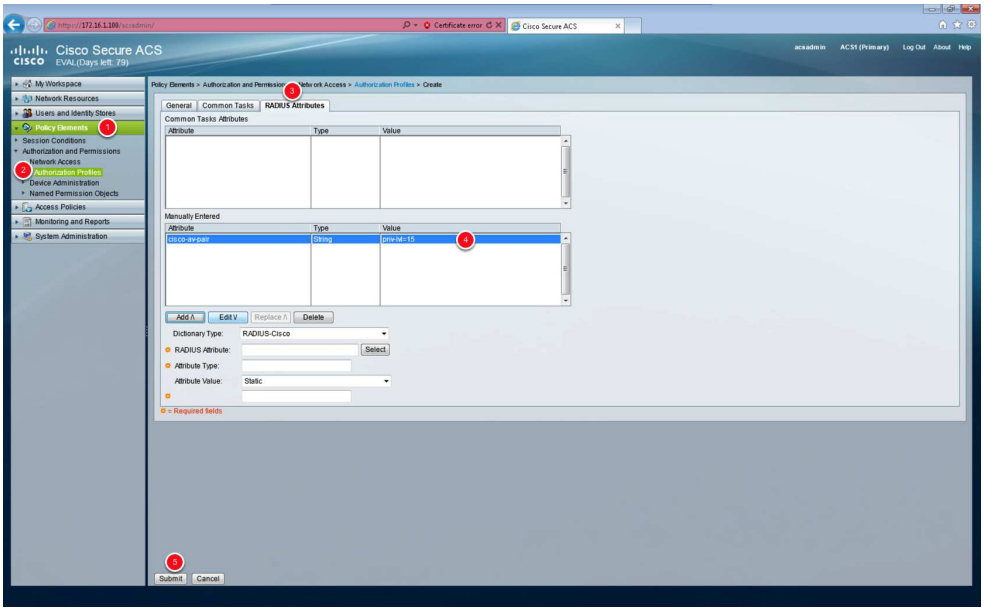


Configure a user in the Internal User store to be used for webauth:

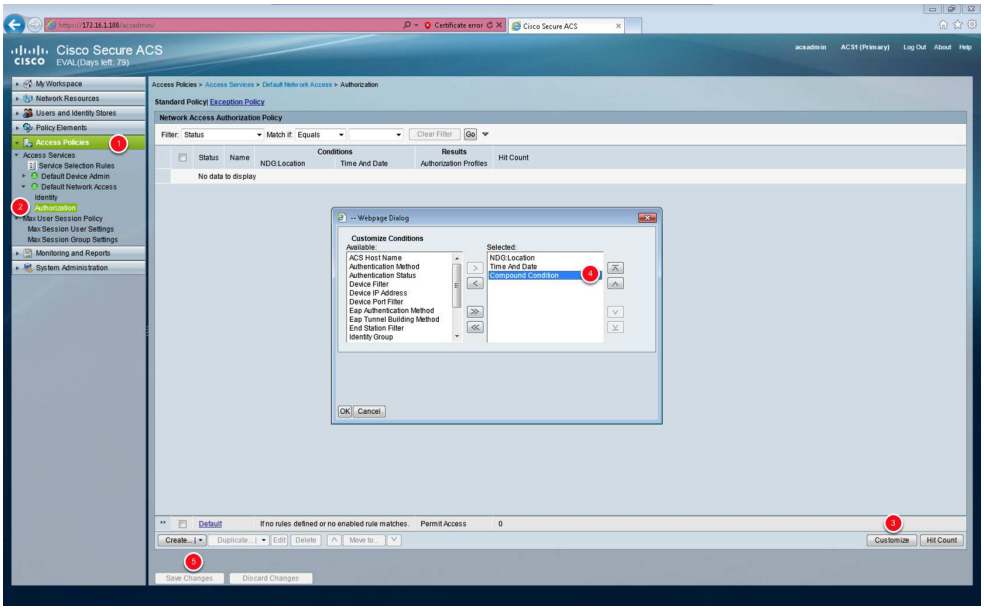


Configure the RADIUS authorization profile, using the Cisco AV Pair with the value of **priv-lvl=15**:

https://t.me/learningnets

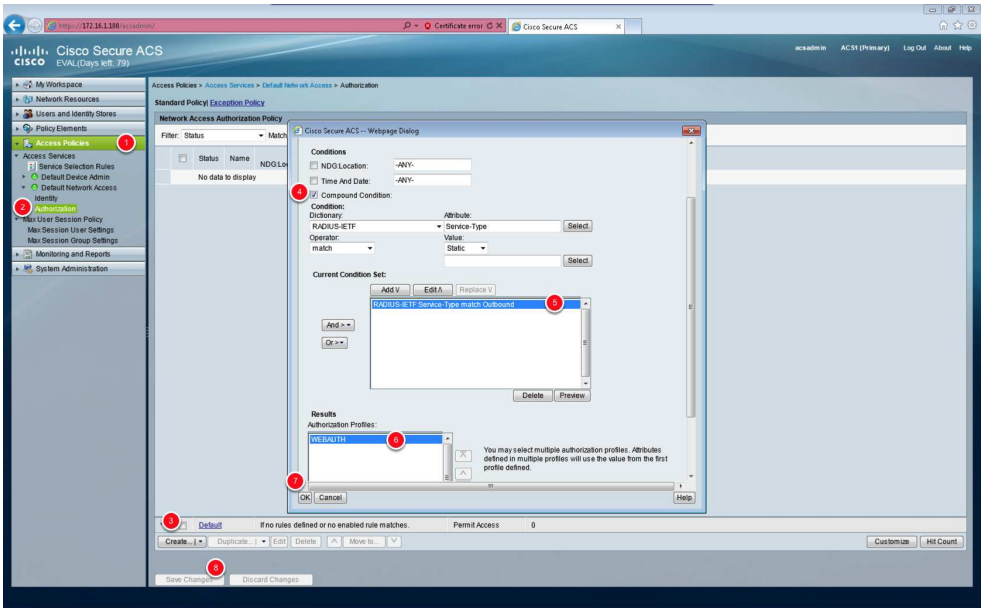


Customize the RADIUS authorization access policies so that a RADIUS **Service-Type** attribute condition can be used:



Configure the RADIUS authorization access policy based on the RADIUS attribute and apply the previously configured profile:

https://t.me/learningnets



Verification Option 2

Note that with the default 802.1x timers, it will take 90 seconds for the switch to fail over to webauth (the port state seems to be authorized by the authentication server, but we see that it actually failed over from 802.1x to webauth, so the state is controlled by webauth now):

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423 IP Address:  136.1.27.52
Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
            Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  0000000000000000030090C0AD
      Acct Session ID:  0x00000005
               Handle:  0x05000003

Runnable methods list:
       Method   State dot1x    Failed over
webauth  Authc Success

----------------------------------------

            Interface:  FastEthernet1/0/5
```

https://t.me/learningnets

```
          MAC Address:  0003.6b3c.35f0 IP Address:  136.1.47.51
Status:  Authz Success
                Domain:  DATA
       Security Policy:  Should Secure
       Security Status:  Unsecure
        Oper host mode:  multi-auth
      Oper control dir:  both Authorized By:  Authentication Server
            Vlan Group:  N/A
       Session timeout:  N/A
          Idle timeout:  N/A
     Common Session ID:  00000000000000040090DA76
       Acct Session ID:  0x00000006
                Handle:  0xCF000004

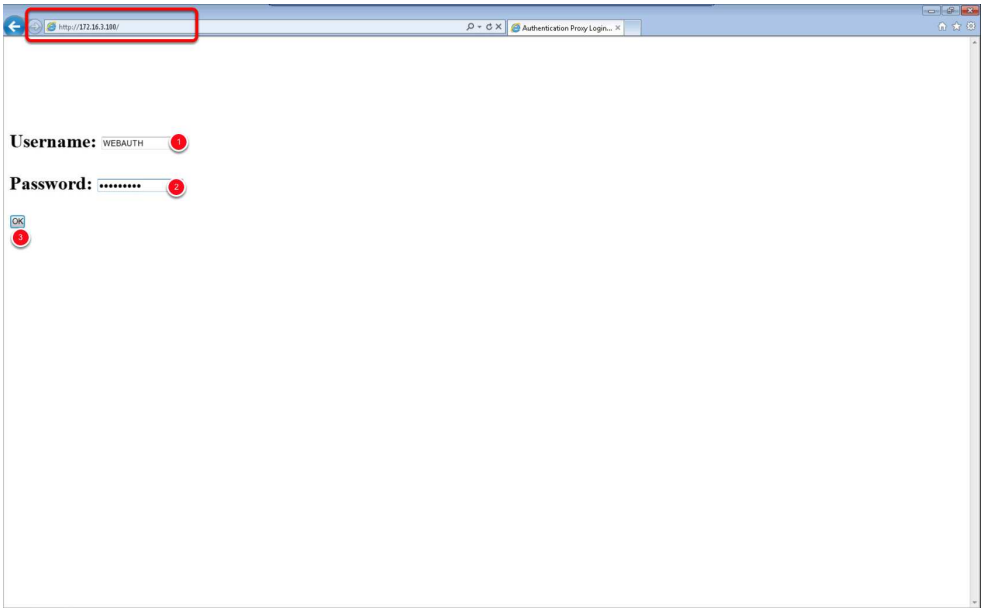
Runnable methods list:
        Method   State dot1x    Failed over
webauth  Authc Success
```

Verify that the ACL controlling access before webauth authentication is the one from the fallback profile. Check that the switch waits for the host to authenticate as the admission session state is **INIT**:

```
SW1#show ip access-lists interface fastEthernet1/0/5
      permit udp host 136.1.27.52 eq bootpc any eq bootps
      permit udp host 136.1.27.52 host 172.16.10.100 eq domain
!
!SW1#show ip admission cache
Authentication Proxy Cache
Total Sessions: 1 Init Sessions: 1 Client IP 136.1.27.52 Port 0, timeout 60, state INIT
```

Using Internet Explorer from **TEST-PC-A**, try browsing to an IP address. The switch will present you with its web portal to authenticate; enter the username/password configured on ACS:

https://t.me/learningnets



When authenticated, verify that the admission session state changed to established, and confirm that ACL controlling access has changed to allow full access only for the authenticated IP address:

```
SW1#show ip access-lists interface fastEthernet1/0/5
 permit ip host 136.1.27.52 any (1 match)
!
!SW1#show ip admission cache
Authentication Proxy Cache
Total Sessions: 1 Init Sessions: 0 Client IP 136.1.27.52 Port 11314, timeout 60, state ESTAB
```

When authenticated, note that there are actually no changes for the port state (TEST-PC-A is authenticated and the IP Phone is not); the change was visible in the admission session state:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423 IP Address:  136.1.27.52
 Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000030090C0AD
      Acct Session ID:  0x00000005
```

https://t.me/learningnets

```
               Handle:  0x05000003

Runnable methods list:
       Method   State dot1x     Failed over
webauth  Authc Success

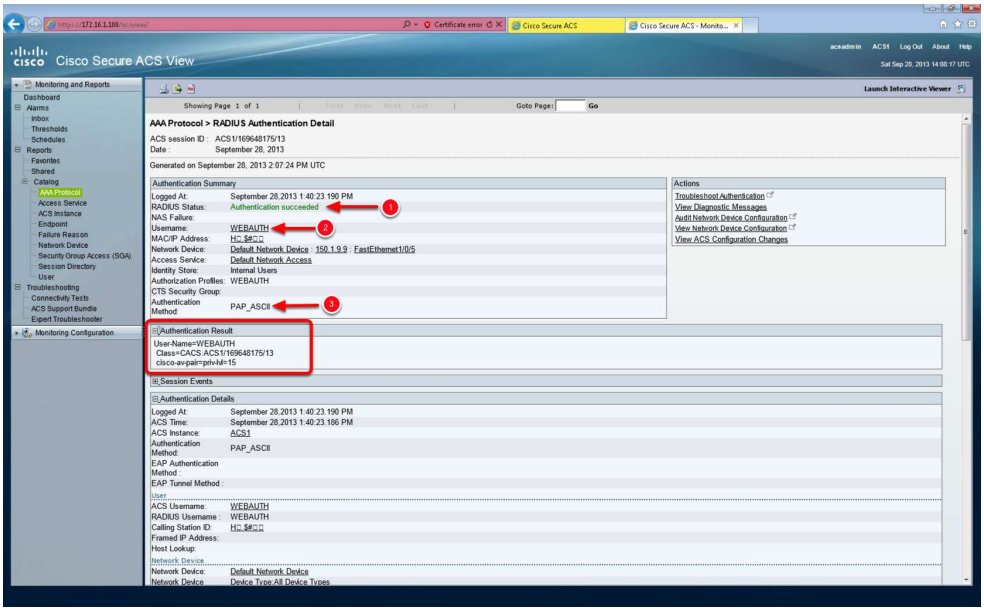
---------------------------------------

            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0 IP Address:  136.1.47.51
Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000040090DA76
      Acct Session ID:  0x00000006
               Handle:  0xCF000004

Runnable methods list:
       Method   State dot1x     Failed over
webauth  Authc Success
```

Verify the RADIUS Authentication Logs from ACS, and note that the authentication method is PAP:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

Local Web Authentication Using ISE

Task

- Configure SW1 and ISE to authorize **TEST-PC-A** using Local Web Authentication:
  - **TEST-PC-A** should use Local Web Authentication in case of MAB failure.
  - SW1 should try authentication through MAB first and webauth second.
  - When successfully authenticated, all TCP and ICMP traffic should be allowed, and authorization received from ISE using Cisco AV Pair.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

There are three possible methods to receive an ACL from the RADIUS server when LWA is used: dACL, Filter-ID ACL, and Proxy-ACL. The first two options were discussed in the MAB Using ISE lab. Proxy-ACL has the following requirements:

- Requires ACL configured on the RADIUS server with the format **proxyacl#n=acl-rule** (for example, **proxyacl#1=permit tcp any any**) to allow all TCP traffic.
- Verified with the `show authentication session interface [interface-name]` and `show epm session ip [IP Address]` commands.
- Does **not** require the `aaa authorization network default group` command to be configured on the switch.

https://t.me/learningnets

- Requires `aaa authorization auth-proxy default group` because actually auth-proxy is used for authorization and not 802.1x.
- Requires `ip device tracking` for LWA to trigger the session state.
- Uses Cisco AV Pair.

Given the supported ACLs for webauth and task requirements, a proxy-ACL must be used. Note that if compared with the ACS task of LWA, when webauth is used as failover from MAB, the command `authentication event fail action next-method` is not required; this is necessary only for 802.1x failover to webauth. Also, because 802.1x is not configured, the global command `dot1x system-auth-control` is not required.

Configuration

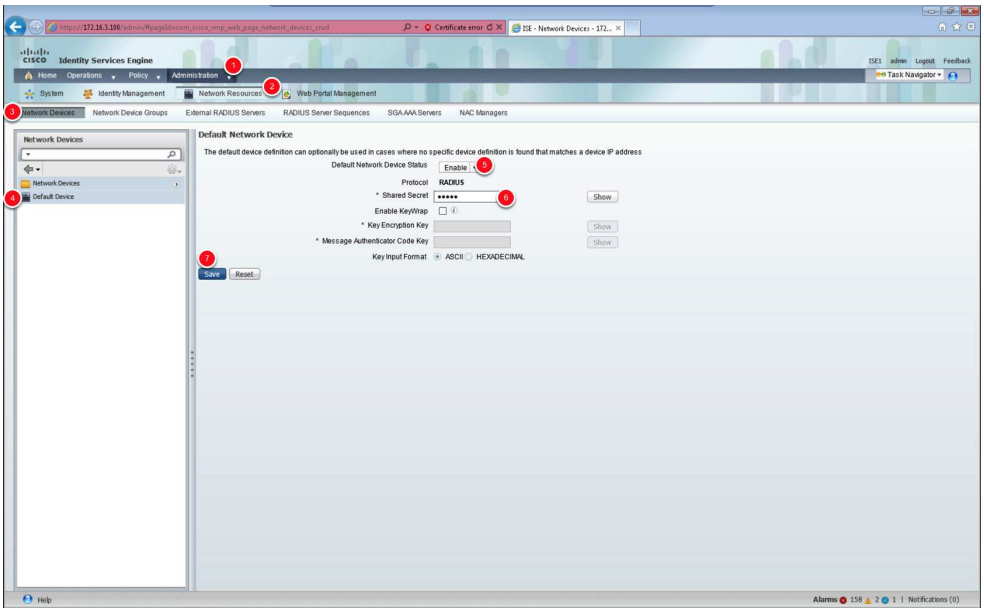
SW1:

```
aaa new-model
aaa authentication login default group ISE
aaa authentication login CONSOLE none
aaa authentication dot1x default group ISE
aaa authorization auth-proxy default group ISE
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Loopback0
!
ip device tracking
!
ip access-list extended PREAUTH_WEB
 permit udp any eq bootpc any eq bootps
 permit udp any host 172.16.10.100 eq domain
!
ip http server
ip http secure-server
!
ip admission name WEBAUTH proxy http
fallback profile WEBAUTH_PROFILE
 ip admission WEBAUTH
 ip access-group PREAUTH_WEB in
!
interface FastEthernet1/0/5
 authentication port-control auto
 authentication host-mode multi-auth
 authentication order mab webauth
 authentication fallback WEBAUTH_PROFILE
```

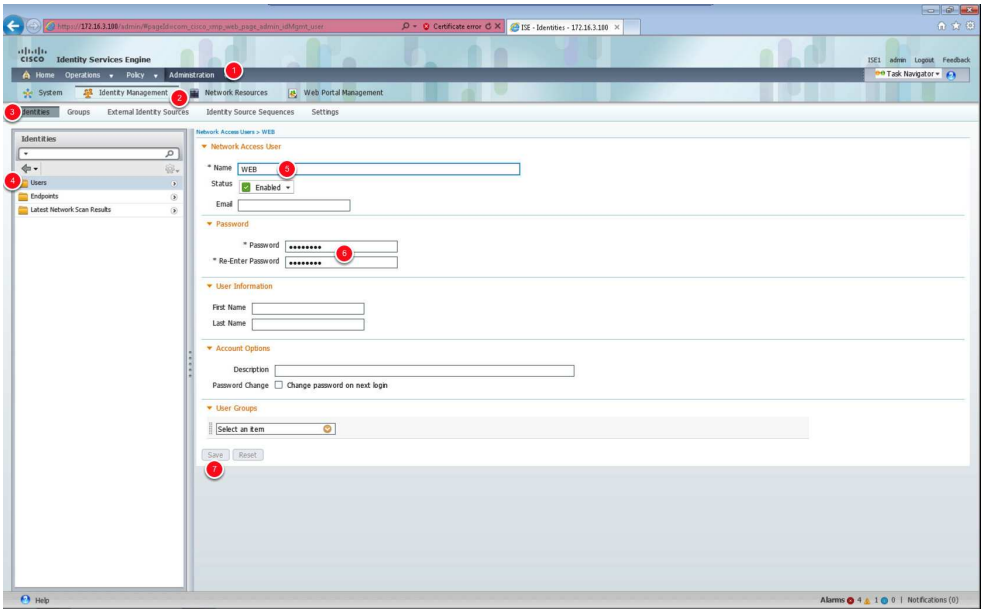
https://t.me/learningnets

```
 mab
!
line console 0
 login authentication CONSOLE
```

Configure RADIUS settings for the **Default Device**; alternatively, configure SW1 specifically:

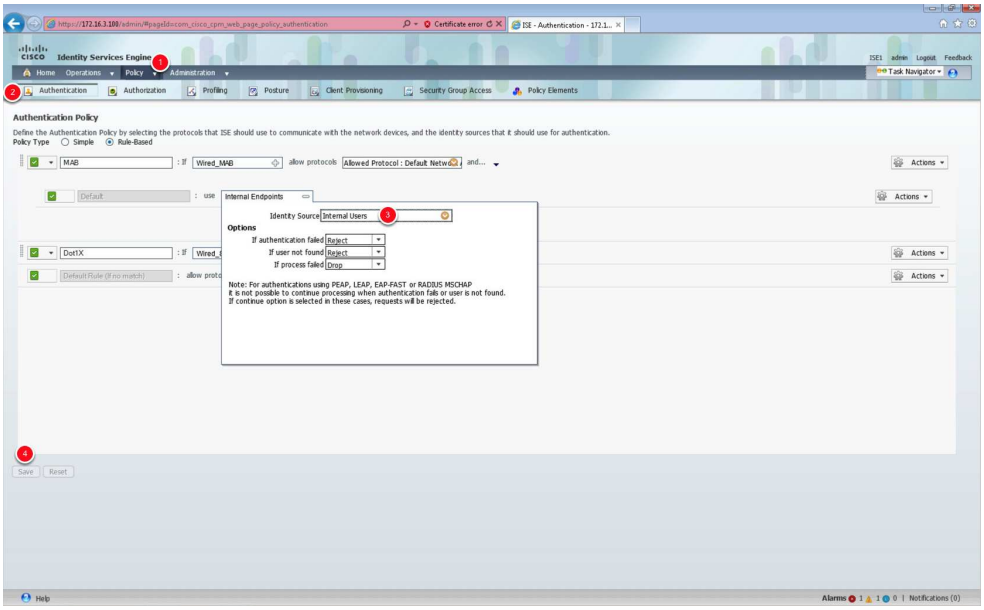


Create a user in the internal Internal Users Identity Store to be used by **TEST-PC-A** when doing the web authentication:

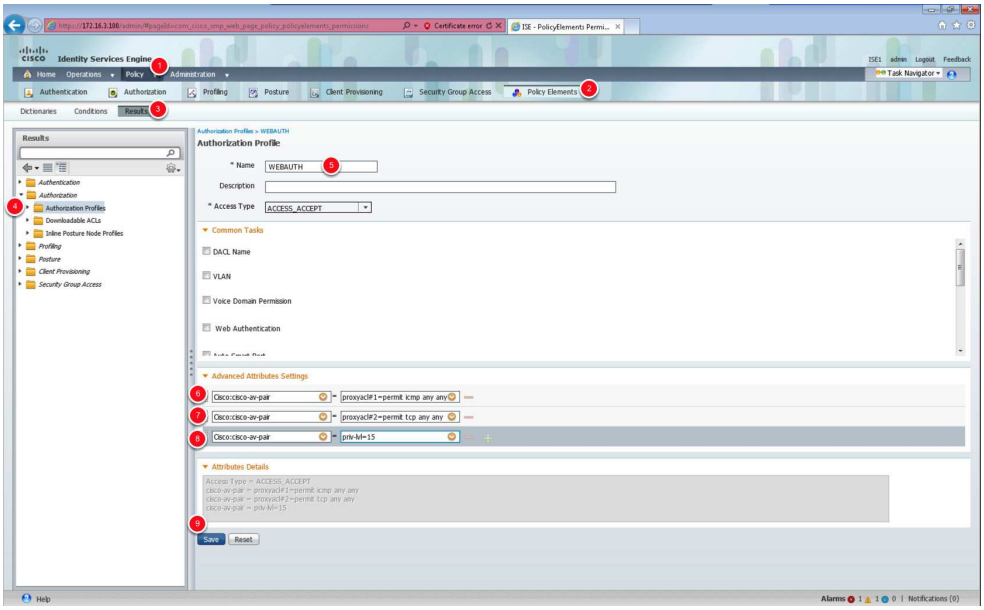


Modify the default MAB authentication policy to check the Internal User Store:

https://t.me/learningnets



Configure the RADIUS authorization profile with the required Cisco AV Pair values:



Apply the authorization profile to the default authorization policy, or create a new authorization policy for this scope.

Verification

Verify that the switch failed over from MAB to webauth (note that the port state seems to be authorized by the authentication server, but we see that it actually failed over from MAB to webauth, so state is controlled by webauth now):

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423 IP Address:  136.1.27.52
```

https://t.me/learningnets

```
Status:  Authz Success

               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000000121D9F
      Acct Session ID:  0x00000001
               Handle:  0xD1000000

Runnable methods list:
       Method   State mab        Not run
webauth  Authc Success


---------------------------------------

            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown
            User-Name:  00036b3c35f0
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000000100122A77
      Acct Session ID:  0x00000002
               Handle:  0x95000001

Runnable methods list:
       Method   State
       mab      Failed over
       webauth  Authc Success
```

Verify that the ACL controlling access before webauth authentication is the one from

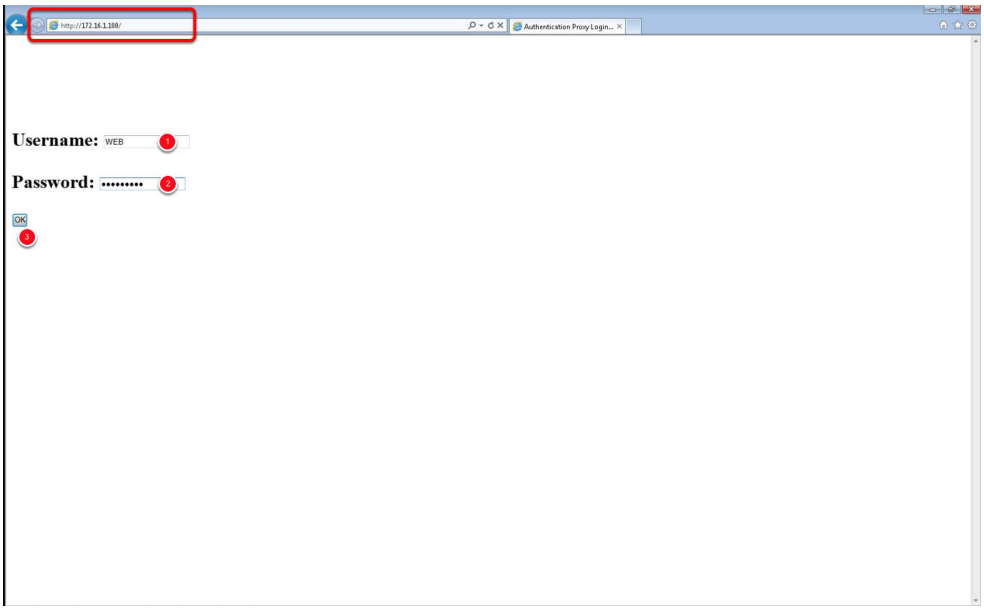
https://t.me/learningnets

fallback profile; check that the switch waits for host to authenticate as the admission session state is **INIT**:

```
SW1#show ip access-lists interface fastEthernet1/0/5
      permit udp host 136.1.27.52 eq bootpc any eq bootps
      permit udp host 136.1.27.52 host 172.16.10.100 eq domain
!
!SW1#show ip admission cache

Authentication Proxy Cache
Total Sessions: 1 Init Sessions: 1 Client IP 136.1.27.52 Port 11932, timeout 60, state INIT
```

Using Internet Explorer from **TEST-PC-A**, try browsing to an IP address. The switch will present you with its web portal to authenticate; enter the username/password configured on ISE:



When authenticated, verify that the admission session state changed to established, and confirm that the ACL controlling access is the one downloaded from ISE:

https://t.me/learningnets

```
SW1#show ip admission cache

Authentication Proxy Cache
Total Sessions: 1 Init Sessions: 0 Client IP 136.1.27.52 Port 12220, timeout 60, state ESTAB

!
!SW1#show epm session ip 136.1.27.52


    Admission feature:  AUTHPROXY Proxy ACL:  permit icmp any any
Proxy ACL:  permit tcp any any
```

When authenticated, note that there are actually no changes for the port state except that proxy-ACL is visible (TEST-PC-A is authenticated and the IP Phone is not); the change was visible in the admission session state:

```
SW1#show authentication sessions interface fastethernet1/0/5
            Interface:  FastEthernet1/0/5
          MAC Address:  48f8.b32e.2423 IP Address:  136.1.27.52
            User-Name:  48f8b32e2423 Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A Proxy ACL:  permit icmp any any
Proxy ACL:  permit tcp any any
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000000121D9F
      Acct Session ID:  0x00000001
               Handle:  0xD1000000

Runnable methods list:
       Method   State mab        Failed over
          webauth  Authc Success


----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown
            User-Name:  00036b3c35f0
```

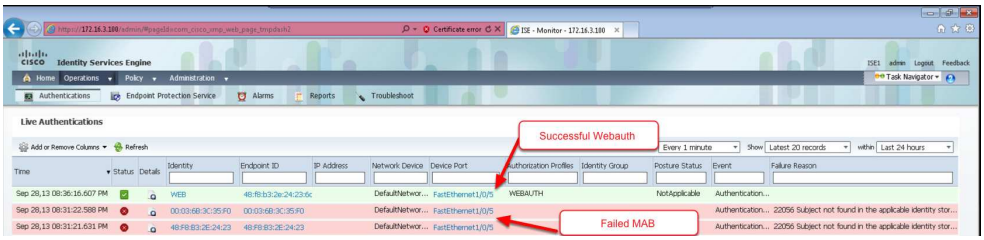
https://t.me/learningnets

```
               Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
        Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  00000000000000000100122A77
      Acct Session ID:  0x00000002
               Handle:  0x95000001

Runnable methods list:
       Method   State

       mab      Failed over
       webauth  Authc Success
```

Verify the Authentication Logs from ISE, and note the two failed MAB authentication requests, one for IP Phone and one for TEST-PC-A:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

ISE CPP and Posture Assessment

Task

- Configure ISE for client provisioning and posture compliance check of **TEST-PC-A** as follows:
  - Provision the NAC Agent only if device was authenticated through MAB.
  - Check if the **firefox.exe** process is running, otherwise start it for device to be compliant.
  - If device is not compliant, allow only DNS and ICMP Echo traffic; if compliant allow all traffic.
  - Posture policy should apply only for 32-bit Windows operating systems.
  - After user authentication, close the agent GUI automatically in 10 seconds.
- Assume that there could be multiple devices behind port Fa1/0/5, all requiring to be authenticated.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.
>
> If NAC Agent is not available in ISE, upload it from a folder named **Lab Files** on Desktop of any **TEST-PC**, browse further to a folder named **Cisco ISE**.

Overview

Posture assessment allows for verification of authenticated device compliance with

https://t.me/learningnets

certain imposed requirements and appropriate authorization. There are three major steps in this process (not including the authentication policy):

1. CPP or Client Provisioning Policy
2. Posture Policy
3. Authorization Policy based on posture status

Client Provisioning Policy

The posture assessment of the endpoint is done through the NAC Agent client, and this can be installed manually by the administrator or automatically through software deployment services. An alternative to this is for ISE to push the client installation after the endpoint is successfully authenticated, a process called CPP. For this to work, ISE will push an authorization policy to the switch with the following attributes:

- A dACL is applied, allowing only access to required services for NAC Agent to be installed.
- A redirect URL is pushed from ISE pointing to an internal webpage used for multiple scopes, such as guest portal, native supplicant configuration, and client provisioning.
- A redirect ACL name is pushed from ISE, and all traffic matching the permit entries from the ACL will be redirected to the ISE portal for CPP.

There are two versions for NAC Agent; it can be persistent, whereby the agent is installed and is automatically loaded each time a user logs in (available for Windows and Mac OS X), or it can be temporal, whereby a web-based agent is dynamically downloaded to the endpoint for each new session and then removed after the posture assessment process (available for Windows only). The main difference between the two options is that the web-based agent has limited remediation capabilities and does not support single sign-on. NAC Agent also facilitates the remediation process and provides an optional acceptable use policy (AUP) to the end user.

Posture Policy

The posture policy defines a set of requirements that an endpoint must meet to be compliant, such as certain file presence, registry key, started processes, applications, Windows versions and updates, and anti-virus/anti-spyware checks. Posture policies are applied against endpoints based on conditions such as user identity and client OS type. The posture status of an endpoint can be:

- Unknown, which means that no data was collected to determine posture state; most likely the NAC agent is not installed on the endpoint.

https://t.me/learningnets

- Noncompliant, which that means a posture assessment was performed, but at least one of the requirements failed the check.
- Compliant, which means that a posture assessment was performed and the endpoint is compliant with all specified requirements.

There can be multiple conditions/requirements specified for compliant status to be achieved; each requirement can be associated with a remediation action that helps endpoints satisfy the requirement, such as Windows system updates.

Authorization Policy

This is the regular authorization policy found in all previous ISE tasks, granting access to certain resources based on VLAN and ACL assignment, for example. However, when posture assessment is involved, different authorization policies will be configured for each of the three possible posture statuses; the most important is the policy for the **unknown** status, which can be used to install NAC Agent on the end system. When the agent is installed but the system is not compliant, you would want to apply an ACL that allows access to the self-remediation process, such as access to antivirus/antispyware and Windows update services. The ACL used for the **unknown** authorization must, at a minimum, permit the following traffic:

- DHCP, if the host requires dynamic IP addressing.
- DNS, to resolve ISE hostname/FQDN from the redirection URL.
- TCP port 8443, for access to ISE portal; required if client provisioning is used (the default port 8443 can be modified).
- TCP and UDP port 8905, used for NAC Agent to discover the ISE node and for keepalives (the NAC Agent can be manually configured with the ISE IP address or FQDN).
- TCP and UDP port 8909, used for NAC Agent wizard installation.

Configuration

SW1:

```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
aaa group server radius ISE
 server-private 172.16.3.100 key CISCO
 ip radius source-interface Loopback0
!
```

https://t.me/learningnets

```
radius-server vsa send authentication
!
aaa server radius dynamic-author
 client 172.16.3.100 server-key CISCO
!
ip http server
ip http secure-server
!
ip device tracking
!
ip access-list extended ACL_REDIRECT
 deny udp any eq bootpc any eq bootps
 deny udp any host 172.16.10.100 eq 53
 deny tcp any host 172.16.3.100 eq 8443
 deny udp any host 172.16.3.100 eq 8905
 deny tcp any host 172.16.3.100 eq 8905
 deny udp any host 172.16.3.100 eq 8909
 deny tcp any host 172.16.3.100 eq 8909
 permit ip any any
!
interface FastEthernet1/0/5
 authentication port-control auto
 authentication host-mode multi-auth
 mab
```

Identify the MAC address of **TEST-PC-A** to be added to the ISE Internal Hosts Store; note that this MAC address is different depending on the rack being assigned:

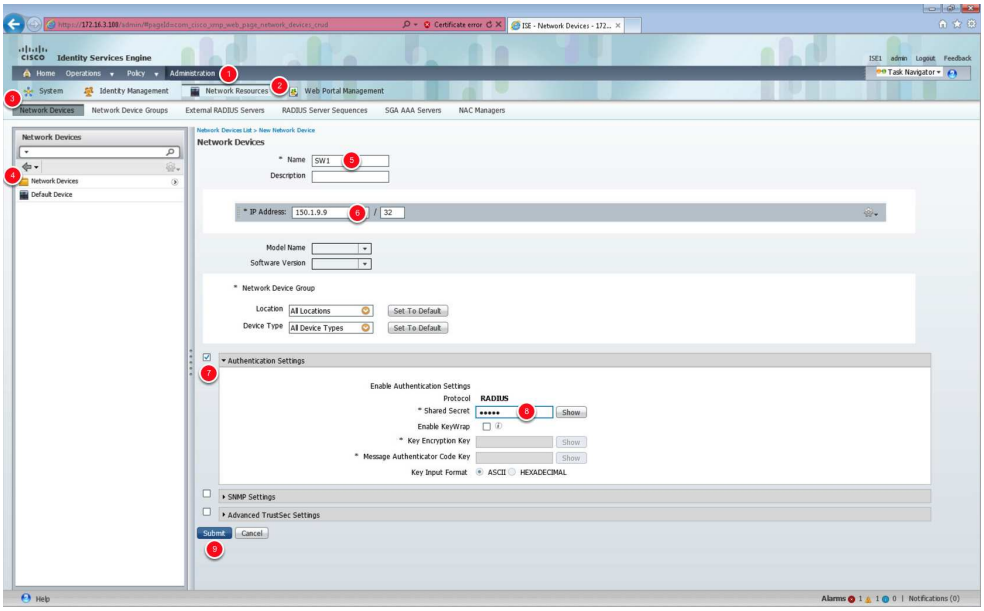
```
SW1#show mac address-table dynamic interface fastEthernet1/0/5
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----  27    48f8.b32e.2423    DYNAMIC     Fa1/0/5

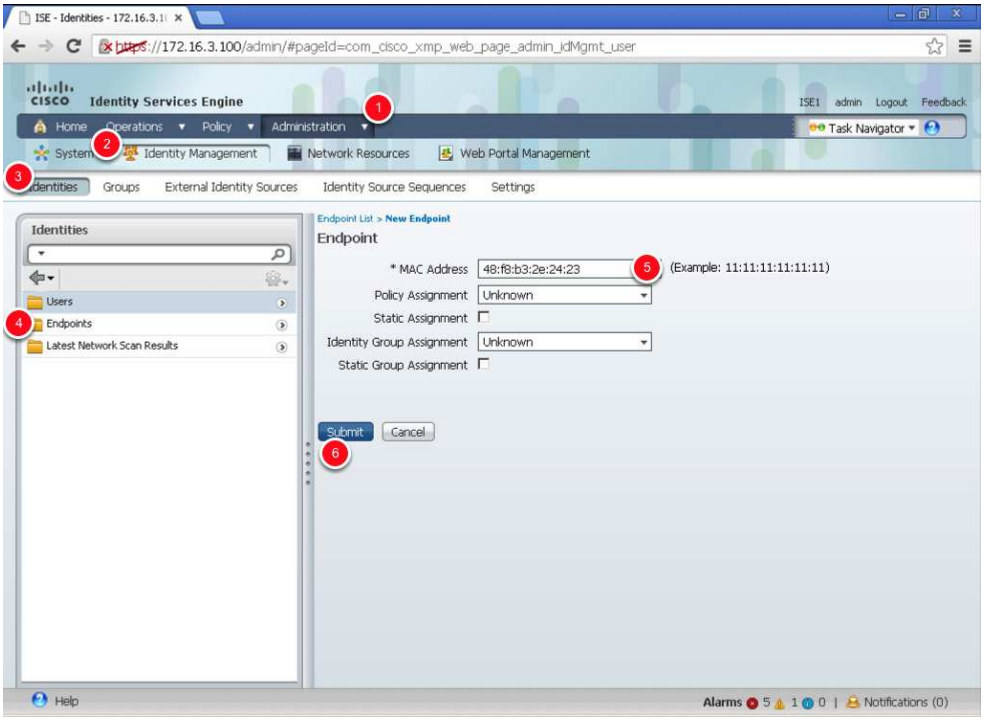

  47    0003.6b3c.35f0    DYNAMIC     Fa1/0/5
```

On ISE, first add SW1 as NAD device:

https://t.me/learningnets

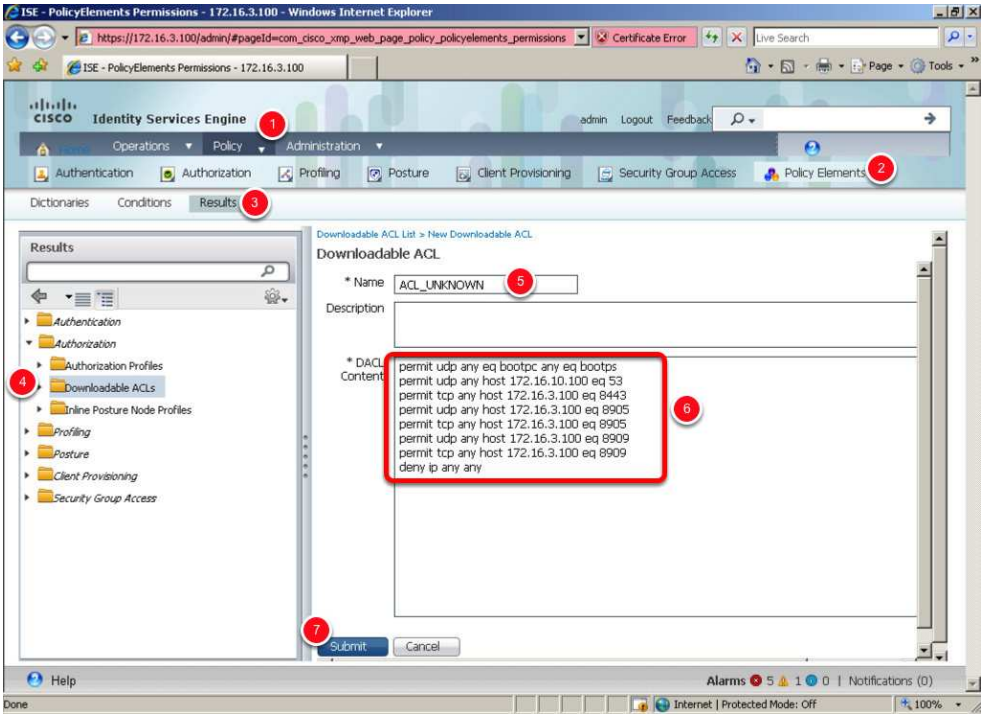


Add the MAC address of **TEST-PC-A** in the internal Endpoint store, to be used for MAB:

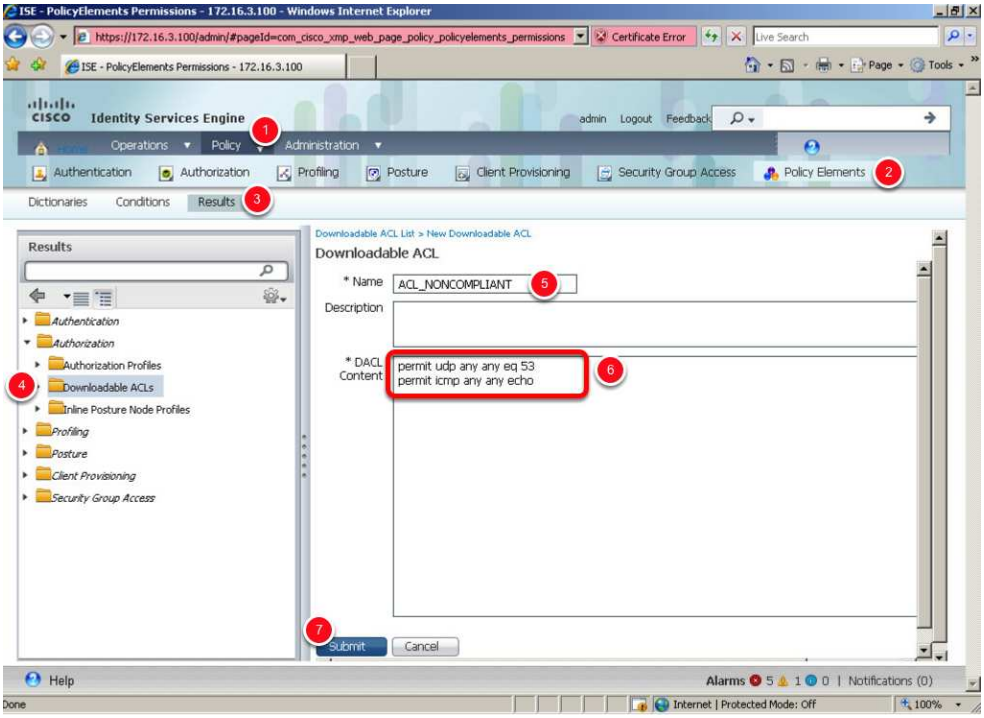


The built-in authentication policies of ISE can be used for MAB, so we will not configure anything for it; however, we need to configure authorization policies. First, configure the three access-lists required for all possible authorization policies. Create the dACL for unknown posture status:

https://t.me/learningnets

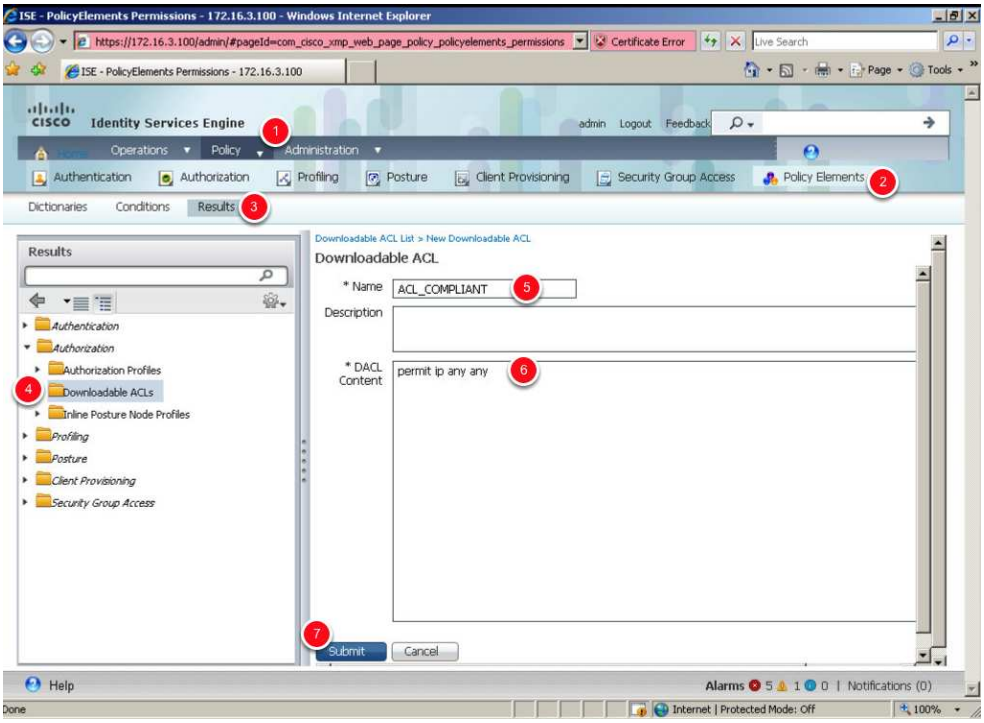


Configure the dACL for noncompliant posture status:

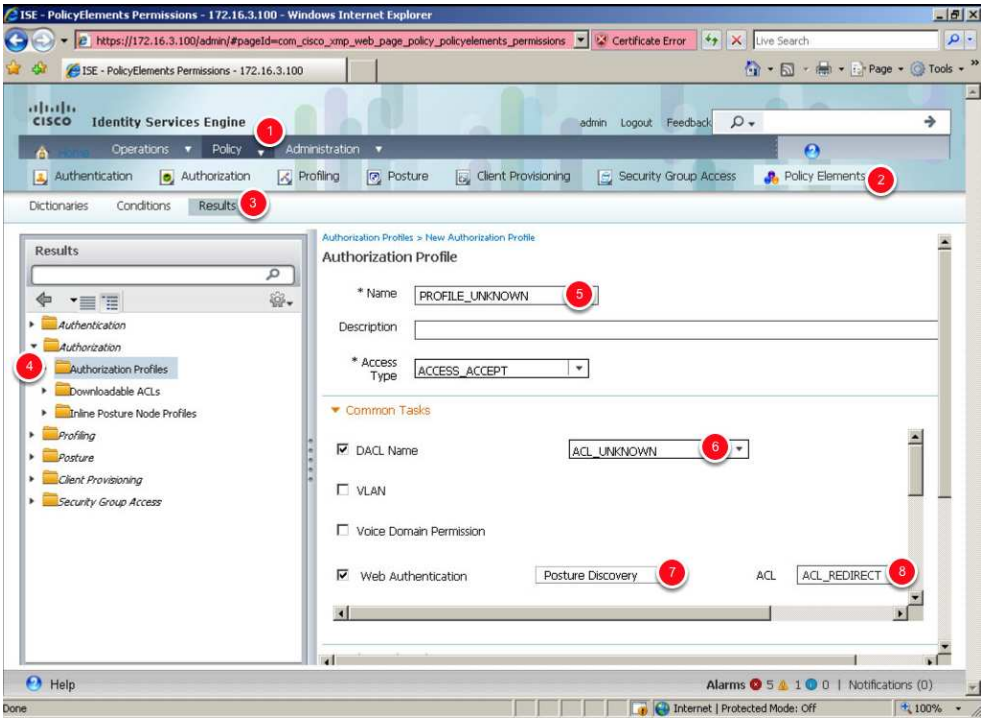


Configure the dACL for compliant posture status:

https://t.me/learningnets

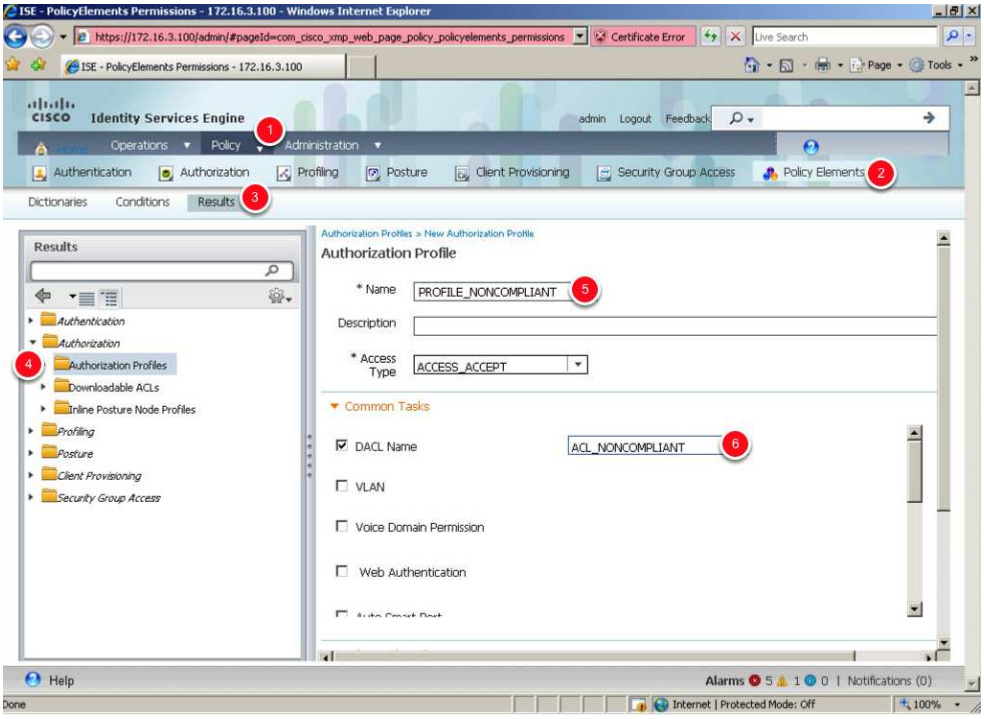


Configure the authorization profile for unknown posture status, specify the dACL name of the redirection ACL as configured on SW1, and configure centralized web authentication for posture discovery and client provisioning:

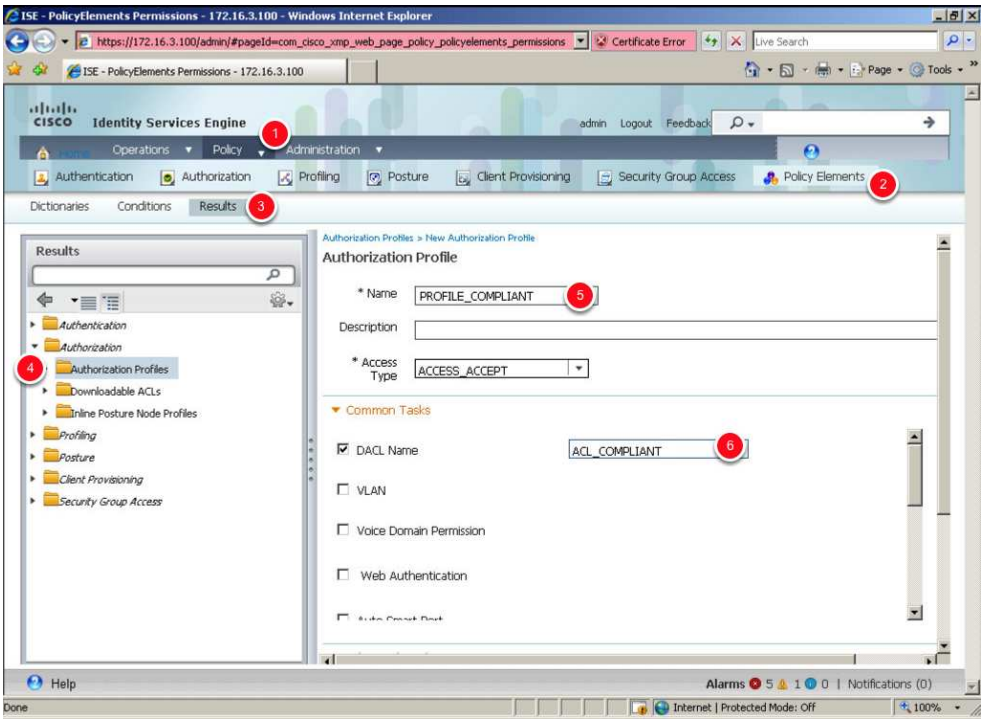


Configure the authorization profile for noncompliant posture status, and specify the dACL:

https://t.me/learningnets

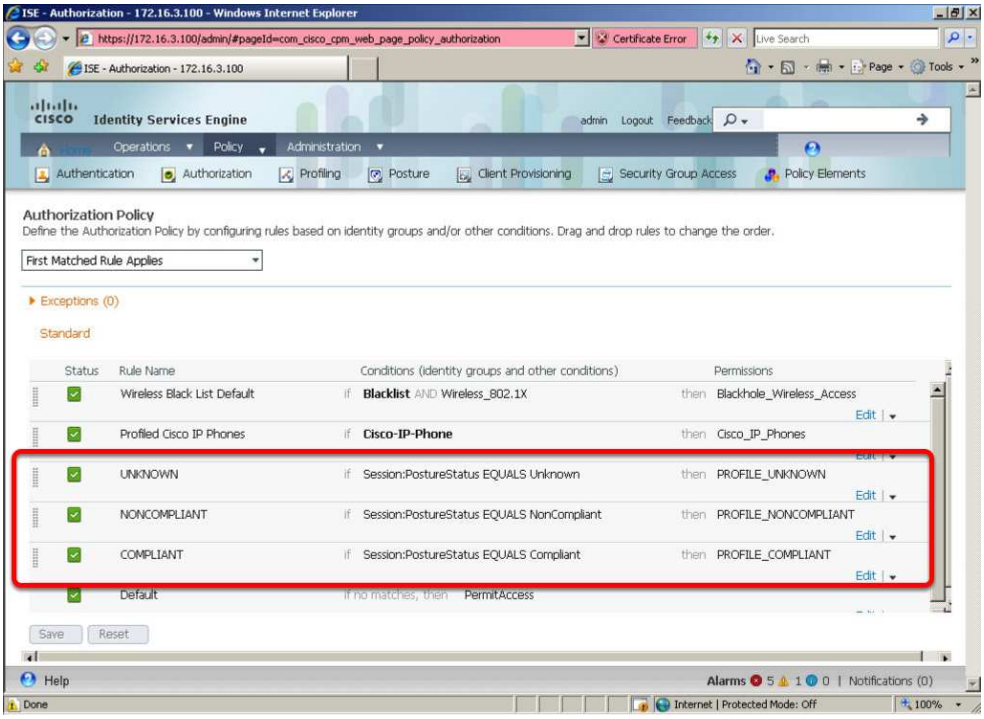


Configure the authorization profile for compliant posture status, and specify the dACL:

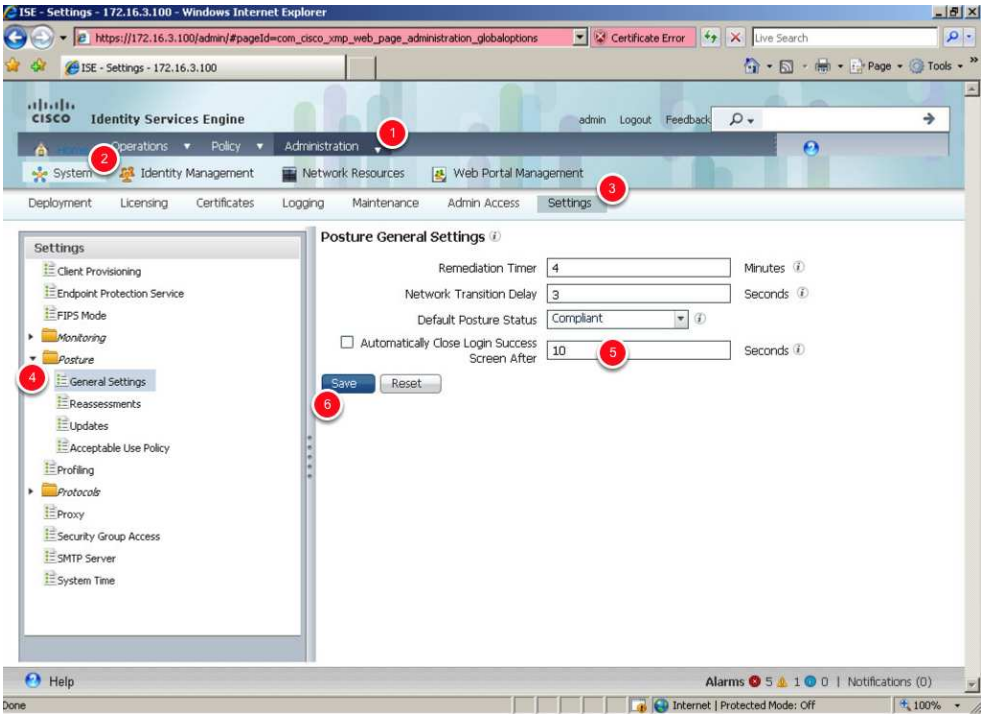


Configure authorization policies for all three posture results (unknown, noncompliant, and compliant):

https://t.me/learningnets

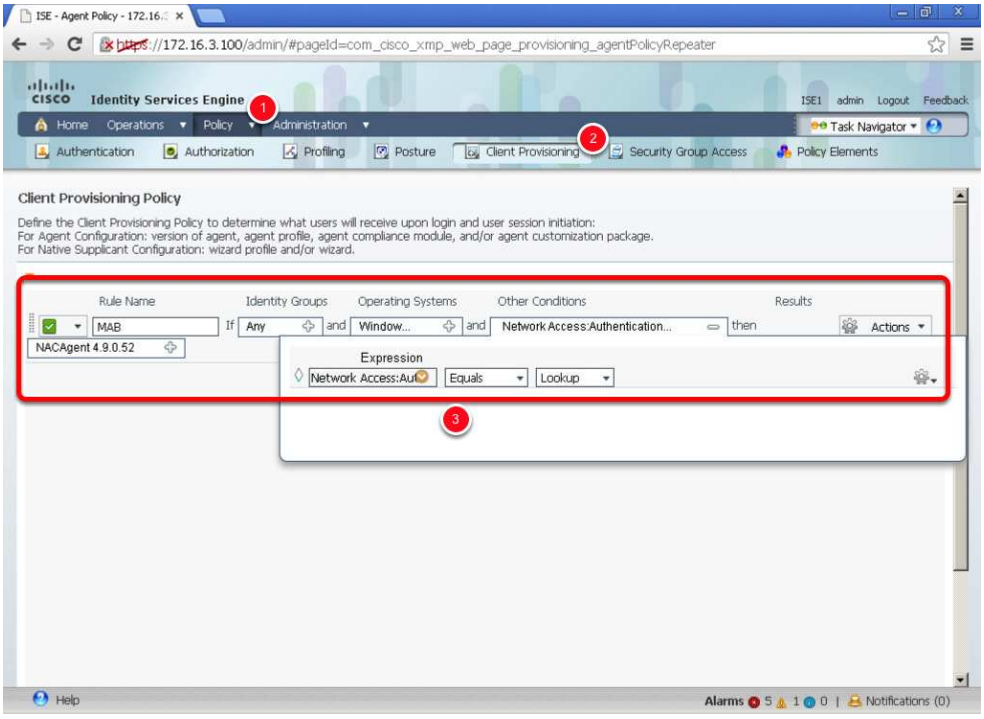


Configure NAC Agent to automatically close 10 seconds after user authentication:

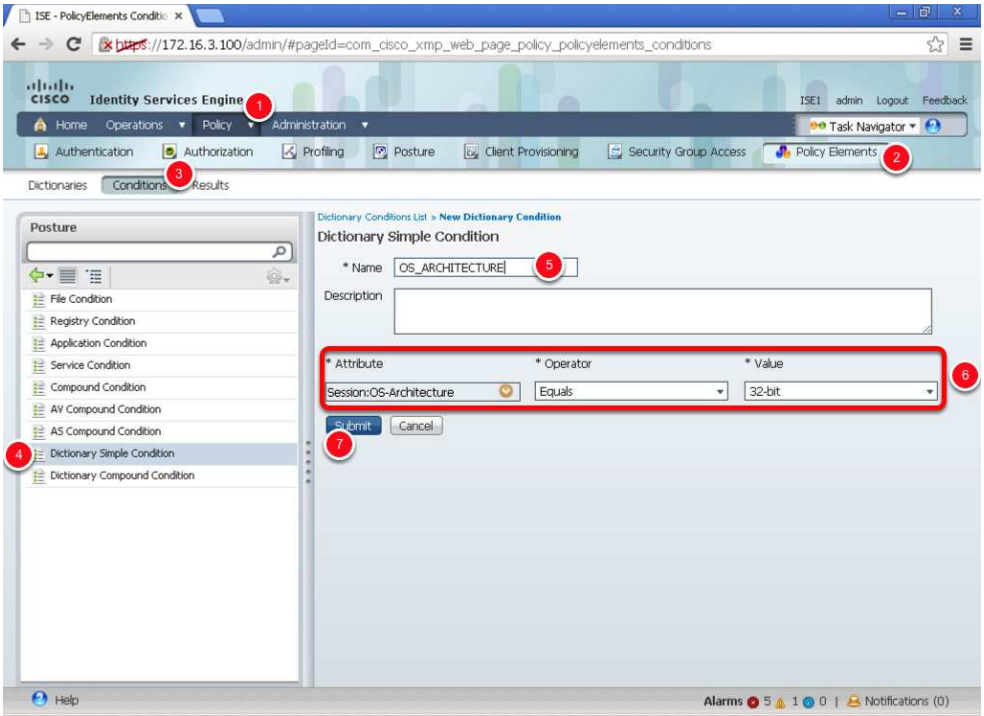


Define a Windows client provisioning policy valid only for MAB authentications (based on the default host lookup mechanism):

https://t.me/learningnets

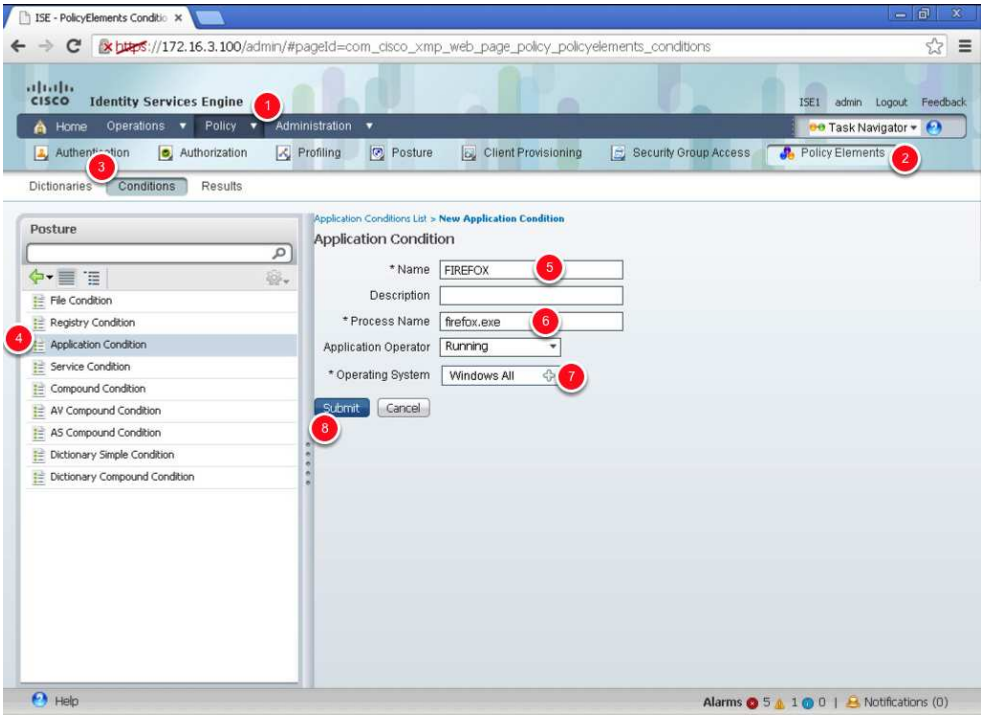


Define a posture condition for the operating system architecture:

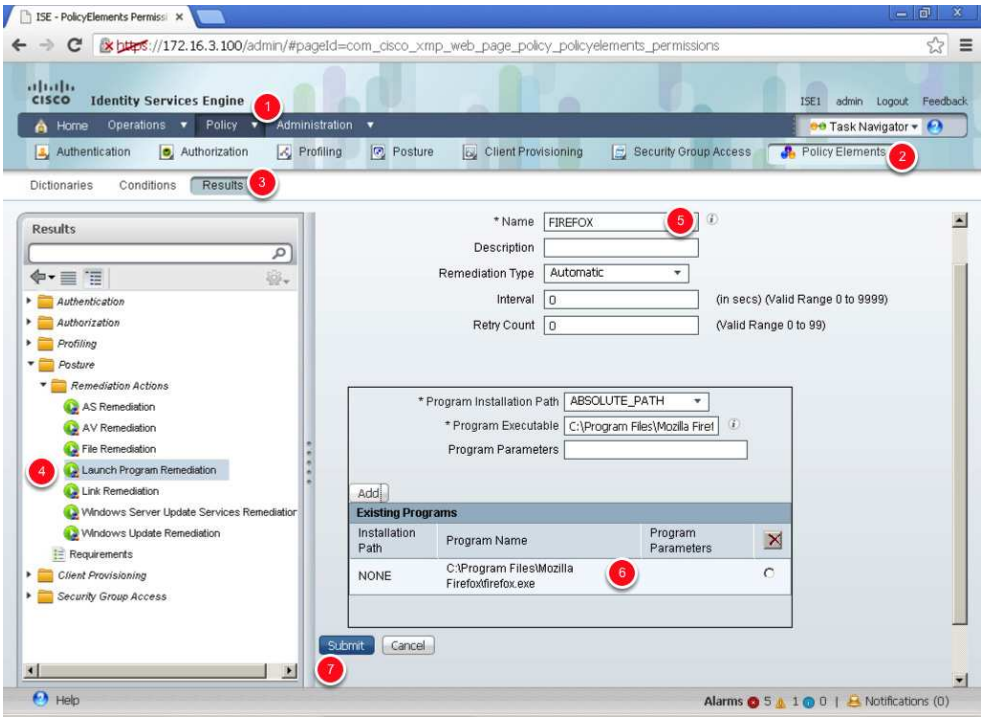


Define a Windows posture condition for the Firefox process:

https://t.me/learningnets

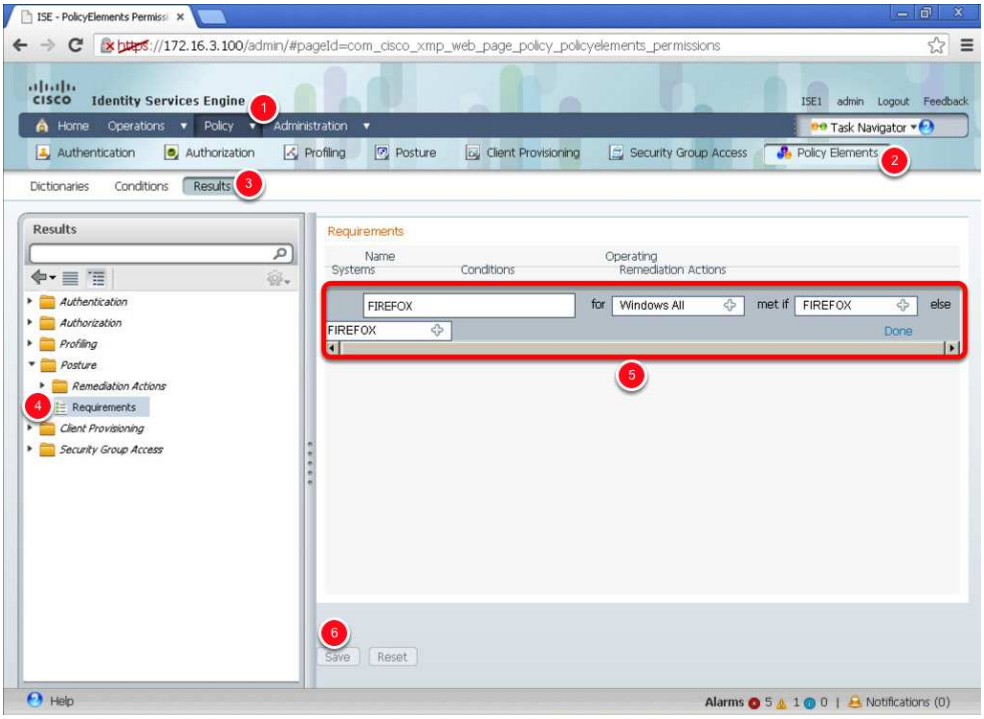


Define a posture remediation action that starts Firefox:

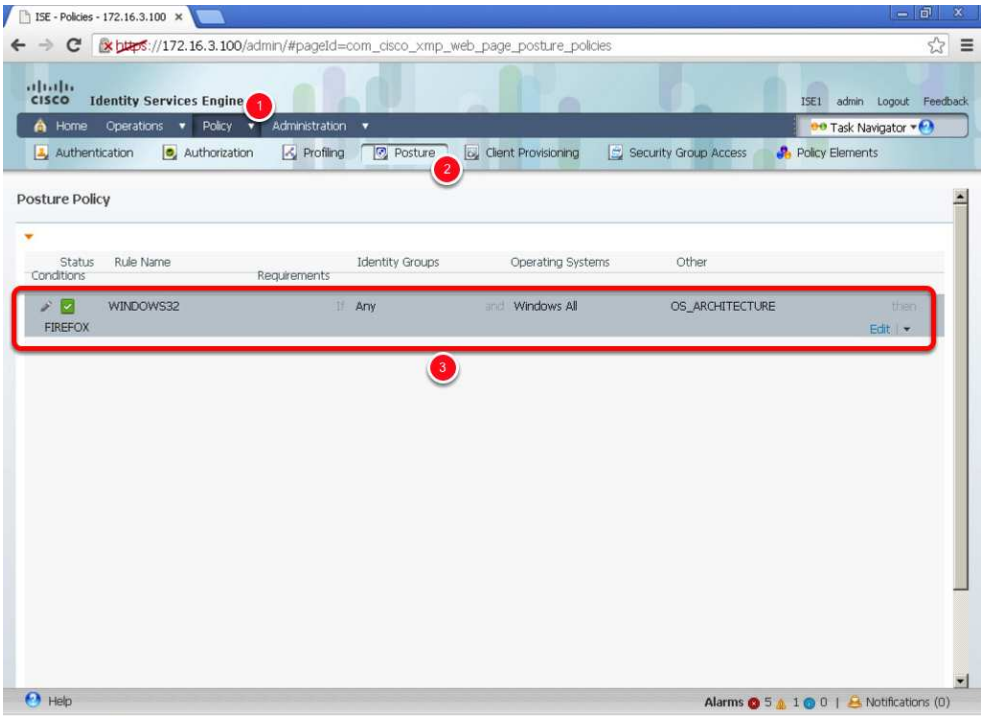


Define a Windows posture requirement that checks whether Firefox has started; if not, it will automatically start it through the previously configured remediation action:

https://t.me/learningnets



Define a Windows posture policy for 32-bit operating systems; to obtain posture-compliant status, the previosuly configured requirement of Firefox must be satisfied:



Verification

Note that the IP Phone will not be authenticated (its MC address was not added to ISE internal Endpoint store) and the switch will continuously try to authenticate it. Check that **TEST-PC-A** was successfully authenticated, matching the unknown posture authorization profile, and therefore the corresponding dACL and redirect

https://t.me/learningnets

URL have been received as part of authorization:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5 MAC Address:  48f8.b32e.2423
           IP Address:  Unknown
            User-Name:  48-F8-B3-2E-24-23 Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A ACS ACL:  xACSACLx-IP-ACL_UNKNOWN-529015b3
 URL Redirect ACL:  ACL_REDIRECT

 URL Redirect:  https://ISE1.inelab.local:8443/guestportal/gateway?sessionId=0000000000000038059E5061&action=cpp
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  0000000000000038059E5061
      Acct Session ID:  0x00000039
               Handle:  0x89000038

Runnable methods list:

       Method   State
       mab      Authc Success

----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown
            User-Name:  00036b3c35f0
               Status:  Authz Failed
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  0000000000000037059CF323
      Acct Session ID:  0x00000038
               Handle:  0xC4000037

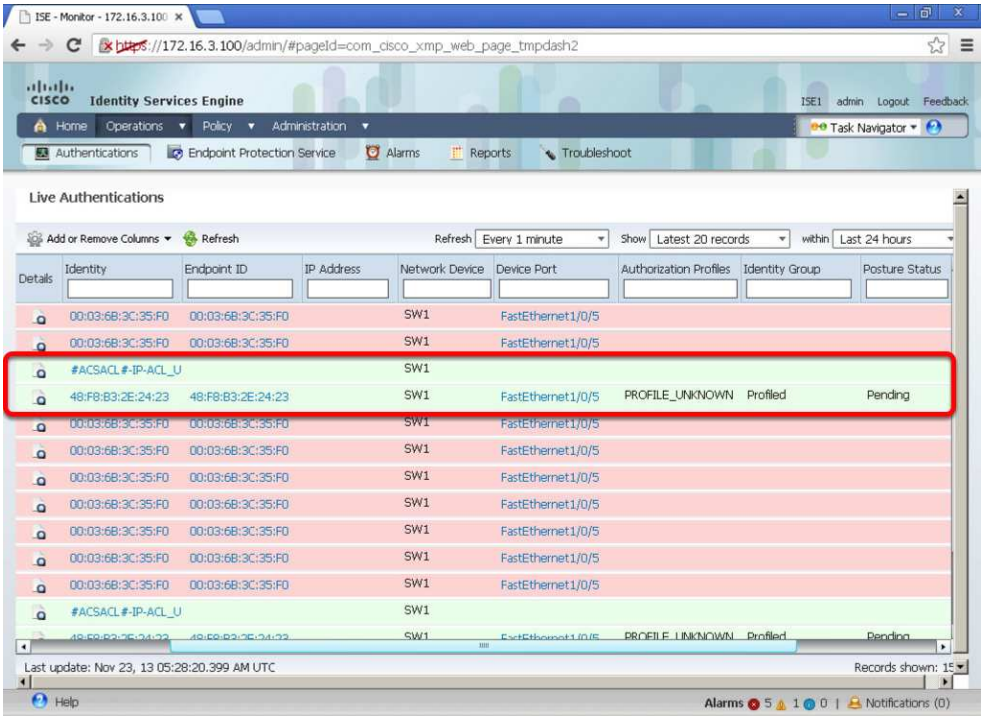
Runnable methods list:
       Method   State
       mab      Failed over
```

https://t.me/learningnets

```
!
!SW1# show ip access-lists xACSACLx-IP-ACL_UNKNOWN-529015b3
Extended IP access list xACSACLx-IP-ACL_UNKNOWN-529015b3 (per-user)
    10 permit udp any eq bootpc any eq bootps
    20 permit udp any host 172.16.10.100 eq domain
    30 permit tcp any host 172.16.3.100 eq 8443
    40 permit udp any host 172.16.3.100 eq 8905
    50 permit tcp any host 172.16.3.100 eq 8905
    60 permit udp any host 172.16.3.100 eq 8909
    70 permit tcp any host 172.16.3.100 eq 8909
    80 deny ip any any
!
!SW1#show ip access-lists interface fastEthernet1/0/5

     permit udp any eq bootpc any eq bootps (2 matches)
     permit udp any host 172.16.10.100 eq domain (53 matches)
     permit tcp any host 172.16.3.100 eq 8443
     permit udp any host 172.16.3.100 eq 8905
     permit tcp any host 172.16.3.100 eq 8905
     permit udp any host 172.16.3.100 eq 8909
     permit tcp any host 172.16.3.100 eq 8909
     deny ip any any (69 matches)
```

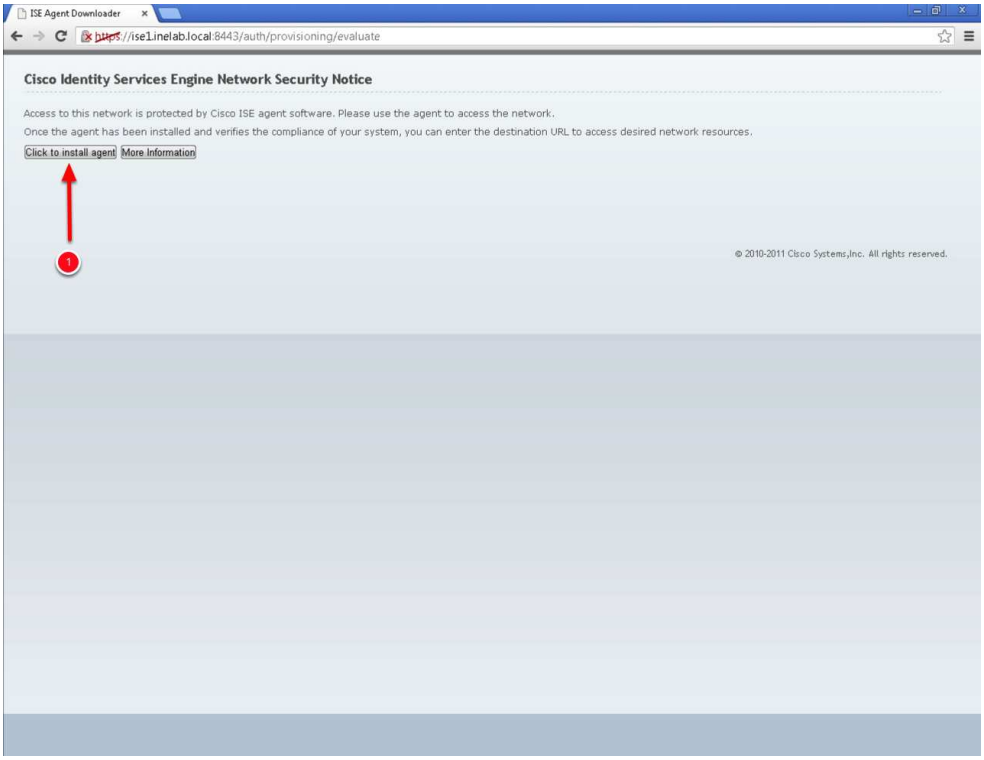
Check the authentication logs on ISE, and note which authorization profile has been matched and that posture status is pending (and therefore unknown):



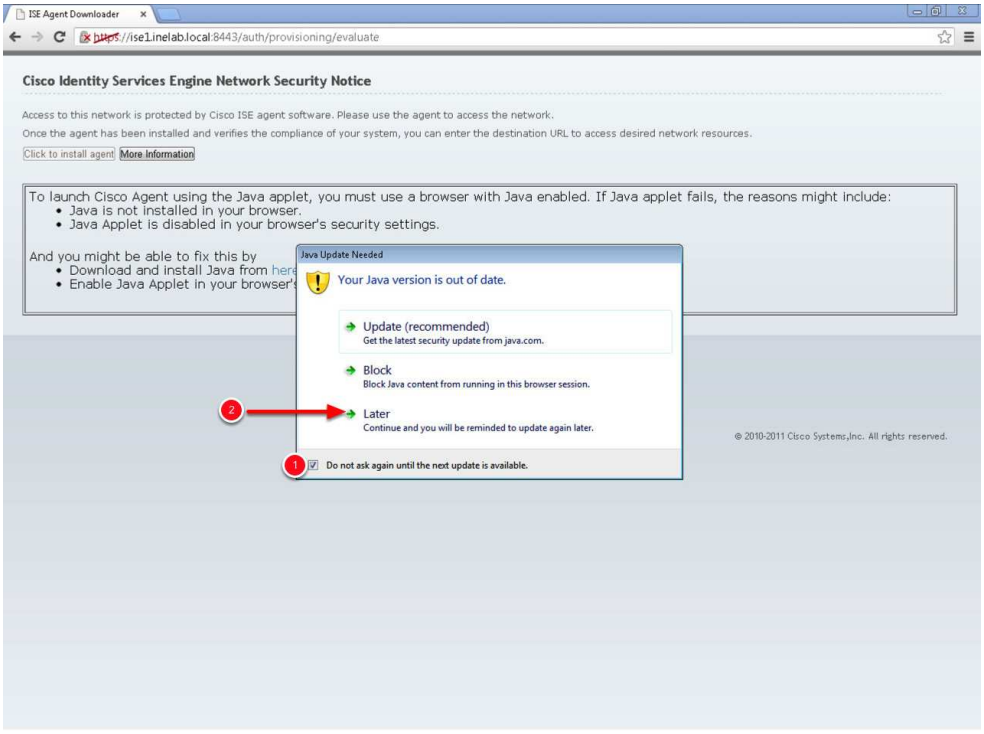
Go to **TEST-PC-A** , open Internet Explorer or Chrome, and try navigating to any

https://t.me/learningnets

destination; traffic will be redirected to the ISE portal for client provisioning based on the redirect ACL. Accept the warning for ISE certificate (it is not trusted because it is using the self-signed certificate), and click to install the NAC Agent:

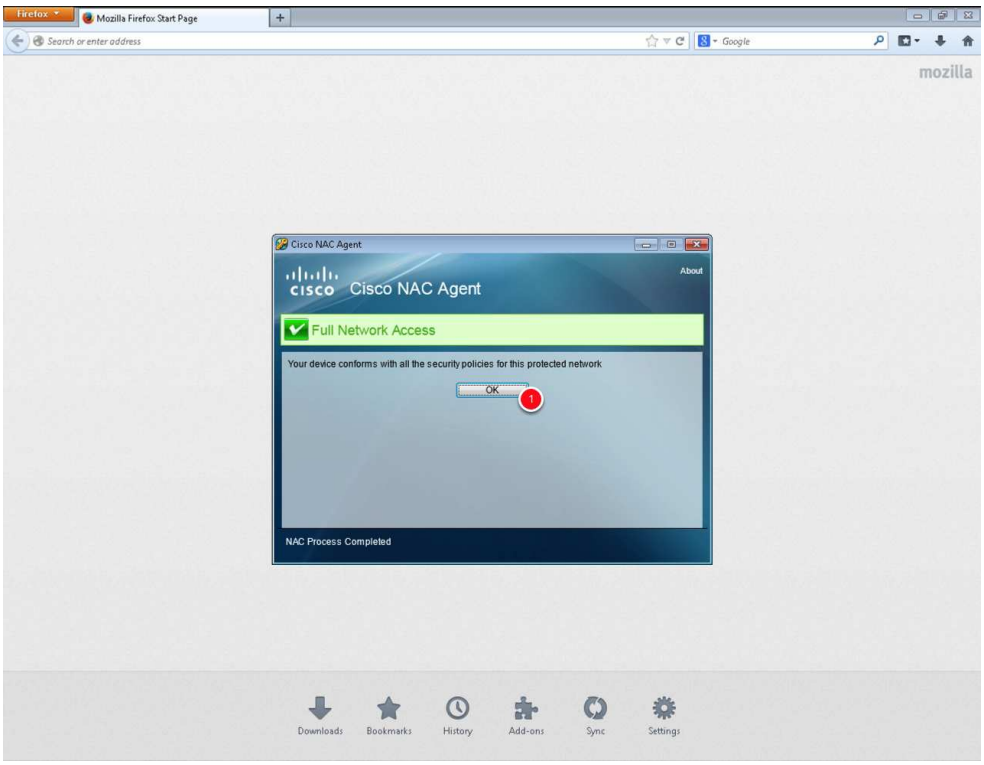


Ignore any errors and accept any warnings so that NAC Agent is installed; for example, if you get a Java warning, select the check box not to be asked again and select later for the update to be performed:

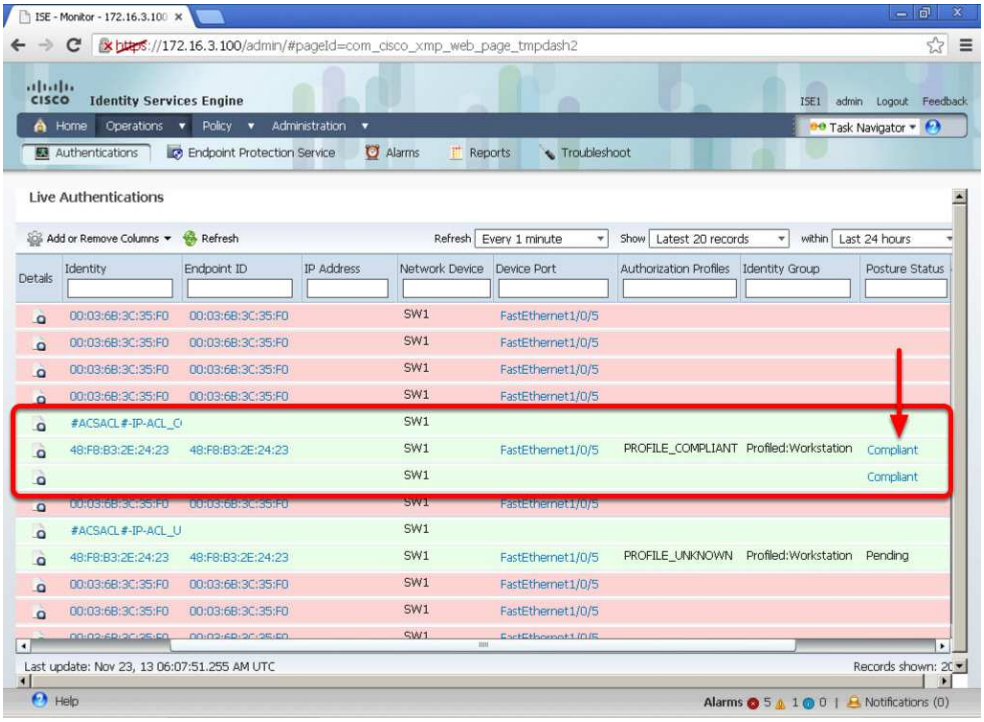


https://t.me/learningnets

After accepting all prompts (it will take few minutes), the NAC Agent installer is presented. Click Next on all screens; NAC Agent will detect that the Firefox process is not started, and it will automatically start it:

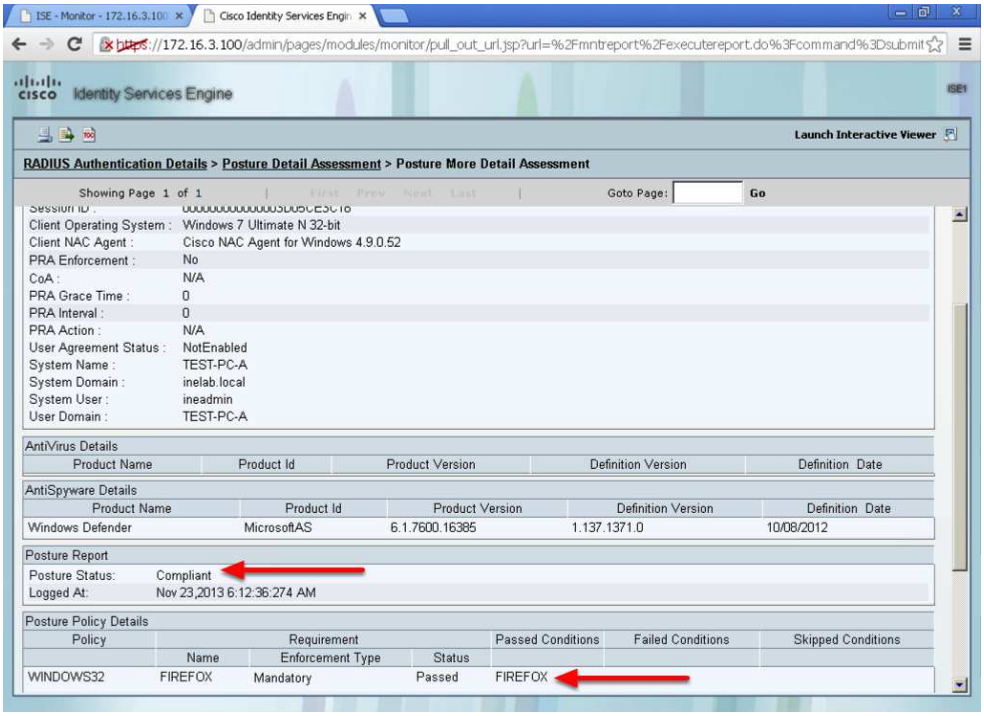


Verify the ISE authentication logs to check the matched authorization profile for posture-compliant status:



Go into details for the authentication logs to identify the posture assessment results:

https://t.me/learningnets



Verify the new authorization policy at switch level:

```
SW1#show authentication sessions interface fastEthernet1/0/5
            Interface:  FastEthernet1/0/5 MAC Address:  48f8.b32e.2423
           IP Address:  136.1.27.51
            User-Name:  48-F8-B3-2E-24-23 Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A ACS ACL:  xACSACLx-IP-ACL_COMPLIANT-529016b3
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000003D05CE3C18
      Acct Session ID:  0x0000003E
               Handle:  0x1600003D

Runnable methods list:
       Method   State
       mab      Authc Success


----------------------------------------
            Interface:  FastEthernet1/0/5
          MAC Address:  0003.6b3c.35f0
           IP Address:  Unknown
            User-Name:  00036b3c35f0
```

https://t.me/learningnets

```
               Status:  Authz Failed
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure
       Oper host mode:  multi-auth
     Oper control dir:  both
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  000000000000003E05CE55C9
      Acct Session ID:  0x0000003F
               Handle:  0xAB00003E

Runnable methods list:
       Method   State
       mab      Failed over
!
!SW1#show ip access-lists xACSACLx-IP-ACL_COMPLIANT-529016b3
Extended IP access list xACSACLx-IP-ACL_COMPLIANT-529016b3 (per-user)
    10 permit ip any any
!
!SW1#show ip access-lists interface fastEthernet1/0/5

     permit ip any any
```

Check IP connectivity with **TEST-PC-A**:

```
R1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.27.51         0148.f8b3.2e24.23       Nov 24 2013 06:12 AM    Automatic
136.1.47.51         0100.036b.3c35.f0       Nov 24 2013 05:17 AM    Automatic
!
!R1#ping 136.1.27.51
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.27.51, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/8 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

MACsec Downlink

Task

- Configure SW3 for MACsec on its Gi1/0/5 port:
  - Allow a single MAC address on the port.
  - Switch policy should not require MACsec, but authorization received from ISE should.
  - If **TEST-PC-B** is not configured for MACsec, allow it access to VLAN 17 if successfully authenticated.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

MACsec(802.11ae) adds security to wired networks by providing confidentiality (encryption) and integrity to data packets. If compared with other data protection technologies (IPsec, for example), MACsec provides it at layer 2 on a hop-by-hop basis; so on each layer 2 hop, packets are decrypted as received so that other features can be applied (such as QoS), and encrypted again before being sent out. MACsec works for both host-to-switch (downlink) and switch-to-switch (uplink) connections, providing linerate (1 or 10 GigabitEthernet) encryption through the use of Galois/Counter Mode Advanced Encryption Standard 128 (GCM-AES-128).

For host-to-switch, encryption can be performed at hardware level if the NIC of the

https://t.me/learningnets

host supports it; otherwise, it can be performed in software using Cisco AnyConnect. MACsec requires the host to first perform 802.1x authentication using an EAP method capable of negotiating the encryption keys: EAP-TLS, PEAP(MS-CHAPv2), EAP-FAST. To negotiate MACsec and derive the keying material, there are two supported protocols:

- Security Association Protocol (SAP), which is Cisco proprietary and used for switch-to-switch MACsec (trunk ports)
- MACsec Key Agreement (MKA), which is standard and used for host-to-switch MACsec (access ports)

For downlink MACsec, the following steps are required:

- Enable MACsec at the port level using the command `macsec .`
- Specify the MKA policy at the port level (the default can be used) with the command `mka default-policy .`
- Enable 802.1x globally on the switch, and configure 802.1x authentication/authorization with ISE/ACS.
- Configure MACsec policies on ISE/ACS.
- Optionally configure the MACsec policy; three options are available: must-secure, should-secure (default), and must-not-secure. Use the port-level command `authentication linksec policy [must-secure|should-secure|must-not-secure]` ; the locally configured setting can be overridden by the authorization received from ISE.

Depending on the host-mode of the port and MACsec policy, it might not work in all cases:

- MACsec is fully supported with host mode of single-mode and multi-domain; however, if the MACsec policy is set to must-secure, a host will not be allowed to pass traffic unless it successfully negotiates MACsec.
- MACsec is partially supported with host mode of multi-host; because only the first MAC address is authenticated (and thus MACsec policy can be negotiated), traffic from any new MAC addresses will be allowed to pass only if the MACsec policy is not set to must-secure.
- MACsec is not supported with host mode of multi-auth, mainly because MACsec is a point-to-point technology, whereas multi-auth is point-to-multipoint.

If the local MACsec policy (configured at switch level) or remote MACsec policy (received from ACS/ISE) is set to must-secure, but the client supplicant is not MACsec capable or has not been configured for MACsec, you can still authorize the host if successfully authenticated by using port-level command `authentication event linksec fail action authorize vlan <vlan_nr>`

https://t.me/learningnets

; this is known as MACsec fallback policy and only works if the port runs in must-secure mode.

Regardless of the MACsec configured policy, the following protocols are allowed in clear-text on the port: CDP, LLDP, EAP, LACP.

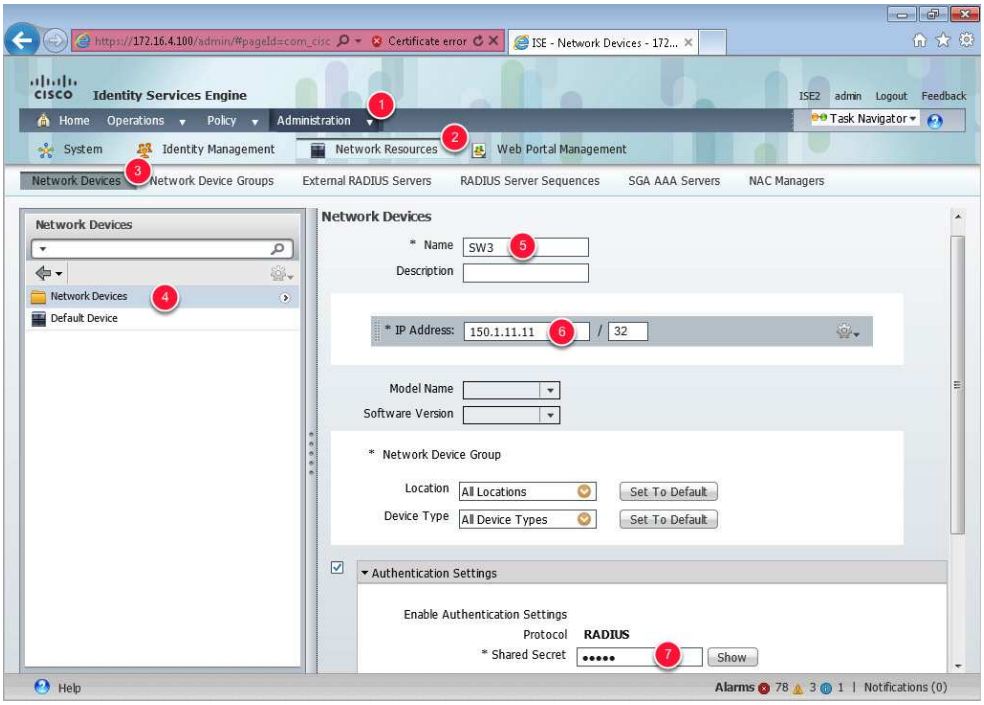
Configuration

SW3:

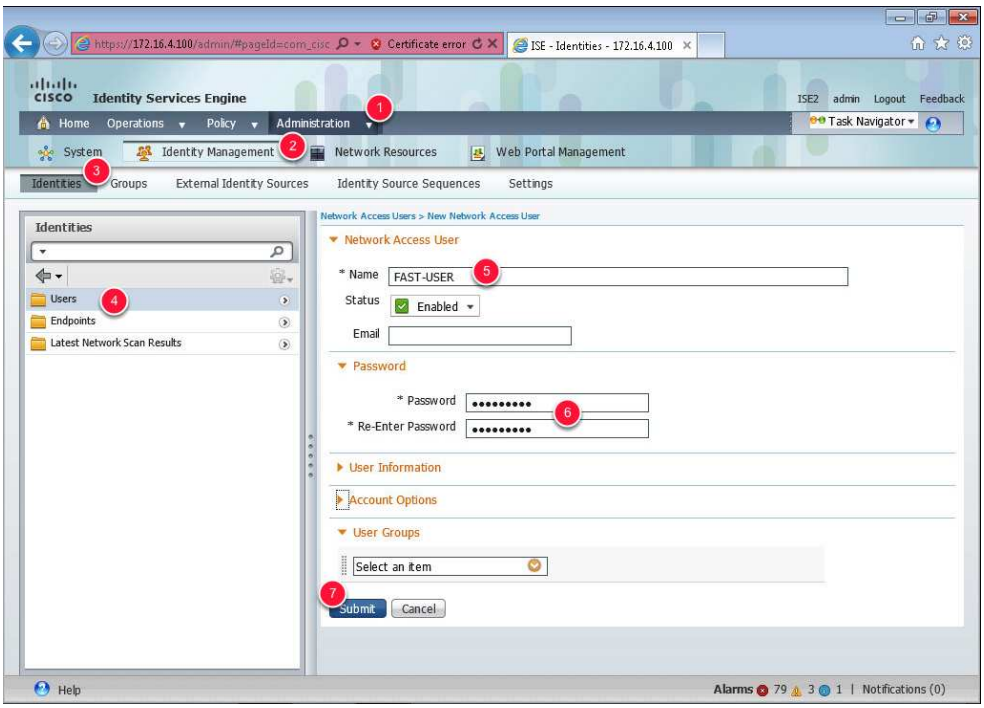
```
aaa new-model
aaa authentication dot1x default group ISE
aaa authorization network default group ISE
!
aaa group server radius ISE
 server-private 172.16.4.100 key CISCO
 ip radius source-interface Loopback0
!
dot1x system-auth-control
!
interface GigabitEthernet1/0/5
 authentication host-mode single-host
 authentication port-control auto
 dot1x pae authenticator
 macsec
 authentication linksec policy must-not-secure
 authentication event linksec fail action authorize vlan 17
 mka default-policy
```

Add SW3 as a NAD device into ISE:

https://t.me/learningnets

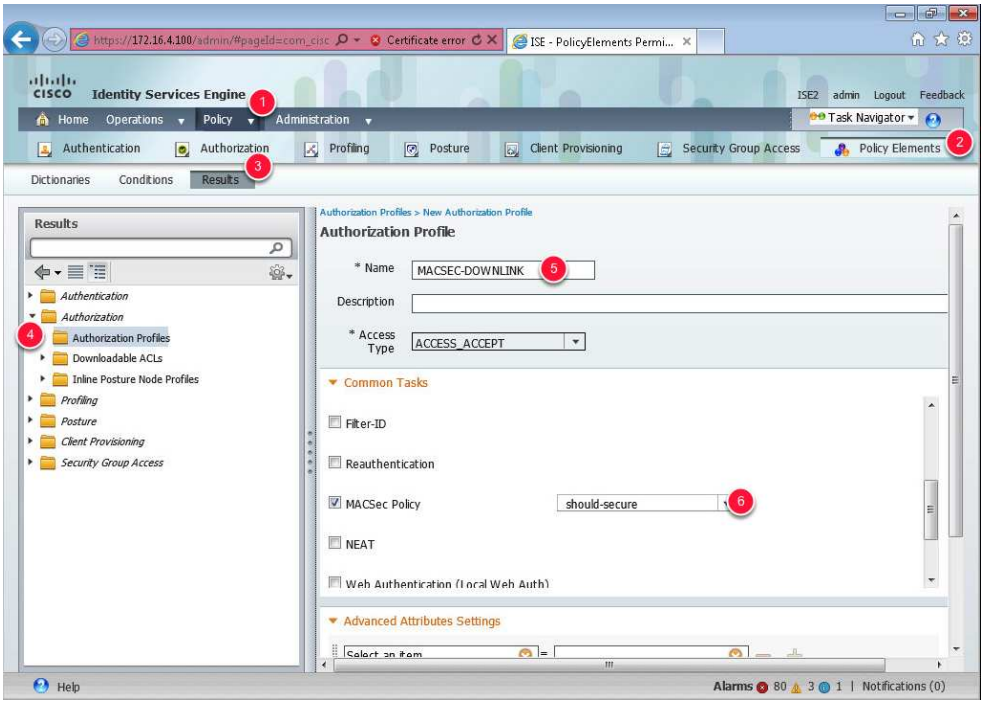


Add a user to the Internal Users identity store to be used for EAP-FAST(EAP-MSCHAPv2):

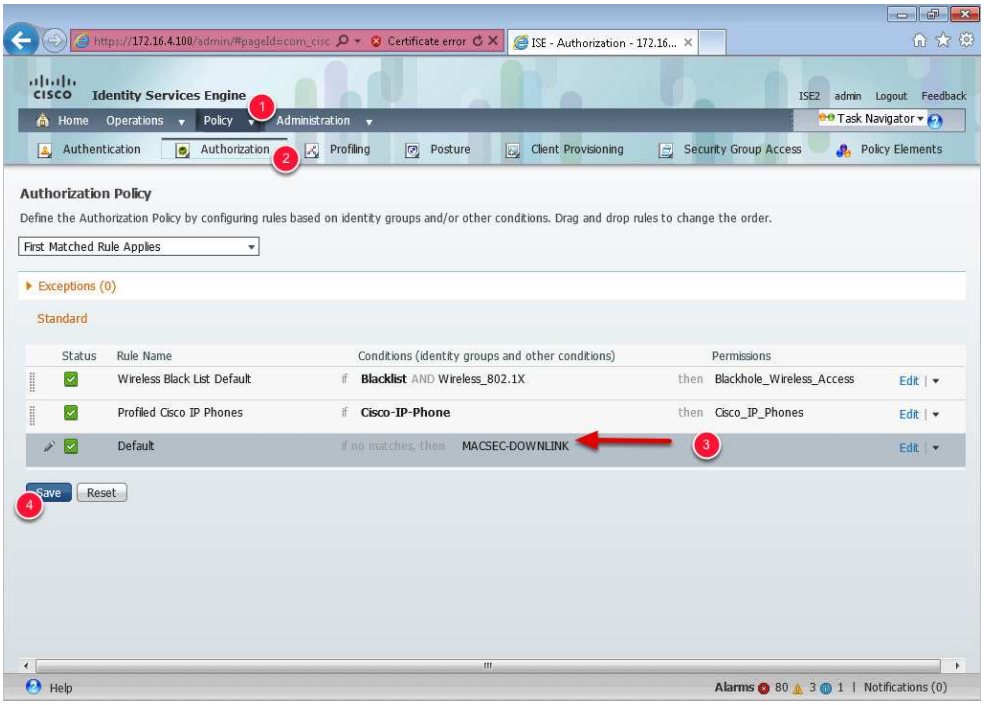


Configure an authorization profile for downlink MACsec, as requested by the task:

https://t.me/learningnets



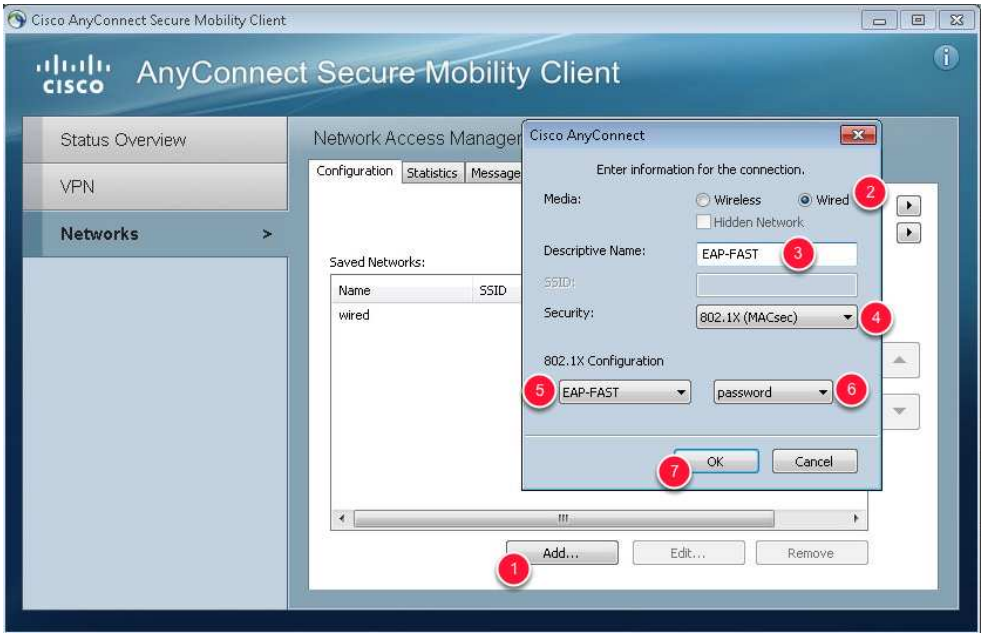
The default authentication and authorization policies will allow the EAP-FAST session to be authenticated and authorized. We must simply change the default authorization policy to apply the previously configured authorization profile (alternatively, a specific authorization policy for MACsec can be created, above the default one):



Because the native Windows supplicant does not support MACsec, and neither does EAP-FAST, AnyConnect must be used. First, enable and start the **Cisco AnyConnect Network Access Manager** service to enable the AnyConnect 802.1x supplicant service. To access the Windows services, type **services.msc**

https://t.me/learningnets

from the command prompt or Run. Open AnyConnect and navigate to the **Advanced** tab to configure an 802.1x MACsec profile for EAP-FAST:



Open AnyConnect and apply the configured profile. When asked by AnyConnect GUI, enter the username/password credentials configured on ISE. You will also be prompted to validate the ISE presented certificate. These credentials, as well as the CA to be trusted, can be preconfigured if the EAP-FAST profile is created using the **Network Access Manager Profile Editor**.

Verification

Verify that **TEST-PC-B** was authenticated and MACsec policy has been enforced by ISE:

```
SW3#show authentication interface gigabitEthernet1/0/5

Client list:
Interface  MAC Address     Method   Domain   Status         Session ID
Gi1/0/5    48f8.b32e.2532  dot1x    DATA     Authz Success  0000000000000000360992EE52

Available methods list:
  Handle  Priority  Name
    3        0      dot1x
Runnable methods list:
  Handle  Priority  Name
    3        0      dot1x
!
!SW3#show authentication sessions interface gigabitEthernet1/0/5
            Interface:  GigabitEthernet1/0/5
          MAC Address:  48f8.b32e.2532
```

https://t.me/learningnets

```
           IP Address:  Unknown User-Name:  FAST-USER
Status:  Authz Success
               Domain:  DATA Security Policy:  Should Secure
Security Status:  Secured
       Oper host mode:  single-host
     Oper control dir:  both Authorized By:  Authentication Server
          Vlan Policy:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  0000000000000000360992EE52
      Acct Session ID:  0x00000006
               Handle:  0x01000037

Runnable methods list:
       Method   State dot1x    Authc Success
```

Verify the MACsec status for the interface toward **TEST-PC-B**:

```
SW3#show macsec summary
 Interface                     Transmit SC          Receive SC
GigabitEthernet1/0/5                 1                    1
 !
!SW3#show macsec interface gigabitEthernet1/0/5
MACsec is enabled
  Replay protect : enabled
  Replay window : 0
  Include SCI : yes
  Cipher : GCM-AES-128
  Confidentiality Offset : 0
 Capabilities
  Max. Rx SA : 16
  Max. Tx SA : 16
  Validate Frames : strict
  PN threshold notification support : Yes
  Ciphers supported : GCM-AES-128 Transmit Secure Channels
SCI : 44D3CA6399850002
   Elapsed time : 00:05:33
   Current AN: 0   Previous AN: -
   SC Statistics
    Auth-only (0 / 0)
    Encrypt (815 / 0) Receive Secure Channels
SCI : 48F8B32E25320000
   Elapsed time : 00:05:33
   Current AN: 0   Previous AN: -
```

https://t.me/learningnets

```
   SC Statistics
    Notvalid pkts 0      Invalid pkts 0
    Valid pkts 416         Late pkts 0
    Uncheck pkts 0       Delay pkts 0
  Port Statistics
   Ingress untag pkts  0        Ingress notag pkts 32767
   Ingress badtag pkts 0        Ingress unknownSCI pkts 0
   Ingress noSCI pkts 0         Unused pkts 0    Notusing pkts 0 Decrypt bytes 132135

   Ingress miss pkts 32767
```

Verify the MKA status for port Gi1/0/5:

```
SW3#show mka sessions interface gigabitEthernet1/0/5 detail

MKA Detailed Status for MKA Session
===================================Status: SECURED - Secured MKA Session with MACsec

Local Tx-SCI..... 44d3.ca63.9985/0002
Interface MAC Address.... 44d3.ca63.9985
MKA Port Identifier...... 2
Interface Name........... GigabitEthernet1/0/5
Audit Session ID......... 000000000000000360992EE52
CAK Name (CKN)........... 546D322AFEA56687736F262DBC3852A1
Member Identifier (MI)... AEA09E4B9891B819F62B1EA1
Message Number (MN)...... 481
Authenticator............ YES
Key Server............... YES

Latest SAK Status........ Rx & Tx
Latest SAK AN............ 0
Latest SAK KI (KN)....... AEA09E4B9891B819F62B1EA100000001 (1)
Old SAK Status........... FIRST-SAK
Old SAK AN............... 0
Old SAK KI (KN).......... FIRST-SAK (0)

SAK Transmit Wait Time... 0s (Not waiting for any peers to respond)
SAK Retire Time.......... 0s (No Old SAK to retire)

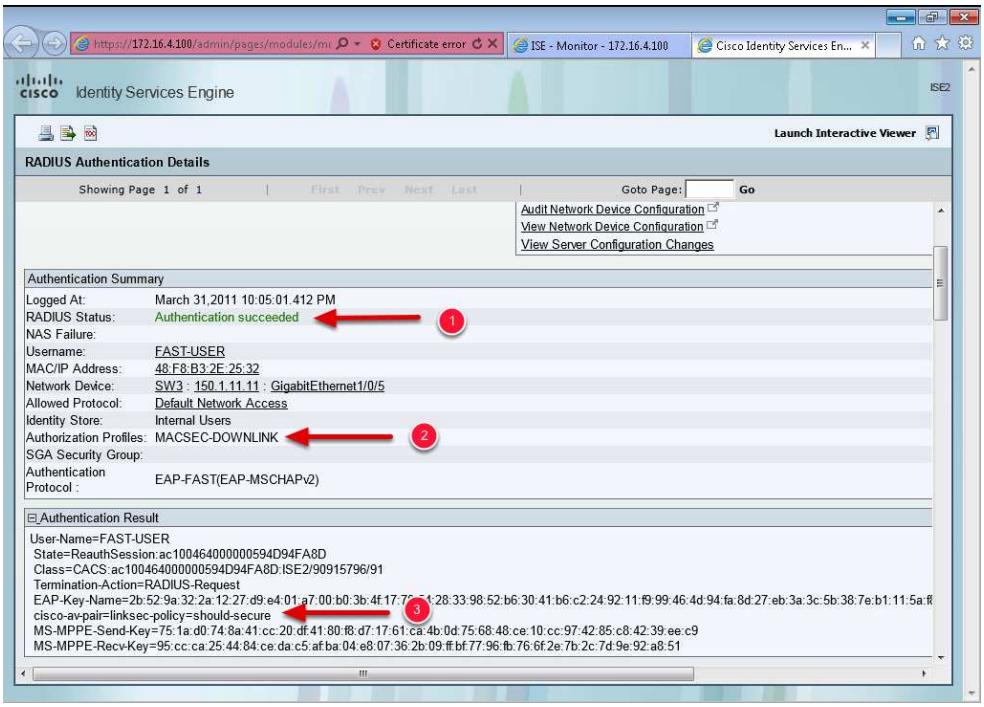
MKA Policy Name.......... *DEFAULT POLICY*
Key Server Priority...... 0
Delay Protection......... NO
Replay Protection........ YES
Replay Window Size....... 0
Confidentiality Offset... 0
```

https://t.me/learningnets

```
Algorithm Agility........ 80C201
Cipher Suite............. 0080020001000001 (GCM-AES-128)
MACsec Capability........ 3 (MACsec Integrity, Confidentiality, & Offset)MACsec Desired.......... YES
 # of MACsec Capable Live Peers............ 1
# of MACsec Capable Live Peers Responded.. 1
 Live Peers List:
   MI                          MN          Rx-SCI (Peer)
   -----------------------------------------------------------
66E2A3F3B9F4900632B9F404  442          48f8.b32e.2532/0000


Potential Peers List:
   MI                          MN          Rx-SCI (Peer)
   -----------------------------------------------------------
```

Verify the ISE Authentication Logs and that MACsec policy has been enforced; CAK keys are sent by ISE to the switch, so that SAK key can be derived:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 4: Identity Management

MACsec Uplink Manual

Task

- Configure SW3 and SW4 to protect packets on the connecting trunk using MACsec:
  - Ensure both integrity and confidentiality.
  - All necessary settings are local to the switches.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 4 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Protecting a link between two switches using MACsec is known as Uplink MACsec, and it can be performed manually (configurations are performed locally on the switches) or dynamically (configurations are partially locally, but the MACsec policy is retrieved from a remote ISE/ACS server). This is very similar to manual IPsec (keying material is statically defined and never expires) and dynamic IPsec (keying material is dynamically negotiated and has a limited lifetime). Negotiation between switches uses SAP protocol.

SAP is a key derivation and exchange protocol based on a draft of IEEE 802.11i, which performs the following functions:

- Negotiates cipher suite for data traffic.
- Derives session keys for data traffic.

https://t.me/learningnets

- Exchanges SCIs (Secure Channel Identifier) that will be used by data traffic.
- Ensures that the exchange is being performed with the same devices that participated in authentication.

For the manual method, you must first enable CTS in manual mode on the interface using the command `cts manual` . Second, specify the PMK and functional mode using the command `sap pmk <pmk_value> mode-list [gcm-encrypt|gmac|no-necap|null]` . Multiple modes can be configured so negotiation occurs between peers, with the most secure method winning. The default is no-encap, which means the interface does not support MACsec, so packets are not modified. The difference between gcm-encrypt and gmac is that gcm-encrypt performs both authentication and encryption, and gmac performs only authentication. The Pairwise Master Key (PMK) is a hexadecimal value that must match on both sides and is used to derive the keying material, similar to how a preshared key is used in IPsec.

> When MACsec is used, an additional layer2 header and trailer are inserted, and the Ethertype value is changed to 0x88e5.

Configuration

```
SW3:
interface GigabitEthernet1/0/22
 cts manual
  sap pmk 101 mode-list gcm-encrypt


SW4:

interface GigabitEthernet1/0/22
 cts manual
  sap pmk 101 mode-list gcm-encrypt
```

Verification

Verify the CTS configuration:

```
SW3#show cts interface summary

Global Dot1x feature is Disabled


CTS Layer2 Interfaces
```

https://t.me/learningnets

```
--------------------
Interface  Mode    IFC-state dot1x-role peer-id    IFC-cache
-----------------------------------------------------------
Gi1/0/22   MANUAL  OPEN      unknown    unknown    invalid

CTS Layer3 Interfaces
--------------------
Interface   IPv4 encap       IPv6 encap       IPv4 policy      IPv6 policy
-------------------------------------------------------------------------
!
!SW3#show cts interface gigabitEthernet1/0/22
Global Dot1x feature is Disabled
Interface GigabitEthernet1/0/22: CTS is enabled, mode:    MANUAL
    IFC state:               OPEN
    Authentication Status:   NOT APPLICABLE
        Peer identity:       "unknown"
        Peer's advertised capabilities: "sap"
    Authorization Status:    NOT APPLICABLE SAP Status:               SUCCEEDED
        Version:             2 Configured pairwise ciphers:
 gcm-encrypt

        Replay protection:      enabled
        Replay protection mode: STRICT
 Selected cipher:        gcm-encrypt


    Propagate SGT:           Enabled
    Cache Info:
        Cache applied to link : NONE

    Statistics:
        authc success:              0
        authc reject:               0
        authc failure:              0
        authc no response:          0
        authc logoff:               0
        sap success:                1
        sap fail:                   0
        authz success:              0
        authz fail:                 0
        port auth fail:             0

    L3 IPM:   disabled.
```

https://t.me/learningnets

Verify the MACsec status for the interface between the switches:

```
SW3#show macsec summary
 Interface                     Transmit SC          Receive SC
GigabitEthernet1/0/22                1                    1
 !
 !SW3#show macsec interface gigabitEthernet1/0/22
 MACsec is enabled
   Replay protect : enabled
   Replay window : 0
   Include SCI : yes
   Cipher : GCM-AES-128
   Confidentiality Offset : 0
  Capabilities
   Max. Rx SA : 16
   Max. Tx SA : 16
   Validate Frames : strict
   PN threshold notification support : Yes
   Ciphers supported : GCM-AES-128 Transmit Secure Channels
 SCI : 44D3CA6399960000
    Elapsed time : 00:07:17
    Current AN: 0   Previous AN: -
    SC Statistics
     Auth-only (0 / 0)
     Encrypt (1261 / 0) Receive Secure Channels
 SCI : 70810539FE960000
    Elapsed time : 00:07:17
    Current AN: 0   Previous AN: -
    SC Statistics
     Notvalid pkts 0      Invalid pkts 0
     Valid pkts 164         Late pkts 0
     Uncheck pkts 0       Delay pkts 0
   Port Statistics
    Ingress untag pkts  0        Ingress notag pkts 13603
    Ingress badtag pkts 0        Ingress unknownSCI pkts 0
    Ingress noSCI pkts 0         Unused pkts 0    Notusing pkts 0 Decrypt bytes 17139

    Ingress miss pkts 13603
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

Section 5 Introduction

> This lab utilizes **Firewall Services**.

Lab Topics

The ASA Firewall is an integral part of the CCIE Security lab exam. From a high level, the topics in the blueprint include:

- Cisco ASA Firewall
  - Basic firewall Initialization
  - Device management
  - Address translation (nat, global, static)
  - Access Control Lists
  - IP routing/Route Tracking
  - Object groups
  - VLANs
  - Configuring Etherchannel
  - High Availability and Redundancy
  - Layer 2 Transparent Firewall
  - Security contexts (virtual firewall)
  - Modular Policy Framework
  - Identity Firewall Services
  - Configuring ASA with ASDM
  - Context-aware services
  - IPS capabilities
  - QoS capabilities

We can further expand these topics in the following checklist as follows:

https://t.me/learningnets

Technology	Mastered?
Initializing the Basic Cisco ASA Firewall (IP Address, Mask, Default Route, etc.)	
Understanding Security Levels (Same Security Interface)	
Understanding Single vs. Multimode	
Understanding Firewall vs. Transparent Mode	
Understanding Multiple Security Contexts	
Understanding Shared Resources for Multiple Contexts	
Understanding Packet Classification in Multiple-Contexts Mode	
VLAN Subinterfaces Using 802.1Q Trunking	
Multiple-Mode Firewall with Outside Access	
Single-Mode Firewall Using the Same Security Level	
Multiple-Mode, Transparent Firewall	
Single-Mode, Transparent Firewall with NAT	
ACLs in Transparent Firewall (for Pass-Through Traffic)	

https://t.me/learningnets

Technology	Mastered?
Understanding How Routing Behaves on the Adaptive Security Appliance (Egress and Next-Hop Selection Process)	
Understanding Static vs. Dynamic Routing	
Static Routes	
RIP with Authentication	
OSPF with Authentication	
EIGRP with Authentication	
Managing Multiple Routing Instances	
Redistribution Between Protocols	
Route Summarization	
Route Filtering	
Static Route Tracking Using an SLA	
Dual ISP Support Using Static Route Tracking	
Redundant Interface Pair	
LAN-Based Active/Standby Failover (Routed Mode)	
LAN-Based Active/Active Failover (Routed Mode)	

https://t.me/learningnets

Technology	Mastered?
LAN-Based Active/Standby Failover (Transparent Mode)	
LAN-Based Active/Active Failover (Transparent Mode)	
Stateful Failover Link	
Device Access Management	
Enabling Telnet	
Enabling SSH	
The nat-control Command vs. no nat-control Command	
Enabling Address Translation (NAT, Global, and Static) Pre & Post 8.4	
NAT Objects	
Context-Aware firewall	
Identity Firewall	
Using ASDM and Cisco Prime	
Policy NAT	
Destination NAT	

https://t.me/learningnets

Technology	Mastered?
Bypassing NAT When NAT Control Is Enabled Using Identity NAT	
Bypassing NAT When NAT Control Is Enabled Using NAT Exemption	
Port Redirection Using NAT	
Tuning Default Connection Limits and Timeouts	
Basic Interface Access Lists and Access Group (Inbound and Outbound)	
Time-Based Access Lists	
ICMP Commands	
Enabling Syslog and Parameters	
NTP with Authentication	
Object Groups (Network, Protocol, ICMP, and Services)	
Nested Object Groups	
URL Filtering	
Java Filtering	
ActiveX Filtering	

https://t.me/learningnets

Technology	Mastered?
ARP Inspection	
Modular Policy Framework (MPF)	
Application-Aware Inspection	
Identifying Injected Errors in Troubleshooting Scenarios	
Understanding and Interpreting Adaptive Security Appliance show and debug Outputs	
Understanding and Interpreting the packet-tracer and capture Commands	

Configurations

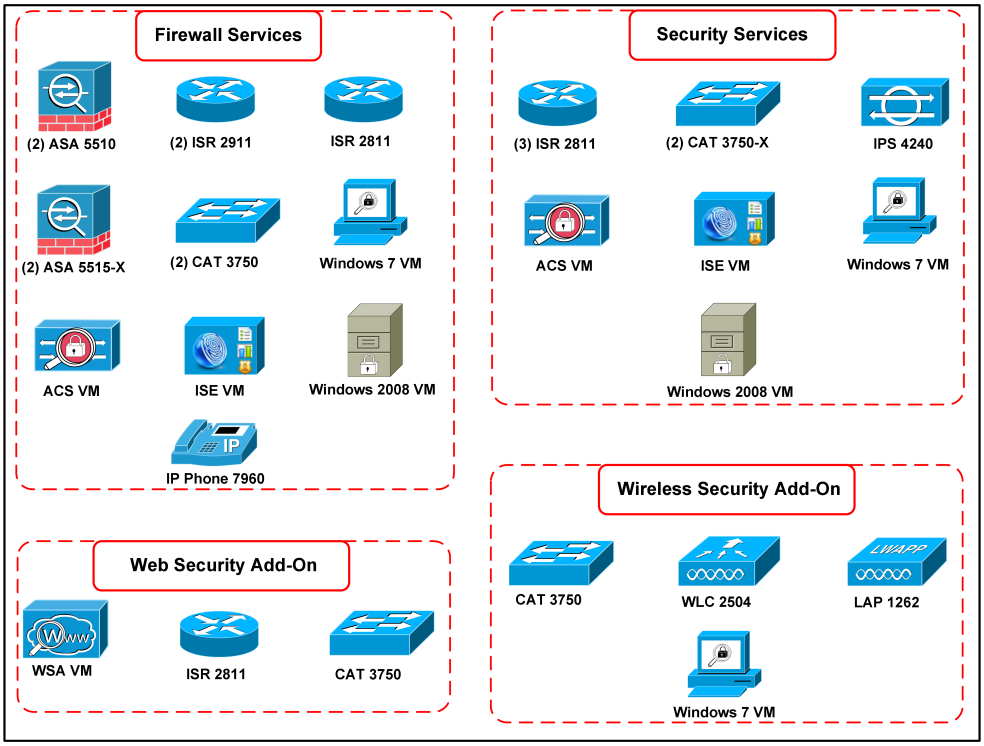
Initial configuration files are provided at the top of this page. Each lab has an initial configuration that is assumed. The initial configuration in some cases is the configuration that the previous lab ended with. If there is only one configuration file, the other devices did not change in the previous lab.

Load the Section 5 Initial Configuration Files to initialize your rack. Pay attention to the instructions in each task, because you may be required to load a different set of initial configuration files.

Section 5 General Lab Topology

The INE CCIE Security Racks are divided into modules as seen in the figure below. This section, Section 5, utilizes Firewall Services.

https://t.me/learningnets



In this section, you will focus on ASA firewalls and how to provide perimeter security services with them. As you work on each task, focus on the small elements of building each technology. Gaining an in-depth knowledge of the workings of small tasks will lead to better troubleshooting of large scenarios in the CCIE Lab exam, as well in as production environments. Pay close attention to the topology, which may change.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124

https://t.me/learningnets

- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Basic Configuration

Task

- Configure ASA3 and ASA4 interfaces as follows:

Device	Hostname	Interface	Nameif	Security-Level	IP Address
ASA3	ASA3	Gi0/0.34	OUTSIDE	10	136.1.34.17/24
ASA3	ASA3	Gi0/3	INSIDE	60	136.1.93.17/24
ASA4	ASA4	Gi0/0.34	VLAN34	DEFAULT	136.1.34.18/24
ASA4	ASA4	Gi0/1	VLAN94	DEFAULT	136.1.94.18/24

- Without configuring any access-list rules, SW1 and SW2 should successfully telnet to R3.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.

Overview

Since version 7.0 of the ASA code, configuring interfaces in the firewall appliance is

https://t.me/learningnets

very similar to configuring interfaces in IOS-based platforms. If the firewall connection to the switch is an 802.1q trunk (the ASA supports 802.1q only, not ISL), you can create sub-interfaces, corresponding to the VLANs carried over the trunk. Do not forget to assign a VLAN number to the sub-interface. The native (untagged) VLAN of the trunk connection maps to the physical interface, and it cannot be assigned to a sub-interface.

When configuring interfaces on the ASA in routed mode, the following can be configured:

- **Nameif** configured with the command `nameif <NAME>` ; it is mandatory and gives the interface a logical name. Without it, even if the interface is in the UP/UP state, it cannot be used for traffic forwarding.
- **IP Address** configured with the command `ip address <IP> <MASK> secondary <IP> <MASK>` ; it is mandatory.
- **Enable Interface** configured with the command `no shutdown` ; it is mandatory.
- **Security Level** configured with the command `security level <0-100>` ; it is optional and assigns the interface a level of trust, based on which there are some implicit firewall rules.
- **Speed and Duplex** configured with the commands `speed [10|100|1000|auto]` and `duplex [half|full|auto]` ; it is optional, and by default all interfaces are set to **auto** negotiating both speed and duplex.

By default, based on the configured **nameif**, ASA assigns the following implicit security-levels to interfaces:

- 100 to a nameif of **inside**.
- 0 to a nameif of **outside**.
- 0 to all other nameifs.

Without any configured access-lists, ASA implicitly allows or restricts traffic flows based on the security-levels:

https://t.me/learningnets

- Traffic from high-security level to low-security level is allowed by default (for example, from 100 to 0, or in our case from 60 to 10).
- Traffic from low-security level to high-security level is denied by default; to allow traffic in this direction, an ACL must be configured and applied (at interface level or global level).
- Traffic between interfaces with an identical security level is denied by default (for example, from 20 to 20, or in our case from 0 to 0); to allow traffic in this direction, the command `same-security-traffic permit inter-interface` must be configured.

The **Management0/0** interface is used by default only for management purposes, so it can only terminate management-plane/control-plane sessions, and not forwarding-plane sessions (no traffic can be passed through the interface); this is because of the default interface-level command `management-only` . On ASA 5500 series, the **Management0/0** interface can be configured to pass traffic through, but not on ASA 5500-X series; this is because outof-band management for the IPS software module or CX module is done through the **Management0/0** interface.

Configuration

SW1:

```
interface FastEthernet1/0/16
 switchport mode access
 switchport access vlan 93
 no shutdown
!
interface FastEthernet1/0/19
 switchport mode access
 switchport access vlan 94
 no shutdown
```

SW2:

```
interface range FastEthernet1/0/16 , FastEthernet1/0/18
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
```

ASA3:

```
hostname ASA3
!
interface GigabitEthernet0/0
```

https://t.me/learningnets

```
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
 nameif OUTSIDE
 security-level 10
 ip address 136.1.34.17 255.255.255.0
!
interface GigabitEthernet0/3
 nameif INSIDE
 security-level 60
 ip address 136.1.93.17 255.255.255.0
 no shutdown
```

ASA4:

```
hostname ASA4
!
interface GigabitEthernet0/0
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
 nameif VLAN34
 ip address 136.1.34.18 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/1
 nameif VLAN94
 ip address 136.1.94.18 255.255.255.0
 no shutdown
!
same-security-traffic permit inter-interface
```

Verification

Verify the nameifs and security-levels:

```
ASA3# show nameif
Interface                Name                    Security
GigabitEthernet0/0.34    OUTSIDE                  10
GigabitEthernet0/1       INSIDE                   60
!
!ASA4# show nameif
```

https://t.me/learningnets

```
Interface                Name                   Security
GigabitEthernet0/0.34    VLAN34                   0
GigabitEthernet0/1       VLAN94                   0
```

Verify the correct IP Address and Subnet Mask assignment; look for the **Current IP Address**; the **System IP Address** will be discussed in the failover section:

```
ASA3# show ip address
System IP Addresses:
Interface                Name                   IP address      Subnet mask     Method
GigabitEthernet0/0.34    OUTSIDE                136.1.34.17     255.255.255.0   manual
GigabitEthernet0/1       INSIDE                 136.1.93.17     255.255.255.0   manual
Current IP Addresses:
Interface                Name                   IP address      Subnet mask     Method
GigabitEthernet0/0.34    OUTSIDE                136.1.34.17     255.255.255.0   manual
GigabitEthernet0/1       INSIDE                 136.1.93.17     255.255.255.0   manual
!
!ASA4# show ip address
System IP Addresses:
Interface                Name                   IP address      Subnet mask     Method
GigabitEthernet0/0.34    VLAN34                 136.1.34.18     255.255.255.0   manual
GigabitEthernet0/1       VLAN94                 136.1.94.18     255.255.255.0   manual
Current IP Addresses:
Interface                Name                   IP address      Subnet mask     Method
GigabitEthernet0/0.34    VLAN34                 136.1.34.18     255.255.255.0   manual
GigabitEthernet0/1       VLAN94                 136.1.94.18     255.255.255.0   manual
```

Verify basic IP connectivity with the both ASA3 and ASA4:

```
SW1#ping 136.1.93.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.93.17, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
!
!SW2#ping 136.1.94.18

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.94.18, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
!
!R3#ping 136.1.34.17

Type escape sequence to abort.
```

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 136.1.34.17, timeout is 2 seconds:.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
!
!R3#ping 136.1.34.18

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.34.18, timeout is 2 seconds:.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Verify that telnet works from both SW1 and SW2, and check the connection state on ASA:

```
SW1#telnet 136.1.34.3
Trying 136.1.34.3 ... Open



User Access Verification

Password: R3>
!
!ASA3# show conn

1 in use, 1 most used
TCP OUTSIDE 136.1.34.3:23 INSIDE 136.1.93.9:21366, idle 0:00:23, bytes 111, flags UIO
!
!SW2#telnet 136.1.34.3

Trying 136.1.34.3 ... Open



User Access Verification

Password: R3>
!
!ASA4# show conn
1 in use, 1 most used
TCP VLAN94 136.1.94.10:44119 VLAN34 136.1.34.3:23, idle 0:00:04, bytes 111, flags UIOB
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Interface High Availability

Task

- Configure ASA3 as follows:

Device	Hostname	Interface	Nameif	Security-Level	IP Address
ASA3	ASA3	Gi0/0 & Gi0/2	OUTSIDE	DEFAULT	136.1.34.17/24
ASA3	ASA3	Gi0/1 & Gi0/3	INSIDE	DEFAULT	136.1.93.17/24

- Interfaces Gi0/0 and Gi0/2 should work in active-standby mode.
- Interfaces Gi0/1 and Gi0/3 should work in active-active mode; allow no more than two interfaces in the bundle and load balance traffic on the ASA based on layer3/layer4 information.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by

https://t.me/learningnets

the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

ASA supports interface high availability and redundancy when operating in routed mode:

- In active-standby mode, where one interface is actively forwarding traffic and one interface is in standby mode; this is the **redundancy** feature.
- In active-active mode, where up to eight interfaces can be actively used for traffic forwarding, sharing the load; this is the **etherchannel** feature.

A logical redundant interface is a pair of one active and one standby physical interface. When the active interface fails, the standby interface becomes active. From the perspective of the firewall applications, this event is completely transparent, because the interface pair is viewed as a single logical interface. You can use redundant interfaces to increase the security appliance reliability. This feature is separate from device-level failover, but you can configure redundant interfaces as well as failover if desired. You can configure up to 8 redundant interfaces. Notice that this method has a limitation; the detection of interface failure is based on simple physical monitoring. When the primary interface loses a carrier, it is declared as failed.

When adding physical interfaces to the redundant pair, make sure they are not configured using the nameif and are in no shutdown state; any ASA interface can be a member of a redundant group, except for the Management0/0 interface. This is just a precaution; the firewall will remove all interface settings when adding the physical interface to a redundant group. The logical redundant interface will take the MAC address of the first interface added to the group, because this will also become the active interface. This MAC address is not changed with the member interface failures, but changes when you swap the order of the physical interfaces added to the pair; optionally, a vMAC can be configured for the redundant interface. With redundant interfaces, the **nameif**, **security-level**, and **IP address** configuration is done at the logical interface level. This feature is not preemptive; for example, if the initial active interface failed and the standby interface became active, when the initial active interface became up again, it would remain in standby state. The active interface can be changed at any time using the command `redundant-interface redundant <nr> active-member <interface>` .

Starting with code version 8.4(1), LACP/802.3ad support was added to the ASA; this

https://t.me/learningnets

is an enhancement to the redundant interface function, because all interfaces within a group (up to eight) can actively forward traffic. ASA supports both **active** and **passive** modes, where **active** initiates the LACP negotiation, and **passive** expects to receive LACP negotiations. Optionally, the etherchannel can be configured in **on** mode, also known as static mode, because there is no negotiation of the etherchannel state between peers. The same restrictions as configuring etherchannels on IOS code apply as well; all interfaces must be of same type and speed and use the same configuration. By default, ASA will load-balance the traffic across etherchannel interfaces based on layer3 information. Optionally, layer2 and layer4 can be taken into consideration, and this is configurable on a per-etherchannel level. The logical portchannel interface will take the MAC address of the lowest number interface from the group; optionally, a vMAC can be configured for the etherchannel interface.

Configuration

SW1:

```
interface FastEthernet1/0/16
channel-group 1 mode active
no shutdown
!
interface FastEthernet1/0/17
channel-group 1 mode active
no shutdown
!
interface Port-Channel 1
switchport mode access
switchport access vlan 93
```

SW2:

```
interface FastEthernet1/0/16
switchport mode access
switchport access vlan 34
no shutdown
!
interface FastEthernet1/0/17
switchport mode access
switchport access vlan 34
no shutdown
```

ASA3:

https://t.me/learningnets

```
hostname ASA3
!
interface GigabitEthernet0/0
no shutdown
!
interface GigabitEthernet0/2
no shutdown
!
interface Redundant 1
member-interface GigabitEthernet0/0
member-interface GigabitEthernet0/2
nameif OUTSIDE
ip address 136.1.34.17 255.255.255.0
!
interface GigabitEthernet0/1
channel-group 1 mode passive
no shutdown
!
interface GigabitEthernet0/3
channel-group 1 mode passive
no shutdown
!
interface Port-Channel 1
lacp max-bundle 2
port-channel load-balance src-dst-ip-port
nameif INSIDE
ip address 136.1.93.17 255.255.255.0
```

Verification

Verify the nameif and IP addressing:

https://t.me/learningnets

```
ASA3# show nameif

Interface                Name                    Security
Port-channel1            INSIDE                  100
Redundant1               OUTSIDE                 0

!
!ASA3# show ip address

System IP Addresses:
Interface                Name                IP address      Subnet mask     Method
Port-channel1            INSIDE              136.1.93.17     255.255.255.0   manual
Redundant1               OUTSIDE             136.1.34.17     255.255.255.0   manual
Current IP Addresses:
Interface                Name                IP address      Subnet mask     Method
Port-channel1            INSIDE              136.1.93.17     255.255.255.0   manual
Redundant1               OUTSIDE             136.1.34.17     255.255.255.0   manual
```

Verify the redundant and etherchannel configuration:

```
ASA3# show interface redundant 1 | b Redundancy
  Redundancy Information: Member GigabitEthernet0/0(Active), GigabitEthernet0/2
        Last switchover at 20:40:01 UTC Nov 4 2013
!
!ASA3# show port-channel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        U - in use      N - not in use, no aggregation/nameif
        M - not in use, no aggregation due to minimum links not met
        w - waiting to be aggregated
Number of channel-groups in use: 1
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Po1(U)            LACP    Gi0/1(P)  Gi0/3(P)
!
!ASA3# show port-channel 1 load-balance
EtherChannel Load-Balancing Configuration:
        src-dst-ip-port

EtherChannel Load-Balancing Addresses UsedPer-Protocol:
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination IP address and TCP/UDP (layer-4) port number
```

https://t.me/learningnets

```
  IPv6: Source XOR Destination IP address and TCP/UDP (layer-4) port number
!
!ASA3# show port-channel 1 brief
Ports: 2  Maxports = 16
Port-channels: 1 Max Port-channels = 48
Protocol: LACP/ passive
Minimum Links: 1Maximum Bundle: 2
```

Load balance: src-dst-ip-port

Verify connectivity with the redundant interface and check that the assigned MAC is the one of the active interface, the one first attached in the group:

```
R3#ping 136.1.34.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.34.17, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R3#show ip arp 136.1.34.17
Protocol  Address          Age (min)  Hardware Addr  Type  Interface
Internet  136.1.34.17            19a493.4caa.f508
  ARPA  FastEthernet0/0
!
!ASA3# show interface gigabitEthernet 0/0 | i MAC
        MAC address a493.4caa.f508
, MTU not set
```

Verify connectivity with the etherchannel interface and that the assigned MAC is the one from the lowest interface in the group:

```
SW1#ping 136.1.93.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.93.17, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
!
!SW1#show ip arp 136.1.93.17
Protocol  Address          Age (min)  Hardware Addr  Type  Interface
Internet  136.1.93.17            44a493.4caa.f505
  ARPA  Vlan93
!
!ASA3# show interface gigabitEthernet 0/1 | i MAC
```

https://t.me/learningnets

```
        MAC address a493.4caa.f505
, MTU 1500
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Reliable Static Routing

Task

- Configure a Loopback2 interface with an IP address of 150.1.111.111/32.
- ASA should have a route for this address as long as the following conditions are met:
  - It has IP connectivity with R1's VLAN 19 IP address.
  - This check should be performed each 5 seconds, with an answer received within 1 second; monitor whether the RTT exceeds 100 ms.
  - Use a total packet payload of 1400 bytes.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

https://t.me/learningnets

ASA supports static routing for both IPv4 and IPv6, but reliable static routing only for IPv4. Reliable static routing means that a condition is attached to the static route and the route is valid (present in the routing table) only if the condition is true. For the condition, ASA supports a subset of the SLA (Service Level Agreement) feature from the IOS code, based solely on ICMP Echo Request packets for which an ICMP Echo Reply packet is expected. Configuration requires the following steps:

1. Configure the SLA, which identifies what IP address is monitored for ICMP connectivity and any optional parameters; it's defined with the command `sla monitor <nr>` , followed by parameters:
   - `type echo protocol icmpecho <IP Address> interface <nameif>`
   - `frequency <seconds>` , how often ICMP Echo Requests packets are generated.
   - `timeout <milliseconds>` , the interval in which an ICMP Echo Reply is expected; otherwise, the monitored destination is considered unreachable.
   - `threshold <milliseconds>` , monitoring the RTT interval; it has no impact on the SLA state.
   - `request-data-size <bytes>` , the payload size for the ICMP Echo Request, useful in identifying the end-to-end MTU (PMTU process); ASA will always add an additional 8 bytes in the payload for the timestamp data.
2. Activate the SLA. Most of the times it's scheduled to always run using the command `sla monitor schedule <nr> life forever start-time now .`
3. Configure the tracking object using the command `track <obj_nr> rtr <sla_nr> reachability` , which will have a value of **TRUE** if the operation return code of the SLA is **OK** (there is IP connectivity with the monitored destination) or a value of **FALSE** if the operation return code of the SLA is **TIMEOUT** (there is no IP connectivity with the monitored destination).
4. Configure tracking static route using the command `route <nameif> <subnet> <mask> <next-hop> track <obj_nr>` , which will be valid only if the tracking object value is **TRUE**.

Basically, SLA monitors connectivity with a given destination (usually a next-hop, directly connected), tracking object monitors the SLA return code, and static route monitors the tracking object state. This feature is useful in dual-ISP scenarios, where, for example, you have two Internet connections with static default routes, but you want to failover from the primary path to the secondary path if the primary next-hop becomes unreachable; so you configure a primary route (lower AD) with track for next-hop reachability, and if the tracker fails, the secondary static route (higher AD) will preempt the primary one and the traffic will flow via the backup path.

ASA supports ECMP (Equal Cost Multi Path) routing only if all next-hops for any given destination are reachable over the same interface. This restriction is because

https://t.me/learningnets

of the internal ASA architecture and how packets processing is implemented, NAT playing a major role. To route packets, ASA does not rely only on the routing table, which behaves differently than the IOS code. To detect the exit interface for a packet, ASA performs the following checks:

- If an active XLATE exists for the destination IP address of the packet, the interface is selected from the respective XLATE; this can be overridden so that a lookup in the routing table is performed only if the `route-lookup` keyword is specified when defining the NAT rule (supported after 8.3 code).
- If no active XLATE exists for the destination IP address of the packet but the flow matches on a configured static NAT entry for the destination IP address, the interface is selected from the respective NAT entry; this can be overridden so that a lookup in the routing table is performed only if the `route-lookup` keyword is specified when defining the NAT rule (supported after 8.3 code).
- If neither of the above are true, it looks in the routing table to find the best match (most specific entry) for the destination IP address, and the interface is selected from the routing entry.

When the exit interface is selected, it needs to find the next-hop IP address, which must be reachable over the selected exit interface or the packet will be dropped. For example, if from the exit interface selection process the packet ends up in the Gi0/0 interface buffer, but from the next-hop IP address selection this is reachable over interface Gi0/1, the packet is dropped.

Configuration

```
R1:
interface Loopback2
 ip address 150.1.111.111 255.255.255.255


ASA2:

sla monitor 2
 type echo protocol ipIcmpEcho 136.1.19.1 interface VLAN19
 request-data-size 1392
 timeout 1000
 threshold 100
 frequency 5
!
sla monitor schedule 2 life forever start-time now
!
```

https://t.me/learningnets

```
track 100 rtr 2 reachability
route VLAN19 150.1.111.111 255.255.255.255 136.1.19.1 track 100
```

Verification

Verify SLA configuration and state:

https://t.me/learningnets

```
ASA2# show sla monitor configuration

SA Agent, Infrastructure Engine-II
Entry number: 2
Owner:
Tag: Type of operation to perform: echo
Target address: 136.1.19.1

Interface: VLAN19
Number of packets: 1Request size (ARR data portion): 1392
Operation timeout (milliseconds): 1000

Type Of Service parameters: 0x0
Verify data: NoOperation frequency (seconds): 5

Next Scheduled Start Time: Start Time already passed
Group Scheduled : FALSE
Life (seconds): Forever
Entry Ageout (seconds): never
Recurring (Starting Everyday): FALSE
Status of entry (SNMP RowStatus): Active
Enhanced History:
!
!ASA2# show sla monitor operational-state

Entry number: 2
Modification time: 14:40:53.300 UTC Mon Feb 3 2003
Number of Octets Used by this Entry: 1660
Number of operations attempted: 7472
Number of operations skipped: 0
Current seconds left in Life: Forever
Operational state of entry: Active
Last time this entry was reset: Never
Connection loss occurred: FALSE
Timeout occurred: FALSEOver thresholds occurred: FALSE
Latest RTT (milliseconds): 1

Latest operation start time: 01:03:28.311 UTC Tue Feb 4 2003Latest operation return code: OK

RTT Values:
RTTAvg: 1       RTTMin: 1       RTTMax: 1
NumOfRTT: 3     RTTSum: 3       RTTSum2: 3
```

https://t.me/learningnets

Verify object state, routing table entry, and IP connectivity:

```
ASA2# show track
Track 100
  Response Time Reporter 2 reachability
  Reachability is Up
  9 changes, last change 10:21:29 Latest operation return code: OK
  Latest RTT (millisecs) 1
  Tracked by:
    STATIC-IP-ROUTING 0
!
!ASA2# show route VLAN19 150.1.111.111 255.255.255.255

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set
S    150.1.111.111 255.255.255.255 [1/0] via 136.1.19.1, VLAN19
!
!ASA2# ping 150.1.111.111
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.111.111, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Perform a **debug icmp trace** and **debug track** on the ASA, and shut down the Gi0/0 interface on R1 to see the packet payload size and object state changes:

```
ASA2:
ASA2# debug icmp trace
debug icmp trace enabled at level 1
ASA2# debug track
!
!R1:
interface gigabitEthernet0/0
 shutdown
!
!ASA2:
ICMP echo request from 136.1.19.16 to 136.1.19.1 ID=8603 seq=37095 len=1400
```

https://t.me/learningnets

```
ICMP echo reply from 136.1.19.1 to 136.1.19.16 ID=8603 seq=37095 len=1400
 ICMP echo request from 136.1.19.16 to 136.1.19.1 ID=27526 seq=53183 len=1400
Track: 100 Change #14 rtr 2, reachability Up->Down
!
!ASA2# show route VLAN19 150.1.111.111 255.255.255.255

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set
ASA2#
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA RIPv2 Routing

Task

- Configure RIPv2 between ASA2, R1, and R2, and advertise all networks.
- Ensure that R1, R2, and ASA2 have specific routes for all prefixes.
- Do not send or accept any RIP updates over VLAN 26 interface.
- ASA generate a default route with a metric of 10.
- Authenticate RIP packets from R1 using the string **WEAK** and RIP packets from R2 using a hash of the string **STRONG**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

ASA supports RIPv1 and RIPv2 dynamic routing (a single process) only for IPv4, with following steps (optional or mandatory):

- **Enable RIP process (mandatory)**; configured with the global command `router rip` , it globally activates RIP. Version1 or version2 can be configured using the process-level command `version [1|2]` .

https://t.me/learningnets

- **Disable auto-summary (mandatory)**; configured with the process-level command `no auto-summary` , it disables the classful network boundary auto-summarization function of RIP, which by default is enabled.
- **Enable interfaces for RIP (mandatory)**; configured with the process-level command `network <classful_network>` , it activates sending and receiving of RIP multicast updates on interfaces matching the command. Also, subnets matched by the command are included in the RIP updates/advertisements.
- **Define passive interfaces (optional)**; configured with the process-level command `passive-interface <nameif>` , it disables sending of RIP updates, but received updates will still be allowed and processed.
- **Enable authentication (optional)**; configured with the interface-level commands `rip authentication mode [text|md5]` and `rip authentication key 0 <key_string> key_id <key_nr>` , it enables clear-text or MD5 authentication for all RIP updates on respective interfaces. A single key can be configured per-interface. If received, RIP updates are not authenticated corerctly; these will be discarded.
- **Originate default route (optional)**; configured with the process-level command `default-information originate [route-map <name>]` , it activates sending of default route in RIP updates. Route-map is optional and can be used to set the metric value.
- **Filter routing updates**; configured with the process-level command `distribute-list <acl_name> [in|out] interface <nameif>` , it filters inbound or outbound updates according to the rules from the ACL. Only standard ACL is allowed, so it can only match on prefix, not on prefix-length.

Configuration

R1:

```
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
!
key chain RIP
 key 1
  key-string 0 WEAK
!
interface gigabitEthernet0/0.19
 ip rip authentication mode text
 ip rip authentication key-chain RIP
```

https://t.me/learningnets

R2:

```
router rip
 version 2
 no auto-summary
 network 136.1.0.0
 network 150.1.0.0
!
key chain RIP
 key 1
  key-string 0 STRONG
!
interface gigabitEthernet0/0
 ip rip authentication mode md5
 ip rip authentication key-chain RIP
```

ASA2:

```
route-map RIP_DEFAULT_ROUTE permit 10
 set metric 10
!
access-list RIP_INBOUND_VLAN26 deny any
!
router rip
 version 2
 no auto-summary
 default-information originate route-map RIP_DEFAULT_ROUTE
 network 136.1.0.0
 network 172.16.0.0
 passive-interface VLAN26
 distribute-list RIP_INBOUND_VLAN26 in interface VLAN26
!
interface Management0/0
 rip authentication mode text
 rip authentication key 0 WEAK key_id 1
!
interface Ethernet0/1.29
 rip authentication mode md5
 rip authentication key 0 STRONG key_id 1
```

https://t.me/learningnets

Verification

Check the routing table on all three devices for RIP-learned routes, and note the default route metric:

https://t.me/learningnets

```
R1#show ip route rip | b Gateway

Gateway of last resort is 136.1.19.16 to network 0.0.0.0
R*    0.0.0.0/0 [120/10] via 136.1.19.16, 00:00:14, GigabitEthernet0/0.19

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
R        136.1.29.0/24
           [120/1] via 136.1.19.16, 00:00:14, GigabitEthernet0/0.19
R        136.1.59.0/24
           [120/2] via 136.1.19.16, 00:00:14, GigabitEthernet0/0.19
      150.1.0.0/32 is subnetted, 4 subnets
R        150.1.2.2 [120/2] via 136.1.19.16, 00:00:14, GigabitEthernet0/0.19
R        150.1.22.22 [120/2] via 136.1.19.16, 00:00:14, GigabitEthernet0/0.19
      172.16.0.0/24 is subnetted, 1 subnets
R        172.16.10.0 [120/1] via 136.1.19.16, 00:00:14, GigabitEthernet0/0.19
!
!R2#show ip route rip | b Gateway

Gateway of last resort is 136.1.29.16 to network 0.0.0.0
R*    0.0.0.0/0 [120/10] via 136.1.29.16, 00:00:05, GigabitEthernet0/0

      136.1.0.0/16 is variably subnetted, 7 subnets, 2 masks
R        136.1.19.0/24 [120/1] via 136.1.29.16, 00:00:05, GigabitEthernet0/0
R        136.1.27.0/24 [120/2] via 136.1.29.16, 00:00:05, GigabitEthernet0/0
R        136.1.49.0/24 [120/2] via 136.1.29.16, 00:00:05, GigabitEthernet0/0
      150.1.0.0/32 is subnetted, 4 subnets
R        150.1.1.1 [120/2] via 136.1.29.16, 00:00:05, GigabitEthernet0/0
R        150.1.11.11 [120/2] via 136.1.29.16, 00:00:05, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
R        172.16.10.0 [120/1] via 136.1.29.16, 00:00:05, GigabitEthernet0/0
!
!ASA2# show route | i R

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       * - candidate default, U - per-user static route, o - ODR
R    136.1.27.0 255.255.255.0 [120/1] via 136.1.19.1, 0:00:03, VLAN19
R    136.1.59.0 255.255.255.0 [120/1] via 136.1.29.2, 0:00:15, VLAN29
R    136.1.49.0 255.255.255.0 [120/1] via 136.1.19.1, 0:00:03, VLAN19
R    150.1.2.2 255.255.255.255 [120/1] via 136.1.29.2, 0:00:15, VLAN29
R    150.1.1.1 255.255.255.255 [120/1] via 136.1.19.1, 0:00:03, VLAN19
R    150.1.11.11 255.255.255.255 [120/1] via 136.1.19.1, 0:00:03, VLAN19
```

https://t.me/learningnets

```
R    150.1.22.22 255.255.255.255 [120/1] via 136.1.29.2, 0:00:15, VLAN29
```

Verify that both routers use authentication for RIP, and verify key chain string:

```
R1#show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 22 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
GigabitEthernet0/0.19 2     2                       RIP

    GigabitEthernet0/0.27 2     2
    GigabitEthernet0/1    2     2
    Loopback0             2     2
    Loopback1             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    136.1.0.0
    150.1.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    136.1.19.16          120      00:00:12
  Distance: (default is 120)
!
!R1#show key chain
Key-chain RIP: key 1 -- text "WEAK"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
!
!R2#show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 22 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
```

https://t.me/learningnets

```
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
GigabitEthernet0/0    2     2                      RIP

    GigabitEthernet0/1    2     2
    Loopback0             2     2
    Loopback1             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    136.1.0.0
    150.1.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    136.1.29.16           120      00:00:04
  Distance: (default is 120)
!
!R2#show key chain RIP
Key-chain RIP: key 1 -- text "STRONG"

        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
```

To verify that RIP packets are authenticated, issue a **debug ip rip** on R1 and R2; you should see messages similar to the following:

```
R1:
RIP: received packet with text authentication WEAK
RIP: received v2 update from 136.1.19.16 on GigabitEthernet0/0.19
     0.0.0.0/0 via 0.0.0.0 in 10 hops
     136.1.29.0/24 via 0.0.0.0 in 1 hops
     136.1.59.0/24 via 0.0.0.0 in 2 hops
     150.1.2.2/32 via 0.0.0.0 in 2 hops
     150.1.22.22/32 via 0.0.0.0 in 2 hops
     172.16.10.0/24 via 0.0.0.0 in 1 hops
!
!R2:
RIP: received packet with MD5 authentication

RIP: received v2 update from 136.1.29.16 on GigabitEthernet0/0
     0.0.0.0/0 via 0.0.0.0 in 10 hops
     136.1.19.0/24 via 0.0.0.0 in 1 hops
     136.1.27.0/24 via 0.0.0.0 in 2 hops
     136.1.49.0/24 via 0.0.0.0 in 2 hops
     150.1.1.1/32 via 0.0.0.0 in 2 hops
```

https://t.me/learningnets

```
150.1.11.11/32 via 0.0.0.0 in 2 hops
172.16.10.0/24 via 0.0.0.0 in 1 hops
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA OSPFv2 Routing

Task

- Enable OSPF on ASA2, R1's VLAN 19, and R2's VLAN 29; ASA2 should use a router-id of 150.1.16.16.
- Make sure that your satisfy following requirements:

Interface	Area	Area Type	DR Selection	Authentication Type	Authentication String
VLAN19	0	Regular	Yes	Clear-Text	WEAK
VLAN26	0	Regular	Yes	-	-
VLAN29	2	Stub	No	MD5	STRONG

- ASA should never be selected as DR; enable area 0 authentication at process level.
- OSPF adjacency between R1 and ASA2 should be declared DOWN if no hello packets are received within 3 seconds.
- R1 should receive an OSPF default route.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config`

https://t.me/learningnets

instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

ASA supports OSPFv2 dynamic routing only for IPv4, with a maximum of two processes and the following steps (optional or mandatory):

- **Enable OSPF process (mandatory)**; configured with the command `router ospf <process_id>`, it globally activates OSPF for the given process ID.
- **Configure OSPF router-ID (optional)**; configured with the process-level command `router-id <A.B.C.D>`, it is a 32-bit number, written in the decimal format of IPv4 address, and uniquely identifies the ASA within the link-state topology. It is optional because ASA will automatically assign a router-id equal to the highest assigned IP address of an active interface.
- **Enable interfaces for OSPF (mandatory)**; configured with the process-level command `network <subnet> <mask> area <area_nr>`, it activates sending and receiving of OSPF hello multicast packets on interfaces matching the command to establish adjacencies. Area ID must match between neighbors to form the adjacency; it can be configured using the IP address of the interface or the subnet address of the interface. Also, subnets matched by the command are included in the OSPF updates/advertisements regardless of how the network command was configured.
- **Configure area types (optional)**; configured with the process-level command `area <nr> stub [no-summary]` or `area <nr> nssa [no-summary]`, it modifies the default behavior of an area, where by default all types of LSAs are allowed. With stub and nssa, ABR restricts type-4 and type-5 LSAs in the area and a default route is injected (for nssa, it is not injected by default). With stub no-summary and nssa no-summary, ABR restricts type-3, type-4 and type-5 LSAs with only a default route injected. An additional difference between stub and nssa is that nssa allows an ASBR in the area (a device performing redistribution).
- **Change Network Type (optional)**; configured with the interface-level command `ospf network point-to-point non-broadcast`, it allows the ASA to run OSPF in unicast mode and restricts the number of neighbors to a maximum of two for the respective interface. To work, it requires an additional process-level command defining the unicast neighbor `neighbor <IP Address> interface <nameif>`. By default, ASA uses the

https://t.me/learningnets

OSPF broadcast network type, where there can be multiple OSPF speakers and a DR is selected, which is in charge with the type-2 LSA generation and has full adjacency with all OSPF speakers on the segment.

- **Change Interface Priority (optional)**; configure with interface-level command `ospf priority <value>` , if a value of zero is specified, ASA is not allowed to participate in the DR selection for that interface; DR is selected based on the highest priority and in case of tie, the device with highest router-ID; by default the interface priority has a value of one.
- **Change OSPF Timers (optional)**; configured with interface-level commands `ospf hello-interval <seconds>` (interval for sending hello packets) and `ospf dead-interval <seconds>` (maximum interval for keeping a neighbor active if no hello packets are received). For OSPF adjacency to be formed, OSPF timers must have the same values between neighbors; the default values are 10 (hello) and 40 (dead).
- **Originate default route (optional)**; configured using the command `default-information originate [always] [route-map <name>]` , it originates an E2 type default route only if it has a default route present in the routing table. If the `always` keyword is specified, it will unconditionally originate the default route, whereas if a route-map is applied, it will originate the default route only if conditions from route-map are satisfied (for example, if a specific route is present in the routing table).

The most important aspect of OSPF configuration from a security standpoint is protocol authentication. OSPF authenticates all OSPF packets (authentication is a part of the OSPF header, and OSPF has the IP protocol number of 89). The OSPF header supports three types of authentication: null (empty), plain-text (clear text password), and secure MD5 hash. You may define various authentication types on different interfaces. First, look at the authentication types:

- **NULL** explicitly states that the OSPF is not authenticated.
- **Plain-text** carries a clear-text password in the OSPF header.
- **MD5-hash** carries a key ID along with the corresponding hash value in the header; multiple keys can be configured on each interface, in which case each OSPF packet is sent several times (authenticated with each key and corresponding hash). This is useful when the authentication string must be changed so the OSPF adjacency does not flap, or when over the same interface different OSPF neighbors have different authentication strings configured.

OSPF authentication can be enabled at the interface level using the command `ospf authentication [null|message-digest]` , or at the process level using the command `area <nr> authentication [message-digest]` . The authentication string and key are configured only at the interface level using the command `ospf authentication-key 0 <string>`

https://t.me/learningnets

Or `ospf message-digest-key <nr> md5 0 <string> .`

OSPF is a highly complex link-state routing protocol, and not all options were included above. ASA supports only a subset of the OSPF features found in IOS code; one important missing feature is the passive-interface. Although the CCIE Security lab should not be focused on deep routing configurations, it is good to know your options for configuring OSPF on ASA.

Configuration

R1:

```
router ospf 10
 network 136.1.19.0 0.0.0.255 area 0
!
interface GigabitEthernet0/0.19
 ip ospf authentication
 ip ospf authentication-key 0 WEAK
 ip ospf hello-interval 1
 ip ospf dead-interval 3
```

R2:

```
router ospf 20
 network 136.1.29.2 0.0.0.0 area 2
 area 2 stub
!
interface GigabitEthernet0/0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 0 STRONG
 ip ospf network point-to-point
```

ASA2:

```
router ospf 1
 router-id 150.1.16.16
 network 136.1.19.16 255.255.255.255 area 0
 network 172.16.10.1 255.255.255.255 area 0
 network 136.1.29.0 255.255.255.0 area 2
 neighbor 136.1.29.2 interface VLAN29
 default-information originate always
 area 0 authentication
 area 2 stub
```

https://t.me/learningnets

```
!
interface Management0/0
 ospf priority 0
 ospf authentication-key 0 WEAK
 ospf hello-interval 1
 ospf dead-interval 3
!
interface Ethernet0/1.26
 ospf priority 0
 ospf authentication null
!
interface Ethernet0/1.29
 ospf network point-to-point non-broadcast
 ospf authentication message-digest
 ospf message-digest-key 1 md5 0 STRONG
```

Verification

Verify OSPF neighbors and configuration:

```
ASA2# show ospf neighbor


Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.22.22       1   FULL/  -        0:00:31     136.1.29.2      VLAN29
150.1.11.11       1   FULL/DR         0:00:02     136.1.19.1      VLAN19
!
!ASA2# show ospf 1
 Routing Process "ospf 1" with ID 150.1.16.16 and Domain ID 0.0.0.1
 Supports only single TOS(TOS0) routes
 Does not support opaque LSA
 It is an area border and autonomous system boundary router
 Redistributing External Routes from,
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x     0
 Number of opaque AS LSA 0. Checksum Sum 0x     0
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 1 normal 1 stub 0 nssa
 External flood list length 0 Area BACKBONE(0)
        Number of interfaces in this area is 2 Area has simple password authentication
        SPF algorithm executed 8 times
        Area ranges are
        Number of LSA 9. Checksum Sum 0x 5580c
```

https://t.me/learningnets

```
        Number of opaque link LSA 0. Checksum Sum 0x     0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0 Area 2
        Number of interfaces in this area is 1 It is a stub area

          generates stub default route with cost 1
        Area has no authentication
        SPF algorithm executed 8 times
        Area ranges are
        Number of LSA 9. Checksum Sum 0x 4cd8a
        Number of opaque link LSA 0. Checksum Sum 0x     0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Verify ASA interface authentication, interface priority, and timers:

```
ASA2# show ospf interface VLAN29

VLAN29 is up, line protocol is up
  Internet Address 136.1.29.16 mask 255.255.255.0, Area 2
  Process ID 1, Router ID 150.1.16.16, Network Type POINT_TO_POINT, Cost: 10
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 0:00:06
  Index 1/3, flood queue length 0
  Next 0x00000000(0)/0x00000000(0)
  Last flood scan length is 4, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.22.22
  Suppress hello for 0 neighbor(s) Message digest authentication enabled
 Youngest key id is 1
!
!ASA2# show ospf interface VLAN19

VLAN19 is up, line protocol is up
  Internet Address 136.1.19.16 mask 255.255.255.0, Area 0
  Process ID 1, Router ID 150.1.16.16, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  Designated Router (ID) 150.1.11.11, Interface address 136.1.19.1
  No backup designated router on this network
```

https://t.me/learningnets

```
Timer intervals configured, Hello 1, Dead 3, Wait 3, Retransmit 5

    Hello due in 0:00:00
  Index 1/1, flood queue length 0
  Next 0x00000000(0)/0x00000000(0)
  Last flood scan length is 2, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 150.1.11.11  (Designated Router)
  Suppress hello for 0 neighbor(s) Simple password authentication enabled
```

Verify that the ABR (ASA) injects a defaut route in the stub area:

```
R2#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 136.1.29.16 to network 0.0.0.0
 O*IA  0.0.0.0/0 [110/2] via 136.1.29.16, 05:16:59, GigabitEthernet0/0

      136.1.0.0/16 is variably subnetted, 6 subnets, 2 masks
O IA     136.1.19.0/24 [110/11] via 136.1.29.16, 05:16:54, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O IA     172.16.10.0 [110/11] via 136.1.29.16, 05:16:54, GigabitEthernet0/0
```

Verify that ASA originates a default route into area 0:

```
R1#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is 136.1.19.16 to network 0.0.0.0
```

https://t.me/learningnets

```
O*E2  0.0.0.0/0 [110/1] via 136.1.19.16, 00:00:04, GigabitEthernet0/0.19

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
O IA     136.1.29.0/24
           [110/11] via 136.1.19.16, 05:18:35, GigabitEthernet0/0.19
      172.16.0.0/24 is subnetted, 1 subnets
O        172.16.10.0 [110/11] via 136.1.19.16, 05:18:45, GigabitEthernet0/0.19
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA EIGRP Routing

Task

- Enable EIGRP on ASA2, R1's VLAN 19, and R2's VLAN 29.
- ASA2 should advertise a 172.0.0.0/8 prefix to R1.
- Ensure that ASA cannot form any EIGRP adjacencies over VLAN 26.
- ASA2 and R2 should use authenticated unicast for the EIGRP adjacency.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

ASA supports EIGRP dynamic routing only for IPv4, and a single process is allowed. As with other routing protocols, ASA implements a subset of the EIGRP functionality found in the IOS code. EIGRP is a sophisticated distributed (diffused) computations-based and scalable protocol. However, configuration is relatively simple and requires just a few steps:

- **Enable EIGRP process (mandatory)**; configured with the command `router eigrp <ASN>` , it globally activates EIGRP within the given autonomous system number and

https://t.me/learningnets

needs to match between EIGRP speakers to become neighbors.

- **Disable auto-summary (mandatory)**; configured with process-level command `no auto-summary` , it disables the classful network boundary auto-summarization function of EIGRP, which by default is enabled.
- **Enable interfaces for EIGRP (mandatory)**; configured with the process-level command `network <subnet> <mask>` , it activates sending and receiving of EIGRP hello multicast packets on interfaces matching the command to establish adjacencies. It can be configured using the IP address of the interface or the subnet address of the interface. Also, subnets matched by the command are included in the EIGRP updates/advertisements regardless of how the network command was configured.
- **Define passive interfaces (optional)**; configured with the process-level command `passive-interface <nameif>` , it disables sending and receiving of EIGRP hello packets, so no neighbors can be formed.
- **Change EIGRP Timers (optional)**; configured with the interface-level commands `hello-interval eigrp <ASN> <seconds>` (interval for sending hello packets) and `hold-interval <ASN> <seconds>` (maximum interval for keeping a neighbor active if no hello packets are received). The default values are 5 (hello) and 15 (hold).
- **Configure Unicast Neighbors (optional)**; configured with the process-level command `neighbor <IP Address> interface <nameif>` , it disables multicast EIGRP for the respective interface and enables unicast EIGRP for the configured neighbors. Both sides must be configured for unicast.
- **Configure manual summary (optional)**; configured with the interface-level command `summary-address eigrp <ASN> <network> <mask>` , it filters any subnets contained within the summary from the outgoing updates and sends only the configured major class.

From a security perspective, authentication is probably the most important aspect. EIGRP supports only MD5 authentication, and a single key can be configured per interface because no key chain is used as in the IOS code.

Authentication is enabled at the interface level using the command `authentication mode eigrp <ASN> md5` . Key/string is configured with the interface-level command `authentication key eigrp <ASN> <string> key-id <nr>` , and both key number and string must match between neighbors.

Configuration

```
R1:
router eigrp 1
 network 136.1.19.1 0.0.0.0
```

https://t.me/learningnets

R2:

```
router eigrp 1
 network 136.1.29.2 0.0.0.0
 neighbor 136.1.29.16 gigabitEthernet0/0
!
key chain EIGRP-CHAIN
 key 1
  key-string EIGRP-PASS
!
interface GigabitEthernet0/0
 ip authentication mode eigrp 1 md5
 ip authentication key eigrp 1 EIGRP-CHAIN
```

ASA2:

```
router eigrp 1
 no auto-summary
 network 136.1.19.0 255.255.255.0
 network 136.1.29.16 255.255.255.0
 network 172.16.10.1 255.255.255.255
 neighbor 136.1.29.2 interface VLAN29
 passive-interface default
 no passive-interface VLAN19
 no passive-interface VLAN29
!
interface Ethernet0/1.29
 authentication key eigrp 1 EIGRP-PASS key-id 1
 authentication mode eigrp 1 md5
!
interface Management0/0
 summary-address eigrp 1 172.0.0.0 255.0.0.0
```

Verification

Verify EIGRP neighbors:

https://t.me/learningnets

```
ASA2# show eigrp neighbors

EIGRP-IPv4 neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   136.1.29.2              Et0/1.29          12  00:03:05 1    200   0   6
0   136.1.19.1              Ma0/0             13  00:05:56 1    200   0   12
```

Verify that authentication is used between ASA2 and R2:

```
R2#show ip eigrp interfaces detail gigabitEthernet0/0
EIGRP-IPv4 Interfaces for AS(1)
                              Xmit Queue   PeerQ        Mean   Pacing Time   Multicast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Gi0/0                    1        0/0       0/0           1       0/0           50           0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 5/1
  Hello's sent/expedited: 108/2
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 7/6
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0 Authentication mode is md5,  key-chain is "EIGRP-CHAIN"
!
!R2#show key chain EIGRP-CHAIN
Key-chain EIGRP-CHAIN: key 1 -- text "EIGRP-PASS"

        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
```

Verify that unicast EIGRP is used between ASA2 and R2:

https://t.me/learningnets

```
R2#show ip eigrp neighbors detail

EIGRP-IPv4 Neighbors for AS(1)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   136.1.29.16             Gi0/0                    10 00:08:00    1   100  0  22 Static neighbor

   Version 8.4/3.0, Retrans: 0, Retries: 0, Prefixes: 2
   Topology-ids from peer - 0
```

Verify that VLAN 26 of ASA2 is passive, so it is not listed as an EIGRP interface:

```
ASA2# show eigrp interfaces

EIGRP-IPv4 interfaces for process 1

                      Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
VLAN19             1        0/0         1       0/1           50           0
VLAN29             1        0/0         1       0/1           50           0
```

Verify that ASA creates the manual summary for VLAN 26 and advertises it out only to R1 (R2 receives the /24 subnet):

```
ASA2# show eigrp topology

EIGRP-IPv4 Topology Table for AS(1)/ID(172.16.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 136.1.29.0 255.255.255.0, 1 successors, FD is 2816
        via Connected, Ethernet0/1.29
P 136.1.19.0 255.255.255.0, 1 successors, FD is 28160
        via Connected, Management0/0P 172.0.0.0 255.0.0.0, 1 successors, FD is 2816
 via Summary (2816/0), Null0
P 172.16.10.0 255.255.255.0, 1 successors, FD is 2816
        via Connected, Ethernet0/1.26
!
!ASA2# show route | i D
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       * - candidate default, U - per-user static route, o - ODR
D     172.0.0.0 255.0.0.0 is a summary, 0:09:07, Null0
```

https://t.me/learningnets

```
!
!R1#show ip route eigrp | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D        136.1.29.0/24
           [90/26112] via 136.1.19.16, 00:12:30, GigabitEthernet0/0.19
D     172.0.0.0/8 [90/26112] via 136.1.19.16, 00:07:11, GigabitEthernet0/0.19
!
!R2#show ip route eigrp | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 6 subnets, 2 masks
D        136.1.19.0/24
           [90/30720] via 136.1.29.16, 00:10:07, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
D        172.16.10.0 [90/28416] via 136.1.29.16, 00:07:39, GigabitEthernet0/0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Route Redistribution

Task

- Enable EIGRP between ASA2 and R1, and advertise Loopback0 subnet.
- Enable OSPF in area 0 between ASA2 and R2, and advertise Loopback0 subnet.
- Without using any `network` commands, ensure that R1 and R2 have EIGRP/OSPF routes for ASA2's VLAN 26 subnet.
- Achieve the following output:

```
R1#ping 150.1.2.2 source loopback0
    Type escape sequence to abort.
    Sending 5, 100-byte ICMP Echos to 150.1.2.2, timeout is 2 seconds:
    Packet sent with a source address of 150.1.1.1 !!!!!

    Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

https://t.me/learningnets

Overview

The CCIE Security lab most likely will not require you to perform advanced routing protocols tuning, but basic routing features should be known by every candidate. This task requires you to redistribute between the routing protocols, namely OSPF, EIGRP, and connected. This is needed to obtain full reachability between the routing domains connected by the firewall. Redistribution can be performed for the following route prefixes: connected, static, RIP, EIGRP, OSPF. There can be one-way redistribution (from RIP to EIGRP, for example) or two-way redistribution (from OSPF to EIGRP and from EIGRP to OSPF). When redistribution is performed from protocol A to protocol B:

- All routing table entries learned through protocol A are advertised into protocol B.
- All locally connected routes advertised into protocol A through `network` commands are advertised into protocol B; however, if connected is also redistributed into protocol B, only prefixes allowed by the connected redistribution will be performed because it's more specific and takes precedence.

The following must be considered for redistribution on the ASA:

- When redistributing into OSPF, if the `subnets` keyword is not specified, no subnets will be redistributed, just major network classes.
- When redistributing into OSPF, by default metric-type is E2 (patch cost is never added to the metric), but it can be changed to E1 (patch cost is always added to the metric), default metric value is 20 in both cases, and it can be changed as well.
- When redistributing into EIGRP or RIP, a metric must be specified; otherwise, prefixes will not be advertised. The exception is for static and connected, where the metric is automatically determined by the ASA.

Configuration

R1:

```
router eigrp 1
 no auto-summary
 network 136.1.19.1 0.0.0.0
 network 150.1.1.1 0.0.0.0
```

R2:

```
router ospf 1
 network 136.1.29.2 0.0.0.0 area 0
```

https://t.me/learningnets

```
 network 150.1.2.2 0.0.0.0 area 0
```

ASA2:

```
prefix-list CONNECTED permit 172.16.10.0/24
prefix-list CONNECTED permit 136.1.19.0/24
prefix-list CONNECTED permit 136.1.29.0/24
!
route-map CONNECTED permit 10
 match ip address prefix-list CONNECTED
!
router eigrp 1
 no auto-summary
 network 136.1.19.0 255.255.255.0
 redistribute connected route-map CONNECTED
 redistribute ospf 1
!
router ospf 1
 network 136.1.29.0 255.255.255.0 area 0
 redistribute connected subnets route-map CONNECTED
 redistribute eigrp 1 subnets
```

Verification

Verify EIGRP and OSPF neighbors:

```
ASA2# show eigrp neighbors
EIGRP-IPv4 neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   136.1.19.1              Ma0/0            10  00:33:55 1    200   0   15
!
!ASA2# show ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.22.22       1   FULL/DR         0:00:34     136.1.29.2      VLAN29
```

Verify that R1 and R2 have a route for ASA2's VLAN 26 prefix:

```
R1#show ip route 172.16.10.0 255.255.255.0
Routing entry for 172.16.10.0/24 Known via "eigrp 1", distance 170, metric 26112, type external
```

https://t.me/learningnets

```
  Redistributing via eigrp 1
  Last update from 136.1.19.16 on GigabitEthernet0/0.19, 00:35:08 ago
  Routing Descriptor Blocks: * 136.1.19.16, from 136.1.19.16, 00:35:08 ago, via GigabitEthernet0/0.19
      Route metric is 26112, traffic share count is 1
      Total delay is 20 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
!
!R2#show ip route 172.16.10.0 255.255.255.0
Routing entry for 172.16.10.0/24
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 136.1.29.16 on GigabitEthernet0/0, 00:35:20 ago
  Routing Descriptor Blocks: * 136.1.29.16, from 172.16.10.1, 00:35:20 ago, via GigabitEthernet0/0

      Route metric is 20, traffic share count is 1
```

Verify that the requested output is matached (there is IP connectivity between R1 and R2 Loopback0 prefixes, so two-way redistribution is functional):

```
R1#ping 150.1.2.2 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.2.2, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Verify that connected redistribution is correctly configured, so VLAN 29 and VLAB 19 are learned by R1 and R2:

```
R1#show ip route eigrp | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks D EX     136.1.29.0/24
 [170/26112] via 136.1.19.16, 00:40:12, GigabitEthernet0/0.19
      150.1.0.0/32 is subnetted, 4 subnets
D EX     150.1.2.2
           [170/1732352] via 136.1.19.16, 00:40:51, GigabitEthernet0/0.19
      172.16.0.0/24 is subnetted, 1 subnets
D EX     172.16.10.0
           [170/26112] via 136.1.19.16, 00:40:12, GigabitEthernet0/0.19
!
!R2#show ip route ospf | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 6 subnets, 2 masks
O E2     136.1.19.0/24 [110/20] via 136.1.29.16, 00:41:17, GigabitEthernet0/0
```

https://t.me/learningnets

```
      150.1.0.0/32 is subnetted, 4 subnets
O E2     150.1.1.1 [110/20] via 136.1.29.16, 00:03:53, GigabitEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O E2     172.16.10.0 [110/20] via 136.1.29.16, 00:40:35, GigabitEthernet0/0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Stub Multicast Routing

Task

- Configure ASA2 to act as a proxy agent for IGMP join/leave messages on its VLAN 29 interface.
- The RPF interface of the ASA for unknown sources should be VLAN 19.
- Match the following output:

```
R1#ping 239.0.0.5 repeat 2
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 239.0.0.5, timeout is 2 seconds:
Reply to request 0 from 136.1.29.2
, 1 msReply to request 1 from 136.1.29.2
, 1 ms
```

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

https://t.me/learningnets

Multicast routing is a fairly complex topic that cannot be fully covered within the scope of the CCIE Security exam, but you must be aware of some basic functionality as it relates to the ASA firewall. Unlike unicast IP traffic (which is one source to one destination), multicast flows have many recipients/subscribers (one source to multiple destinations). This requires a completely different packet delivery method, as opposed to hop-by-hop destination-based forwarding. Multicast traffic is distinguished by the destination IP address in the range 224.0.0.0-239.255.255.255 (every such address is called a multicast group). When a network is configured for multicast routing, it accepts such packets and floods them toward all subscribed members of the particular group. This flooding procedure is the core of the multicast traffic delivery.

When a new member wants to join a multicast group, it uses a special protocol called IGMP (Internet Group Membership Protocol) to signal to its default router that it wants to join the group. The router, which is a part of the multicast-capable network, will further interact with other routers through the use of PIM (Protocol Independent Multicast) to let them know that a new member needs the multicast feed. As a multicast packet is forwarded through the network, every router performs the so-called RPF (Reverse Path Forwarding) check on the packet, to prevent routing loops. The RPF procedure takes the source IP address of the packet and verifies that the shortest path toward the source IP address is across the interface that packet was received on (based on the current state of the unicast routing table); otherwise, the packet will be dropped. This ensures that the packet is not looped but correctly received from the respective upstream node. If the RPF check is successful, the firewall inspects the list of the downstream interfaces that received subscription signals for the packet's multicast group. If there are any, the packet is replicated and forwarded downstream. This procedure is repeated until the packet is delivered to all terminal (leaf) nodes in the network, which represent the multicast subscribers.

In this task, we configure the ASA firewall as a stub multicast router, where stub means that the firewall does not participate as a layer 3 hop in the multicast tree (it cannot interpret any IGMP messages received from attached hosts and does not participate in the PIM tree), but only acts as a blind relay for IGMP messages and multicast packets; however, it still performs RPF checks.

In our scenario, R2 is the subscriber, sending IGMP messages to the ASA. The firewall in turn forwards the messages to R1, which is the real multicast router. This procedure makes R1 think that the ASA needs the multicast feeds (that are actually subscribed by R2) and R1 will flood them down to the firewall. The command to enable IGMP message forwarding is `igmp forward interface <nameif>` and is applied to the interface connected to the stub multicast segment (R2). Note that the stub multicast routing requires the `multicast-routing` command to be entered in the global configuration mode first, to enable multicast packet forwarding and IGMP;

https://t.me/learningnets

when this command is used, all ASA interfaces will have IGMP enabled. For a layer3 device to accept IGMP join messages, it must have PIM enabled on the interface (sparse, dense, or sparse-dense mode); in our case it's R1.

When the firewall receives the multicast packets, it first determines whether they are permitted by security-levels or any applied ACLs. If the check was successful, the firewall will then do the RPF check described earlier, to make sure the packet is not looped. The ASA would use its routing table to do the source IP address lookup. In some cases, if you need to tweak the RPF check result (you receive the multicast traffic inbound on the correct interface but the unicast routing table does not point to the same interface for the multicast source), you can configure static multicast routes using the command `mroute <src_ip> <src_mask> [<nameif>|<rpf_neighbor_ip>]` ; when the RPF check is done, first the ASA consults the multicast routing table, and if no entries are found for the source, the unicast routing table is used—so it basically overrides the unicast routing table entries.

ICMP Echo Request packets are sent as multicast, whereas ICMP Echo Reply packets will be unicast. With ICMP inspection activated, ASA considers ICMP a stateful protocol, and because source/destination IP address from the Echo Request does not match the Echo Reply traffic, it is interpreted as a security violation, and Echo Reply traffic is dropped by the ASP engine; this can be viewed with the command `show asp drop` . For ping to work, ICMP inspection must be disabled.

Configuration

```
R1:
interface GigabitEthernet0/0.19
 ip pim sparse-dense-mode



R2:
interface GigabitEthernet0/0
 ip igmp join 239.0.0.5



ASA2:
multicast-routing
mroute 0.0.0.0 0.0.0.0 VLAN19
!
access-list PERMIT permit ip any any
access-group PERMIT in interface VLAN29
!
```

https://t.me/learningnets

```
interface Ethernet0/1.29
 igmp forward interface VLAN19
!
policy-map global_policy
 class inspection_default
  no inspect icmp
```

Verification

Verify that ASA2 has joined the multicast group, proxying for R2:

```
R1#show ip igmp membership
Flags: A  - aggregate, T - tracked
       L  - Local, S - static, V - virtual, R - Reported through v3
       I - v3lite, U - Urd, M - SSM (S,G) channel
       1,2,3 - The version of IGMP, the group is in
Channel/Group-Flags:
       / - Filtering entry (Exclude mode (S,G), Include mode (G))
Reporter:
       <mac-or-ip-address> - last reporter if group is not explicitly tracked
       <n>/<m>      - <n> reporter in include mode, <m> reporter in exclude

 Channel/Group                  Reporter        Uptime   Exp.  Flags  Interface
*,239.0.0.5                     136.1.19.16     00:00:24 02:35 2A     Gi0/0.19

 *,224.0.1.40                   136.1.19.1      13:54:49 02:07 2LA    Gi0/0.19
```

Confirm that R2 joined the multicast group, on ASA2:

```
ASA2# show igmp interface VLAN29
VLAN29 is up, line protocol is up
  Internet address is 136.1.29.16/24 IGMP is enabled on interface
  Current IGMP version is 2
  IGMP query interval is 125 seconds
  IGMP querier timeout is 255 seconds
  IGMP max query response time is 10 seconds
  Last member query response interval is 1 seconds
  Inbound IGMP access group is:
  IGMP limit is 500, currently active joins: 1
  Cumulative IGMP activity: 1 joins, 0 leaves IGMP forwarding on interface VLAN19
  IGMP querying router is 136.1.29.16 (this system)
!
!ASA2# show igmp groups
```

https://t.me/learningnets

```
IGMP Connected Group Membership
Group Address    Interface            Uptime    Expires   Last Reporter
239.0.0.5        VLAN29               14:03:36  00:02:12  136.1.29.2
```

Check the multicast routing table before generating traffic; there are no active flows:

```
ASA2# show mroute

Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State
(*, 239.0.0.5), 00:31:51/never, RP 0.0.0.0, flags: DC
  Incoming interface: Null
  RPF nbr: 0.0.0.0 Immediate Outgoing interface list:
VLAN29, Forward, 00:31:51/never
```

Generate traffic from R1; it should be functional. Verify the multicast routing table again on ASA2:

```
R1#ping 239.0.0.5 repeat 2
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 239.0.0.5, timeout is 2 seconds:
Reply to request 0 from 136.1.29.2
, 4 msReply to request 1 from 136.1.29.2
, 1 ms
!
!ASA2# show mroute 136.1.19.1 239.0.0.5

Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State
(136.1.19.1, 239.0.0.5), 00:02:32/00:02:57, flags: DJT
  Incoming interface: VLAN19
  RPF nbr: 136.1.19.1 Inherited Outgoing interface list:
```

https://t.me/learningnets

```
VLAN29, Forward, 00:37:21/never
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA PIM Multicast Routing

Task

- Configure ASA2 for multicast routing, acting as the RP for only the group 239.0.0.10.
- Limit the number of IGMP states on ASA2's VLAN 29 to 15.
- Configure ASA2 to allow only R1 as its PIM neighbor on VLAN 19 interface; ensure that R1 is the DR.
- On R2, join the VLAN 29 interface to group 239.0.0.10 and make sure R1 can ping it.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

In addition to acting as a stub multicast router, the ASA firewall is capable of becoming a regular layer 3 multicast-capable router, but ASA cannot function in both modes at the same time. In this mode, the ASA will establish PIM (Protocol Independent Multicast) adjacencies with neighboring multicast routers. This would allow the firewall to signal building of multicast distribution trees with other routers. The ASA supports only PIM Sparse Mode, which is a scalable multicast routing protocol. The key feature of PIM SM is the use of a special router called

https://t.me/learningnets

Rendezvous Point (RP), which functions as the meeting point of multicast sources and receivers. The sources register with the RP, and the subscribers build initial distribution trees to the same RP. This allows them to meet, and this is why RP is so critical to the PIM SM network.

When configured for PIM Multicast routing mode, the ASA firewall accepts and processes IGMP messages by itself. There is no need to relay them to any helper router. When the firewall receives the proper IGMP message, it initiates building of multicast subscription tree toward the RP, just as a normal multicast router. However, there is a difference; Cisco IOS routers support automatic learning of the RP information via protocols such as BSR (bootstrap router) or auto-RP. However, the ASA only supports static manual RP configuration using the command `pim rp-address <IP Address> list <ACL>`. When configuring your firewall for PIM multicast routing, do not forget to enter this command, or the multicast routing will not work.

The following steps are required to enable PIM SM Multicast routing in the ASA firewall:

1. Enable multicast-routing command globally with the command `multicast-routing`; this enables IGMP on all interfaces and allows to further enable PIM.
2. Configure the interface connected to the multicast network for PIM, using the `pim` command.
3. Configure the subscriber-facing interface for IGMP, possibly tuning IGMP settings using the `igmp` command.
4. Configure the RP for the network, using the command `pim rp-address <IP Address> list <ACL>`; ACL is optional, only standard is supported, and it can restrict the groups that configured RP is responsible for.

If ASA is the RP, it can be configured to filter which sources are allowed to register with the RP and for which multicast groups. RP registration is done by the PIM DR router of the segment, but all PIM Register messages contain the initial multicast packet, which includes the IP address of the multicast source and the multicast group address. If the RP rejects the registration, it sends back to the DR a PIM Register-Stop message and does not build the SPT toward the source. This is configured with the global command `pim accept-register [list <name>| route-map <name>]`. Only extended ACL is supported, source matching on the source of the multicast feed and destination matching on the multicast group. This filtering does not work if the RP is also the DR performing the PIM registration.

ASA can limit the number of multicast groups allowed to be joined on any given interface using the interface-level command `igmp limit <nr>`. It can also limit which groups can be joined on any given interface using interface-level command `igmp access-group <acl>`

https://t.me/learningnets

. If standard ACL is used, multicast groups allowed to be joined must match a permit entry. If extended ACL is used, use any as the source and configure the multicast group as the destination. Usually extended ACL is used for IGMP version 3, in order to match on both the source of the multicast stream the multicast group address, but IGMP version 3 is not supported by the ASA yet.

Configuration

ASA2:

```
multicast-routing
!
access-list RP-GROUP standard permit host 239.0.0.10
pim rp-address 136.1.19.16 RP-GROUP
!
access-list PIM-NEIGHBOR-FILTER standard permit host 136.1.19.1
!
interface Management0/0
 pim
 pim dr-priority 0
 pim neighbor-filter PIM-NEIGHBOR-FILTER
!
access-list IGMP-GROUP standard permit host 239.0.0.10
!
interface Ethernet0/1.29
 pim
 igmp limit 15
 igmp access-group IGMP-GROUP
!
access-list PERMIT permit ip any any
access-group PERMIT in interface VLAN29
!
policy-map global_policy
 class inspection_default
  no inspect icmp
```

R1:

```
interface GigabitEthernet0/0.19
 ip pim sparse-mode
!
ip pim rp-address 136.1.19.16
```

R2:

https://t.me/learningnets

```
interface GigabitEthernet0/0
 ip igmp join 239.0.0.10
```

Verification

Verify PIM neighbors; R1 is the DR:

```
ASA2# show pim neighbor

Neighbor Address  Interface          Uptime    Expires DR pri Bidir
136.1.19.1        VLAN19             14:51:37  00:01:34 1 (DR)
!
!ASA2# show pim interface VLAN19

Address          Interface          PIM  Nbr    Hello  DR         DR
                                         Count  Intvl  Prior
136.1.19.16      VLAN19             on   1      30     0          136.1.19.1
```

Verify IGMP settings on the VLAN 29 interface:

```
ASA2# show igmp interface VLAN29
VLAN29 is up, line protocol is up
  Internet address is 136.1.29.16/24 IGMP is enabled on interface
  Current IGMP version is 2
  IGMP query interval is 125 seconds
  IGMP querier timeout is 255 seconds
  IGMP max query response time is 10 seconds
  Last member query response interval is 1 seconds Inbound IGMP access group is: IGMP-GROUP
IGMP limit is 15, currently active joins: 1
  Cumulative IGMP activity: 2 joins, 1 leaves
  IGMP querying router is 136.1.29.16 (this system)
 !
 !ASA2# show igmp groups
IGMP Connected Group Membership
Group Address    Interface            Uptime    Expires   Last Reporter
239.0.0.10       VLAN29               16:08:23  00:03:50  136.1.29.2

224.0.1.40       VLAN19               16:29:31  00:03:26  136.1.19.1
```

Verify the RP configuration:

https://t.me/learningnets

```
ASA2# show pim group-map


Group Range         Proto Client Groups RP address      Info

224.0.1.39/32*      DM    static 0      0.0.0.0
224.0.1.40/32*      DM    static 1      0.0.0.0
239.0.0.10/32*      SM    config 1      136.1.19.16     RPF: Tu1,136.1.19.16 (us)

224.0.0.0/24*       L-Localstatic 1      0.0.0.0
232.0.0.0/8*        SSM   config 0      0.0.0.0
224.0.0.0/4*        SM    static 0      0.0.0.0         RPF: ,0.0.0.0
```

Verify that multicast traffic is functional, and look at the flags check that SPT was joined:

```
R1#ping 239.0.0.10 repeat 2
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 239.0.0.10, timeout is 2 seconds:
Reply to request 0 from 136.1.29.2, 4 ms
Reply to request 1 from 136.1.29.2, 1 ms
!
!ASA2# show mroute

Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(*, 224.0.1.40), 15:44:28/never, RP 0.0.0.0, flags: DPC
  Incoming interface: Null
  RPF nbr: 0.0.0.0
  Immediate Outgoing interface list:
    VLAN19, Null, 15:44:28/never

(*, 239.0.0.10), 15:44:28/never, RP 136.1.19.16, flags: SCJ
  Incoming interface: Tunnel1
  RPF nbr: 136.1.19.16
  Immediate Outgoing interface list:
    VLAN29, Forward, 15:44:28/never
(136.1.19.1, 239.0.0.10), 00:00:41/00:03:18, flags: SJT
```

https://t.me/learningnets

```
  Incoming interface: VLAN19 RPF nbr: 136.1.19.1
  Inherited Outgoing interface list: VLAN29, Forward, 15:44:28/never
```

The DR established a tunnel with the RP for the PIM Register messages:

```
R1#show interfaces tunnel0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel Description: Pim Register Tunnel (Encap) for RP 136.1.19.16
  Interface is unnumbered. Using address of GigabitEthernet0/0.19 (136.1.19.1)
  MTU 17912 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set Tunnel source 136.1.19.1 (GigabitEthernet0/0.19), destination 136.1.19.16
   Tunnel Subblocks:
      src-track:
         Tunnel0 source tracking subblock associated with GigabitEthernet0/0.19
          Set of tunnels with source GigabitEthernet0/0.19, 1 member (includes iterators), on interface <OK>
Tunnel protocol/transport PIM/IPv4
  Tunnel TOS/Traffic Class 0xC0,  Tunnel TTL 255
  Tunnel transport MTU 1472 bytes Tunnel is transmit only

  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Basic Access-Lists

Task

- Configure ASA2 to allow the following traffic inbound on its VLAN 29 interface; use a global ACL:
  - PING from R2's Loopback1 address.
  - SYSLOG/HTTP/RDP towards **TEST-SRV-A**.
  - ensure that initial firewall policies, based on security-levels, are still valid.
- Configure ASA1 to allow following traffic inbound on its VLAN 49 interface; use an interface ACL:
  - TELNET/SSH for R2's Loopback0 address.
  - FTP/TFTP/HTTPS from R1's Loopback0 address.
  - Return traffic for UNIX style traceroute.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Access-lists are commonly used to implement traffic filtering in the ASA firewalls. By default, ASA permits sessions to be initiated from a higher security level interface to

https://t.me/learningnets

a lower security level interface. When traffic is allowed one way, a session state is created in the state/conn table and return traffic will be automatically allowed back in; thus, access-lists or security-level does not apply to the return traffic:

- This behavior only works for protocols that do not negotiate secondary connections (such as FTP and TFTP).
- For protocols negotiating secondary channels, it works only if the respective protocol is inspected by MPF engine, but this is discussed in another lab.

At a minimum, using access-lists allows you to do the following:

- Permit access from the lower security level interfaces to higher security level interfaces.
- Restrict access between interfaces with same security-level if this is allowed (by default, it is not and requires the global command `same-security-traffic permit inter-interface`).
- Permit return traffic for sessions that are not inspected by the ASA firewall (such as for ICMP, which is not inspected by default, or for the traceroute command).
- Permit return traffic for multicast traffic, which cannot be inspected by the firewall because traffic is multicast one-way and unicast for return.
- Filter routing updates for OSPF and RIP routing processes.

ASA supports both IPv4 and IPv6 access-lists, both standard and extended. Just like on a router, you can apply a maximum of one access-list per protocol, per interface, per direction, so the maximum allowed is:

- One IPv4 and one IPv6 ACL applied inbound on an interface; configured with the command `access-group <ACL> in interface <nameif> .`
- One IPv4 and one IPv6 ACL applied outbound on an interface; configured with the command `access-group <ACL> out interface <nameif> .`

Starting with code 8.3(1), you can also apply one IPv4 and one IPv6 ACL globally, configured with the command `access-group <ACL> global` ; the global access-list is always interpreted as an inbound ACL. When the global ACL is applied, the implicit `deny ip any any` rule is removed from any ACL applied inbound on any interfaces, so that the global ACL is inspected. Assuming all ACLs are configured, and assuming a packet enters the **inside** interface and exits the **outside** interface, the rules are:

- Traffic is matched against the inbound ACL applied on the **inside** interface. If there is a match, the action is taken; if not, it will match on the default **deny ip any any** unless global ACL is configured.

https://t.me/learningnets

- Traffic is matched against the global ACL. If there is a match, the action is taken; if not, it will match on the default **deny ip any any** rule.
- Traffic is matched against the outbound ACL applied on the **outside** interface. If there is a match, the action is taken; if not, it will match on the default **deny ip any any** rule.

When an ACL is applied inbound on an interface, all inbound traffic is subject to the ACL rules, and the default security-level rule no longer applies; when an ACL is applied globally, all inbound traffic from all interfaces is subject to the global ACL, and the security-level rules no longer apply. We used a global ACL on ASA2, so after configuring the required rules, we need to further allow traffic inbound on VLAN 19 and VLAN 26 because it was initially based on the security-levels (only Loopback1 interface subnets of R1/R2 are routed through ASA2; Loopback0 interfaces of R1/R2 are routed through ASA1).

ASA also supports ethertype ACLs, which are discussed in the transparent mode task for the ASA, and webtype ACLs, which are discussed in the VPN section. Note that for FTP and TFTP, we only allow the control session; the secondary session will automatically be allowed, because the MPF engine of the ASA inspects FTP and TFTP by default. This applies for ICMP as well, but only because ICMP inspection was enabled from the initial configuration files; by default it is not enabled. For UNIX style traceroute, you must allow the returning ICMP unreachable and time-exceeded messages.

Configuration

ASA1:

```
access-list VLAN49_INBOUND permit tcp any host 150.1.2.2 range 22 23
access-list VLAN49_INBOUND permit tcp host 150.1.1.1 any eq 21
access-list VLAN49_INBOUND permit tcp host 150.1.1.1 any eq 443
access-list VLAN49_INBOUND permit udp host 150.1.1.1 any eq 69
access-list VLAN49_INBOUND permit icmp any any unreachable
access-list VLAN49_INBOUND permit icmp any any time-exceeded
!
access-group VLAN49_INBOUND in interface VLAN49
```

ASA2:

```
access-list GLOBAL remark *** PERMIT TASK TRAFFIC ***
access-list GLOBAL extended permit icmp host 150.1.22.22 any echo
access-list GLOBAL extended permit tcp any host 172.16.10.100 eq 80
access-list GLOBAL extended permit tcp any host 172.16.10.100 eq 3389
```

https://t.me/learningnets

```
access-list GLOBAL extended permit udp any host 172.16.10.100 eq 514
access-list GLOBAL remark *** PERMIT RESTRICTED TRAFFIC DUE TO USING GLOBAL ACL ***
access-list GLOBAL permit ip 172.16.10.0 255.255.255.0 136.1.29.0 255.255.255.0
access-list GLOBAL permit ip 172.16.10.0 255.255.255.0 host 150.1.22.22
access-list GLOBAL permit ip 136.1.19.0 255.255.255.0 any
access-list GLOBAL permit ip 136.1.27.0 255.255.255.0 any
access-list GLOBAL permit ip host 150.1.11.11 any
!
access-group GLOBAL global
```

Verification

Test that telnet traffic is allowed toward R2's Loopback0 address; verify the match/hitcount on the configured ACL:

```
R1#telnet 150.1.2.2
Trying 150.1.2.2 ... Open


User Access Verification

Password: R2>
!
!ASA1# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list VLAN49_INBOUND; 6 elements; name hash: 0x1dc0bbb8
access-list VLAN49_INBOUND line 1 extended permit tcp any host 150.1.2.2 range ssh telnet (hitcnt=1) 0xc0f6e585

access-list VLAN49_INBOUND line 2 extended permit tcp host 150.1.1.1 any eq ftp (hitcnt=0) 0xe77ee376
access-list VLAN49_INBOUND line 3 extended permit tcp host 150.1.1.1 any eq https (hitcnt=0) 0x2da2195a
access-list VLAN49_INBOUND line 4 extended permit udp host 150.1.1.1 any eq tftp (hitcnt=0) 0x7bb827ad
access-list VLAN49_INBOUND line 5 extended permit icmp any any unreachable (hitcnt=0) 0xfe2f5f42
access-list VLAN49_INBOUND line 6 extended permit icmp any any time-exceeded (hitcnt=0) 0xd46954bb
```

Verify other traffic (for example, FTP using the packet-tracer tool), and check the match/hitcount on the configured ACL:

```
ASA1# packet-tracer input VLAN49 tcp 150.1.1.1 20000 150.1.2.2 21

Phase: 1
Type: ROUTE-LOOKUP
Subtype: input
Result: ALLOW
```

https://t.me/learningnets

```
Config:
Additional Information:
in   150.1.2.2       255.255.255.255 VLAN59

Phase: 2
Type: ACCESS-LIST
Subtype: log
Result: ALLOW
Config:access-group VLAN49_INBOUND in interface VLAN49
access-list VLAN49_INBOUND extended permit tcp host 150.1.1.1 any eq ftp
Additional Information:

Phase: 3
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 4
Type: INSPECTSubtype: inspect-ftp
Result: ALLOW
Config:
class-map inspection_default
 match default-inspection-traffic
policy-map global_policy
 class inspection_default
  inspect ftp
service-policy global_policy global
Additional Information:

Phase: 5
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 6
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Config:
Additional Information:
New flow created with id 22, packet dispatched to next module
```

https://t.me/learningnets

```
Result:
input-interface: VLAN49
input-status: up
input-line-status: up
output-interface: VLAN59
output-status: up
output-line-status: upAction: allow
!
!ASA1# packet-tracer input VLAN49 tcp 136.1.49.1 20000 150.1.2.2 21

Phase: 1
Type: ROUTE-LOOKUP
Subtype: input
Result: ALLOW
Config:
Additional Information:
in   150.1.2.2       255.255.255.255 VLAN59

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: DROP
Config:Implicit Rule
Additional Information:

Result:
input-interface: VLAN49
input-status: up
input-line-status: up
output-interface: VLAN59
output-status: up
output-line-status: up
Action: dropDrop-reason: (acl-drop) Flow is denied by configured rule
!
!ASA1# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list VLAN49_INBOUND; 6 elements; name hash: 0x1dc0bbb8
access-list VLAN49_INBOUND line 1 extended permit tcp any host 150.1.2.2 range ssh telnet (hitcnt=1) 0xc0f6e585
access-list VLAN49_INBOUND line 2 extended permit tcp host 150.1.1.1 any eq ftp (hitcnt=1) 0xe77ee376

access-list VLAN49_INBOUND line 3 extended permit tcp host 150.1.1.1 any eq https (hitcnt=0) 0x2da2195a
access-list VLAN49_INBOUND line 4 extended permit udp host 150.1.1.1 any eq tftp (hitcnt=0) 0x7bb827ad
access-list VLAN49_INBOUND line 5 extended permit icmp any any unreachable (hitcnt=0) 0xfe2f5f42
```

https://t.me/learningnets

```
access-list VLAN49_INBOUND line 6 extended permit icmp any any time-exceeded (hitcnt=0) 0xd46954bb
```

Verify that traceroute works, and verify the match/hitcount on the configured ACL; because R1 is the next-hop but also the last, there is no ICMP Exceeded message, just ICMP Port Unreachable:

```
R2#traceroute 150.1.1.1
Type escape sequence to abort.
Tracing the route to 150.1.1.1
VRF info: (vrf in name/id, vrf out name/id) 1 136.1.49.1 0 msec *  0 msec
!
!ASA1# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list VLAN49_INBOUND; 6 elements; name hash: 0x1dc0bbb8
access-list VLAN49_INBOUND line 1 extended permit tcp any host 150.1.2.2 range ssh telnet (hitcnt=1) 0xc0f6e585
access-list VLAN49_INBOUND line 2 extended permit tcp host 150.1.1.1 any eq ftp (hitcnt=1) 0xe77ee376
access-list VLAN49_INBOUND line 3 extended permit tcp host 150.1.1.1 any eq https (hitcnt=0) 0x2da2195a
access-list VLAN49_INBOUND line 4 extended permit udp host 150.1.1.1 any eq tftp (hitcnt=0) 0x7bb827ad
access-list VLAN49_INBOUND line 5 extended permit icmp any any unreachable (hitcnt=2) 0xfe2f5f42
access-list VLAN49_INBOUND line 6 extended permit icmp any any time-exceeded (hitcnt=0) 0xd46954bb
```

Verify that PING traffic from R2's Loopback1 address is functional:

```
R2#ping 150.1.11.11 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.11.11, timeout is 2 seconds:
Packet sent with a source address of 150.1.22.22 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R2#ping 150.1.11.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.11.11, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!
ASA2# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list GLOBAL; 9 elements; name hash: 0xbaf7cf02
access-list GLOBAL line 1 remark *** PERMIT TASK TRAFFIC ***
access-list GLOBAL line 2 extended permit icmp host 150.1.22.22 any echo (hitcnt=1) 0x01304b27

access-list GLOBAL line 3 extended permit tcp any host 172.16.10.100 eq www (hitcnt=0) 0xac48b904
```

https://t.me/learningnets

```
access-list GLOBAL line 4 extended permit tcp any host 172.16.10.100 eq 3389 (hitcnt=0) 0x3294785b
access-list GLOBAL line 5 extended permit udp any host 172.16.10.100 eq syslog (hitcnt=0) 0xa694a198
access-list GLOBAL line 6 remark *** PERMIT RESTRICTED TRAFFIC DUE TO USING GLOBAL ACL ***
access-list GLOBAL line 7 extended permit ip 172.16.10.0 255.255.255.0 136.1.29.0 255.255.255.0 (hitcnt=0) 0x8a69891
access-list GLOBAL line 8 extended permit ip 172.16.10.0 255.255.255.0 host 150.1.22.22 (hitcnt=0) 0x70da529b
access-list GLOBAL line 9 extended permit ip 136.1.19.0 255.255.255.0 any (hitcnt=0) 0xff304bca
access-list GLOBAL line 10 extended permit ip 136.1.27.0 255.255.255.0 any (hitcnt=0) 0x2a46c5e3
access-list GLOBAL line 11 extended permit ip host 150.1.11.11 any (hitcnt=0) 0xaa5f99fe
```

Check that traffic allowed initially by the security-level rules is still functional:

```
R1#telnet 136.1.29.2
Trying 136.1.29.2 ... Open

User Access Verification

Password: R2>
!
!ASA2# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list GLOBAL; 9 elements; name hash: 0xbaf7cf02
access-list GLOBAL line 1 remark *** PERMIT TASK TRAFFIC ***
access-list GLOBAL line 2 extended permit icmp host 150.1.22.22 any echo (hitcnt=1) 0x01304b27
access-list GLOBAL line 3 extended permit tcp any host 172.16.10.100 eq www (hitcnt=0) 0xac48b904
access-list GLOBAL line 4 extended permit tcp any host 172.16.10.100 eq 3389 (hitcnt=0) 0x3294785b
access-list GLOBAL line 5 extended permit udp any host 172.16.10.100 eq syslog (hitcnt=0) 0xa694a198
access-list GLOBAL line 6 remark *** PERMIT RESTRICTED TRAFFIC DUE TO USING GLOBAL ACL ***
access-list GLOBAL line 7 extended permit ip 172.16.10.0 255.255.255.0 136.1.29.0 255.255.255.0 (hitcnt=0) 0x8a69891
access-list GLOBAL line 8 extended permit ip 172.16.10.0 255.255.255.0 host 150.1.22.22 (hitcnt=0) 0x70da529b
access-list GLOBAL line 9 extended permit ip 136.1.19.0 255.255.255.0 any (hitcnt=1) 0xff304bca

access-list GLOBAL line 10 extended permit ip 136.1.27.0 255.255.255.0 any (hitcnt=0) 0x2a46c5e3
access-list GLOBAL line 11 extended permit ip host 150.1.11.11 any (hitcnt=0) 0xaa5f99fe
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Object-Group Access-Lists

Task

- Create the following object groups on ASA2:
  - **R2**, including R2's Loopback1 and VLAN 29 subnets.
  - **SUBNETS**, including VLAN 19 and VLAN 26.
  - **R1_LOOPBACK1**, including R1's Loopback1 subnet.
  - **TELNET**, including port for telnet.
  - **TFTP**, including port for tftp.
  - **OTHER_PORTS**, including ports for FTP and NTP.
- Configure ASA2 to allow following traffic inbound on its VLAN 29 interface:
  - TELNET/TFTP/FTP/NTP traffic from **R2** to **SUBNETS** and **R1_LOOPBACK1** .
- The ACL should use exactly three entries.
  - Additional object-groups can be configured if required.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

https://t.me/learningnets

Overview

Object groups allow simplifying of large access-list configuration. You can group objects of similar nature (such as a group of networks and host, a collection of TCP ports, or a bunch of ICMP message types) and reference them in access-lists. Thus, instead of working with addresses and port numbers, you can work with higher-level objects that reflect the logical structure of your network. For example, you may have an object group named **PUBLIC_HOSTING** that lists the publically accessible servers and an object group named **MANAGEMENT_SEGMENT** that lists the management stations, along with an object group named **PUBLIC_PORTS** that lists the FTP, WWW, and HTTPS ports. By building your access-list out of objects groups, you make them more readable and manageable, because you don't need to add new ACL entries for every new public server, management server, or port.

ASA supports the following object group types:

- **Network**, matching on IPv4/IPv6 hosts or subnets.
- **Protocol**, matching on layer3/layer4 IP protocols.
- **ICMP-Type**, matching on ICMP Types.
- **Service**, matching on TCP/UDP ports; a service object group supports multiple sub-types: TCP, UDP, TCP-UDP, Generic (matching on a mixture of source or destination TCP/UDP ports).

Additionally, ASA supports object-group nesting, where multiple object groups of the same type (and sub-type service object groups) can be further aggregated by another object group. You can build ACLs by using only object-groups, the general command being `access-list <name> permit [<service_object_group>|<protocol_object_group>] object-group <network_for_source> object-group <service_for_source> object-group <network_for_destination> object-group <service_for_destination>` . However, you must ensure that configuration is consistent; for example, you cannot use a protocol object-group matching on ip and attach a service object-group for source or destination.

Configuration

ASA2:

```
object-group network R2
 network-object host 150.1.22.22
 network-object 136.1.29.0 255.255.255.0
!
object-group network SUBNETS
```

https://t.me/learningnets

```
 network-object 136.1.19.0 255.255.255.0
 network-object 172.16.10.0 255.255.255.0
!
object-group network R1_LOOPBACK1
 network-object host 150.1.11.11
!
object-group network ALL_DESTINATIONS
 group-object SUBNETS
 group-object R1_LOOPBACK1
!
object-group service TELNET tcp
 port-object eq 23
!
object-group service TFTP udp
 port-object eq 69
!
object-group service OTHER_PORTS
 service-object tcp destination eq 21
 service-object udp destination eq 123
!
access-list VLAN29_INBOUND extended permit tcp object-group R2 object-group ALL_DESTINATIONS object-group TELNET
access-list VLAN29_INBOUND extended permit udp object-group R2 object-group ALL_DESTINATIONS object-group TFTP
access-list VLAN29_INBOUND extended permit object-group OTHER_PORTS object-group R2 object-group ALL_DESTINATIONS
!
access-group VLAN29_INBOUND in interface VLAN29
```

Verification

View the access-lists in both forms:

```
ASA2# show running-config access-list
access-list VLAN29_INBOUND extended permit tcp object-group R2 object-group ALL_DESTINATIONS object-group TELNET
access-list VLAN29_INBOUND extended permit udp object-group R2 object-group ALL_DESTINATIONS object-group TFTP
access-list VLAN29_INBOUND extended permit object-group OTHER_PORTS object-group R2 object-group ALL_DESTINATIONS
!
!ASA2# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list VLAN29_INBOUND; 24 elements; name hash: 0xd515f651
access-list VLAN29_INBOUND line 1 extended permit tcp object-group R2 object-group ALL_DESTINATIONS object-group TEL
  access-list VLAN29_INBOUND line 1 extended permit tcp host 150.1.22.22 136.1.19.0 255.255.255.0 eq telnet (hitcnt=
  access-list VLAN29_INBOUND line 1 extended permit tcp host 150.1.22.22 172.16.10.0 255.255.255.0 eq telnet (hitcnt
  access-list VLAN29_INBOUND line 1 extended permit tcp host 150.1.22.22 host 150.1.11.11 eq telnet (hitcnt=0) 0x5f4
  access-list VLAN29_INBOUND line 1 extended permit tcp 136.1.29.0 255.255.255.0 136.1.19.0 255.255.255.0 eq telnet
  access-list VLAN29_INBOUND line 1 extended permit tcp 136.1.29.0 255.255.255.0 172.16.10.0 255.255.255.0 eq telnet
```

https://t.me/learningnets

```
  access-list VLAN29_INBOUND line 1 extended permit tcp 136.1.29.0 255.255.255.0 host 150.1.11.11 eq telnet (hitcnt=
access-list VLAN29_INBOUND line 2 extended permit udp object-group R2 object-group ALL_DESTINATIONS object-group TFT

  access-list VLAN29_INBOUND line 2 extended permit udp host 150.1.22.22 136.1.19.0 255.255.255.0 eq tftp (hitcnt=0)
  access-list VLAN29_INBOUND line 2 extended permit udp host 150.1.22.22 172.16.10.0 255.255.255.0 eq tftp (hitcnt=0
  access-list VLAN29_INBOUND line 2 extended permit udp host 150.1.22.22 host 150.1.11.11 eq tftp (hitcnt=0) 0x2dd40
  access-list VLAN29_INBOUND line 2 extended permit udp 136.1.29.0 255.255.255.0 136.1.19.0 255.255.255.0 eq tftp (h
  access-list VLAN29_INBOUND line 2 extended permit udp 136.1.29.0 255.255.255.0 172.16.10.0 255.255.255.0 eq tftp (
  access-list VLAN29_INBOUND line 2 extended permit udp 136.1.29.0 255.255.255.0 host 150.1.11.11 eq tftp (hitcnt=0)
access-list VLAN29_INBOUND line 3 extended permit object-group OTHER_PORTS object-group R2 object-group ALL_DESTINAT

  access-list VLAN29_INBOUND line 3 extended permit tcp host 150.1.22.22 136.1.19.0 255.255.255.0 eq ftp (hitcnt=0)
  access-list VLAN29_INBOUND line 3 extended permit tcp host 150.1.22.22 172.16.10.0 255.255.255.0 eq ftp (hitcnt=0)
  access-list VLAN29_INBOUND line 3 extended permit tcp host 150.1.22.22 host 150.1.11.11 eq ftp (hitcnt=0) 0x40fa51
  access-list VLAN29_INBOUND line 3 extended permit tcp 136.1.29.0 255.255.255.0 136.1.19.0 255.255.255.0 eq ftp (hi
  access-list VLAN29_INBOUND line 3 extended permit tcp 136.1.29.0 255.255.255.0 172.16.10.0 255.255.255.0 eq ftp (h
  access-list VLAN29_INBOUND line 3 extended permit tcp 136.1.29.0 255.255.255.0 host 150.1.11.11 eq ftp (hitcnt=0)
  access-list VLAN29_INBOUND line 3 extended permit udp host 150.1.22.22 136.1.19.0 255.255.255.0 eq ntp (hitcnt=0)
  access-list VLAN29_INBOUND line 3 extended permit udp host 150.1.22.22 172.16.10.0 255.255.255.0 eq ntp (hitcnt=0)
  access-list VLAN29_INBOUND line 3 extended permit udp host 150.1.22.22 host 150.1.11.11 eq ntp (hitcnt=0) 0xc0fdae
  access-list VLAN29_INBOUND line 3 extended permit udp 136.1.29.0 255.255.255.0 136.1.19.0 255.255.255.0 eq ntp (hi
  access-list VLAN29_INBOUND line 3 extended permit udp 136.1.29.0 255.255.255.0 172.16.10.0 255.255.255.0 eq ntp (h
  access-list VLAN29_INBOUND line 3 extended permit udp 136.1.29.0 255.255.255.0 host 150.1.11.11 eq ntp (hitcnt=0)
```

Test, for example, telnet traffic:

```
R2#telnet 150.1.11.11 /source-interface loopback1
Trying 150.1.11.11 ... Open



User Access Verification

Password: R1>
!
!ASA2# show access-list | i hitcnt=1

access-list VLAN29_INBOUND line 1 extended permit tcp host 150.1.22.22 host 150.1.11.11 eq telnet (hitcnt=1) 0x5f4ae
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Identity Firewall

Task

- Configure ASA2 to allow user **USER1** and all users from AD group **CCIELAB** access to any resources.
- For LDAP integration, use **ldapuser** with the password **1qaz@WSX**; use the same password for user **USER1**.
- Read Active Directory for user-to-group mappings once each hour and consider a user as idle if no traffic is received for 30 minutes.
- If the domain controller or the AD agent is unreachable, disable the IP-to-user mappings.
- Apply filtering inbound on VLAN 19 interface.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Identity Firewall functionality was added into ASA code 8.4(2), and it allows configuring access-list rules based on users or groups from Active Directory (only for the source; the destination is still an IP address or subnet); it works for both IPv4

https://t.me/learningnets

and IPv6 access-lists, regardless of the firewall mode (single or multiple context, routed or transparent). This is a major improvement that allows for a much more convenient way of configuring firewall rules within an Active Directory environment, moving the decision away from the IP address paradigm toward a role (user/group from AD). Because now a user can have multiple devices connected over different media types (wired, wireless, VPN), multiple IP addresses represent the same user; configuring access-lists by referencing the user will optimize the ACL.

For configuration, use the following steps:

1. Install and configure the AD agent.
2. Configure the ASA to communicate with both AD server (LDAP) and AD Agent.
3. Enable identity firewall functionality and configure optional parameters.
4. Configure identity firewall rules.

ASA communicates with the LDAP server to retrieve the user-to-group mappings, and communicates with the AD Agent (through RADIUS) to retrieve the user-to-IP mappings. Through the Identity firewall configuration, the users are not authenticated by the ASA on-demand but transparent, because the AD Agent reads the domain logs for successful domain user authentication and passes this information to the ASA in the form of user-to-IP mappings.

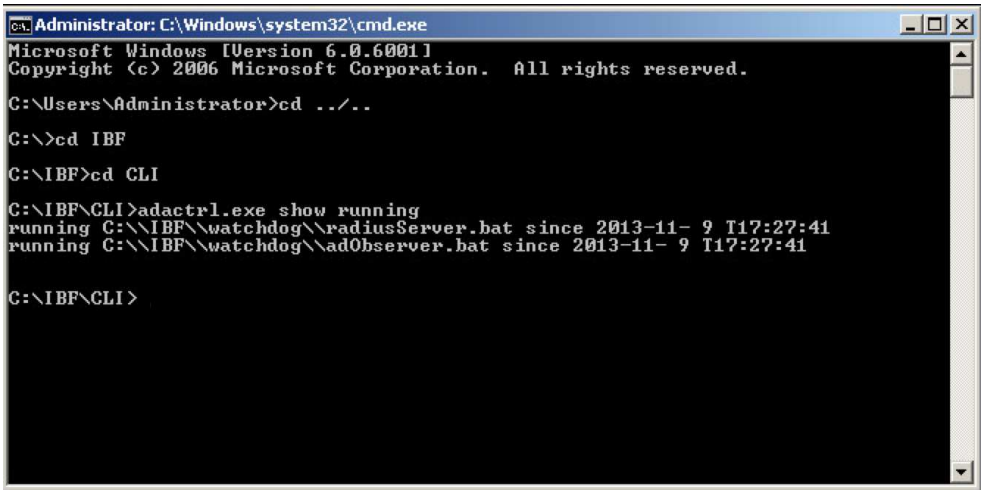
For the scope of identity firewall, another object group of type **user** can be used, allowing to match on users and groups from Active Directory; alternatively, users and groups can be directly specified within the ACL rules. Note that where it's required to specify the domain name in ASA configuration, you must specify the NETBIOS name, not the FQDN, and this is case-sensitive.

Configuration

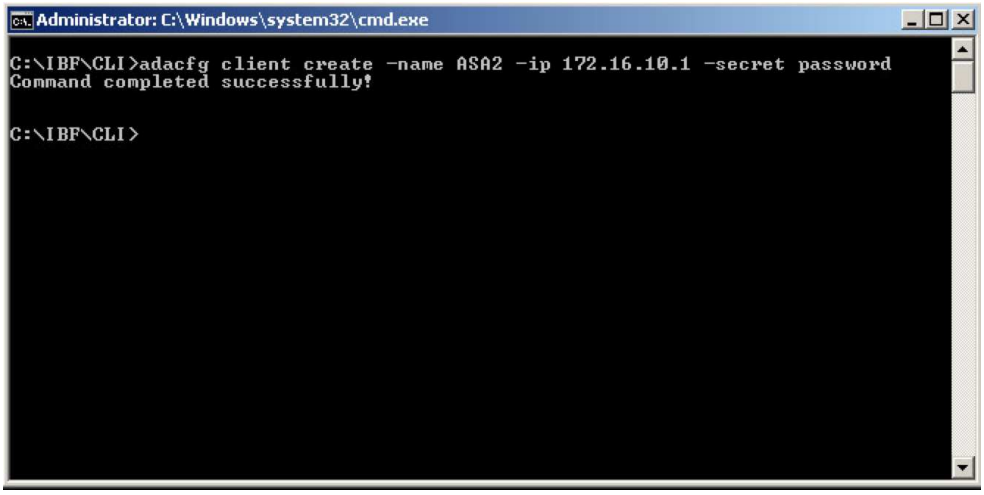
The AD Agent should already be installed on **TEST-SRV-A**. Verify the existence of the **C:\IBF** directory. If it's not present, install the AD Agent found in the **Lab Files** folder from the desktop of **TEST-SRV-A**; if the installer indicates that it's already installed, re-install it and choose **yes** on all prompts. Restart **TEST-SRV-A** after this.

Verify that the AD Agent is running, using the command `adactrl.exe show running` from the **C:\IBF\CLI** directory; if output is different, make sure that you restarted the Windows machine, or alternatively use the command `adactrl.exe restart` to restart the service:

https://t.me/learningnets

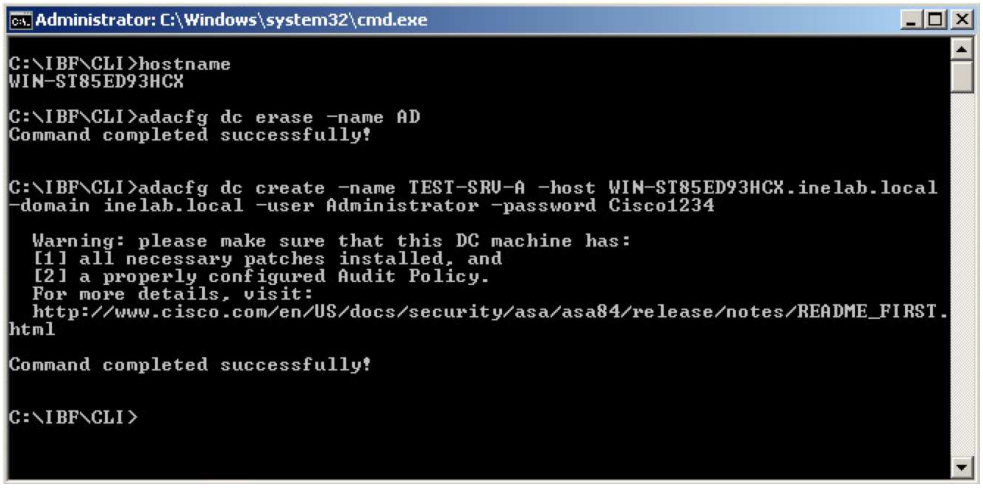


Configure ASA2 as a client for the AD Agent using the command **`adacfg client create -name <string> -ip <IP Address> -secret <string>`** ; to remove a client entry, use the command **`adacfg client erase - name <string>`** :

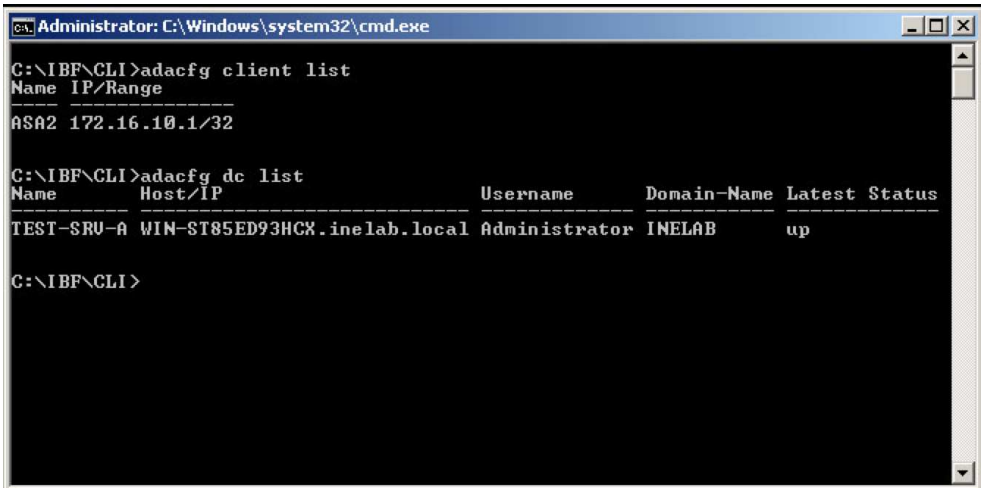


Configure a connection between the AD Agent and Domain Controller using the command **`adacfg dc create -name <string> -host <dc_netbios_name> -domain <DC_domain_name> -user <username> -password <string>`** . To remove a DC entry, use the command **`adacfg dc erase -name <string>`** . If a connection for the domain is already present, it must be removed first; although the command requires the NETBIOS name to be used, because of some current AD Agent limitations, the FQDN should be used instead:

https://t.me/learningnets



Use the commands **`adacfg client list`** and **`adacfg dc list`** to verify your configuration; the status of the connection with the Domain Controller should be **UP**:



ASA2:

```
aaa-server IDENTITY protocol radius
 ad-agent-mode
!
aaa-server IDENTITY (VLAN26) host 172.16.10.100
 key password
!
aaa-server AD protocol ldap
!
aaa-server AD (VLAN26) host 172.16.10.100
 ldap-base-dn dc=inelab,dc=local
 ldap-scope subtree
 ldap-naming-attribute sAMAccountName
 ldap-login-password 1qaz@WSX
 ldap-login-dn cn=ldapuser,cn=Users,dc=inelab,dc=local
 server-type microsoft
!
```

https://t.me/learningnets

```
user-identity enable
user-identity default-domain inelab
user-identity domain INE aaa-server AD
user-identity ad-agent aaa-server IDENTITY
!
user-identity inactive-user-timer minutes 30
user-identity poll-import-user-group-timer hours 1
!
user-identity action domain-controller-down INE disable-user-identity-rule
user-identity action ad-agent-down disable-user-identity-rule
!
object-group user FIREWALL
 user inelab\USER1
 user-group inelab\\CCIELAB
!
access-list VLAN19_INBOUND permit ip object-group-user FIREWALL any any
access-group VLAN19_INBOUND in interface VLAN19
```

Verification

Verify that ASA communication with the AD Agent is in the **UP** state:

```
ASA2# show user-identity ad-agent
Primary AD Agent: Status                    up (registered)
 Mode:                      full-download
 IP address:                172.16.10.100
 Authentication port:       udp/1645
 Accounting port:           udp/1646
 ASA listening port:        udp/3799
 Interface:                 VLAN26
 Up time:                   6 mins 8 secs
 Average RTT:               0 msec

AD Domain Status: Domain INELAB:            up
!
!ASA2# test aaa-server ad-agent IDENTITY
Server IP Address or name: 172.16.10.100
INFO: Attempting Ad-agent test to IP address <172.16.10.100> (timeout: 12 seconds)
INFO: Ad-agent Successful
```

Query the AD server for the required usernames and groups:

https://t.me/learningnets

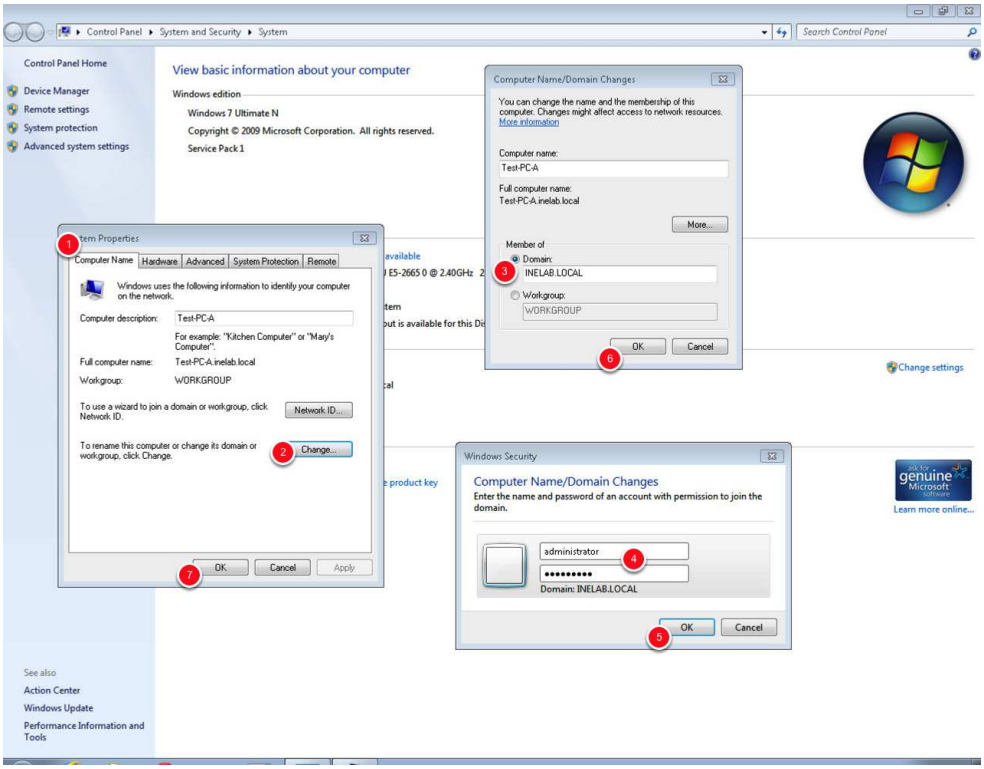
```
ASA2# show user-identity ad-groups INE filter CCIELAB


Domain:INE      AAA Server Group: AD
Group list retrieved successfully
Number of Active Directory Groups: 1 dn: CN=CCIELAB,OU=CCIELAB,DC=inelab,DC=local
sAMAccountName: CCIELAB
!
!ASA2# show user-identity ad-users INE filter USER3

Domain:INE      AAA Server Group: AD
User list retrieved successfully
Number of Active Directory Users: 1 dn: CN=USER3,OU=CCIELAB,DC=inelab,DC=local

sAMAccountName: USER3
```

To test, join **TEST-PC-A** to the Active Directory domain, right-click **Computer**, select **Properties**, and use **Administrator** the password **Cisco1234** for domain joining:



After **TEST-PC-A** restarts, authenticate in the Active Directory domain using the user **USER1** and the password **1qaz@WSX**. Initiate a telnet session from **TEST-PC-A** to R2's VLAN 29 IP address; it should be functional:

```
ASA2# show user-identity ip-of-user USER1
inelab\136.1.27.51 (Login)
```

https://t.me/learningnets

```
!
!ASA2# show user-identity user active user inelab\USER1 list detail
  inelab\user1: 5 active conns; idle 0 mins 136.1.27.51: login 11 mins, idle 0 mins, 5 active conns
!
!ASA2# show conn user inelab\USER1
6 in use, 53 most used
TCP VLAN29 136.1.29.2:23 VLAN19 (inelab\user1) 136.1.27.51:33788, idle 0:01:34, bytes 103, flags UIO

TCP VLAN26 172.16.10.100:49156 VLAN19 (inelab\user1) 136.1.27.51:1117, idle 0:44:05, bytes 3554, flags UIO
TCP VLAN26 172.16.10.100:135 VLAN19 (inelab\user1) 136.1.27.51:1115, idle 0:44:05, bytes 660, flags UIO
TCP VLAN26 172.16.10.100:49156 VLAN19 (inelab\user1) 136.1.27.51:1051, idle 0:44:11, bytes 5807, flags UIO
TCP VLAN26 172.16.10.100:135 VLAN19 (inelab\user1) 136.1.27.51:1050, idle 0:44:11, bytes 816, flags UIO
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Static NAT and PAT

Task

- Configure static NAT on ASA1:
  - Map R1's Loopback0 address to 136.1.59.101.
  - HTTP sessions toward ASA1's VLAN 59 address should be redirected to R1's VLAN 49 address.
  - TCP session toward ASA1's VLAN 49 address on port 2023 should be redirected to R2's VLAN 59 address on telnet service.
- Configure static NAT on ASA2:
  - Map R2's Loopback1 address to 136.1.19.101.
  - HTTP sessions toward ASA2's VLAN 19 address should be redirected to R2's VLAN 29 address.
  - TCP session toward ASA2's VLAN 29 address on port 2023 should be redirected to R1's VLAN 19 address on telnet service.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

https://t.me/learningnets

Overview

NAT architecture on ASA was completely redesigned starting with code 8.3(1); therefore, based on the ASA software versions present in the current CCIE Security Blueprint, both the old way and new way of NAT will be tested. In the following tasks, we'll explore configurations for all NAT options. Given that traffic could match multiple NAT rules, it's important to understand the order of operation.

For before 8.3(1), the following is the order of operation:

- NAT Exemption
- Static NAT (identity, regular and policy); rules are processed in order until first match; processing order is given by configuration order, thus is important
- Dynamic Policy NAT; rules are processed in order until first match; processing order is given by configuration order, thus is important
- Dynamic NAT (identity and regular); it works on a best match basis, thus order of configuration is not important

For 8.3(1) and later, the following is the order of operation:

- Manual NAT or Twice NAT; rules are processed in order until first match; processing order is given by configuration order, thus is important
- Static Object NAT or Auto NAT; rules are ordered based on several factors: number of IP addresses, lowest numerical IP addresses, lowest alphabetical name of object
- Dynamic Object NAT or Auto NAT; rules are ordered based on several factors: number of IP addresses, lowest numerical IP addresses, lowest alphabetical name of object
- After-auto Manual NAT or Twice NAT; rules are processed in order until first match; processing order is given by configuration order, thus is important

Old NAT Syntax

A static NAT rule forms a bidirectional one-to-one mapping between local and global IP addresses, unlike a dynamic rule which usually maps a number of local addresses to a number of global address. Static rules is always present in the NAT table of the ASA, so a session can be initiated from both sides, as long as the default security-level rules or applied access-lists allow the session. Static translations are used for the following purposes (assume **inside** is the trusted network and **outside** is the untrusted network):

-

https://t.me/learningnets

To allow accessing a server located on the inside using a fixed global outside IP, using the command syntax `static (<real_interface>,<mapped_interface>) [<mapped_IP>|interface] <real_IP>` . For example, with `static (inside,outside) 136.1.59.100 172.16.10.100` , all sessions initiated from local 172.16.10.100 on the inside interface will be source translated to global 136.1.59.100 when exiting on the outside interface, and all sessions destined to global 136.1.59.100 on the outside interface will be destination translated to local 172.16.10.100 and forced to exit the inside interface. ASA interface can be used as well (for example, `static (inside,outside) interface 172.16.10.100`), but it is not recommened because all traffic destined to the outside interface of the ASA will be redirected to 172.16.10.100.

- To redirect a particular port at the global IP to the port at the inside IP (aka static PAT or port redirection), using the command syntax `static (<real_interface>,<mapped_interface>) [tcp|udp] [<mapped_IP>|interface] <mapped_port> <real_IP> <real_port>` . For example, with `static (inside,outside) tcp 136.1.59.100 4545 172.16.10.100 45` , all TCP sessions destined to global 136.1.59.100 on port 4545 will be redirected (destination translated) to local 172.16.10.100 on port 45. Using static PAT instead of static NAT allows for publishing multiple services found on different real IP addresses using a single mapped IP address. ASA interface is commonly used for static PAT, because it does not disable all services on the ASA for that interface. It still can be reached except on the ports used in PAT; for example, `static (inside,outside) tcp interface 4545 172.16.10.100 45` will redirect all TCP port 4545 traffic toward ASA on the outside interface to 172.16.10.100 on port 45, whereas traffic for all other ports and IP protocols is still destined to ASA.
- For reverse redirection. Sometimes you may want to redirect a connection going to an inside host to some outside destination. This is often called routing simplification, because it might be used in situations where inside hosts lack routing information (they have no default gateway set, or they have the default route pointing toward some router other than the firewall). For example, imagine that you have a DNS server on the outside with the IP address 136.1.122.200 but you want the inside hosts to use the IP 172.16.10.200 to access the DNS server (for example, the hosts have no default gateway set). The following “reverse” construct (local and global

https://t.me/learningnets

interfaces are swapped) would achieve the goal: `static (outside,inside) 172.16.10.200 136.1.122.200` . ASA will do Proxy-ARP for the IP address 172.16.10.200 on its inside interface, provided that 172.16.10.0 is the inside's interface subnet; otherwise, this subnet should be explicitly routed to the firewall. Remember that the inside hosts most likely need their sources translated when accessing the outside DNS. Thus, you may need to add a dynamic NAT rule as well:

- To block/subnet translation. For example, the command `static (inside,outside) 172.16.10.0 192.168.0.0 netmask 255.255.255.0` will establish bidirectional mapping between the subnets 172.16.10.0/24 and 192.168.0.0/24 mapping 172.16.10.1 to 192.168.0.1, and so on.

With the static command, ASA has some additional optional parameters:

https://t.me/learningnets

- `tcp <max_conn> [<max_embryonic>]` is used to prevent DoS attacks against the corresponding mapped/global IP address. This is a legacy way to configure the anti-DoS settings and is now configured through MPF. However, **max_conn** is the maximum number of concurrent TCP connections established to the mapped IP address. The other parameter, **max_embryonic**, specifies the maximum number of allowed embryonic connections, TCP connections that have not yet finished the 3-way TCP handshake. The goal is to prevent resource starvation on the protected server and protect it against SYN-flooding attack. The default limit is set to zero for both parameters, which means an unlimited number of connections.
- `udp <max_conn>` has the same purpose as on TCP, but for UDP sessions; this has also been migrated to MFP configuration.
- `norandomseq` . By default, ASA inspects any TCP sessions and modifies TCP Initial Sequence Number (ISN) to a truly random value. This is meant to prevent TCP connection hijacking when a hacker might guess the ISN and inject seemingly correct packets into TCP flow. In some cases you may need to disable this option; this is most commonly necessary if the TCP session header is authenticated in some way, such as by using the MD5 hash option. A good example is an authenticated BGP peering session. The other scenario is that there is another firewall that has already done the ISN randomization. This option might be selectively enabled and disabled using the MPF as well.

New NAT Syntax

With the new NAT syntax, **static** command has been removed and everything is configured using the **nat** command through object-nat or twice-nat. Between twice-nat and after-auto twice-nat there is no functional difference, it just modifies the NAT order of processing as specified in the task overview. The main differences between object-nat and twice-nat are:

- With object-nat the `nat` command is configured under the defined object, whereas with twice-nat the `nat` command is configured at the global level.
- Object-nat does not allow both source and destination NAT within a single rule, whereas twice-nat does.
- Within object-nat rules, static rules are processed first, before dynamic rules, whereas twice-nat rules are processed in order based on the line/sequence number

https://t.me/learningnets

(so dynamic rules may be processed ahead of static rules).
- Object-nat can only reference network objects, whereas twice-nat can reference both network and service objects.

Although both object-nat and twice-nat can be used to configure static NAT or static PAT, from a practical point of view it makes more sense to use object-nat. An **object** is a new addition to the 8.3(1) code, required for configuring object-NAT. There are two types of objects, **network** and **service**, and although an **object** is very similar to **object groups**, there is one important difference: an **object** can contain only a single entry:

- **network** can contain one of the following: a host, a subnet, or a range of IP addresses.
- **service** can contain one of the following: a source UDP or TCP port, a destination UDP or TCP port, a source and destination UDP port, a source and destination TCP port, or an IP protocol.

To configure static object-nat, follow these steps:

1. Create a network object using the command **`object network <name> .`**
2. Configure a single entry within the object, defining the real IP address, real subnet, or IP address range using one of these commands: **`host <IP Address> , subnet <network> <mask> , range <start_ip> <end_ip>`** . If multiple entries are configured, the latest one will always override the previous one.
3. Configure static object-nat using the object-level command **`nat ([<real_interface>|any] [<mapped_interface>|any]) static [<mapped_IP>|interface] service [tcp|udp] <real_svc> <mapped_svc>`** ; the **any** keyword can be used for source and/or destination interface, which makes the NAT rule valid regardless of which interface the traffic is coming from or going to.

Because the static NAT scope has been already discussed, let's study some examples of mapping the old NAT to the new NAT:

```
! Old NAT Syntax
!
static (inside,outside) 136.1.59.100 172.16.10.100
!
!New Object NAT Syntax
!
object network NAT1
 host 172.16.10.100
 nat (inside,outside) static 136.1.59.100
!
```

https://t.me/learningnets

```
!New Twice NAT Syntax
!
object network REAL
 host 172.16.10.100
!
object-group network MAPPED
 network-object host 136.1.59.100
!
nat (inside,outside) source static REAL MAPPED
!
!New Twice NAT Syntax
!
object network REAL
 host 172.16.10.100
!
object-group network MAPPED
 network-object host 136.1.59.100
!
nat (outside,inside) source static any any destination static MAPPED REAL
```

Except for the static NAT configuration, depending on the security-levels and direction of traffic flow, applying an ACL might be required to allow the session. Based on the ASA version:

- If ASA uses a code version prior to 8.3(1), thus using the old NAT syntax, in the ACL you must match on the global/mapped IP address.
- If ASA uses a code version 8.3(1)or later, thus using the new NAT syntax, in the ACL you must match on the local/real IP address.

In the configurations below, we could have used only object-NAT or only twice-NAT; both types were used just to show the configuration differences.

Configuration

ASA1:

```
static (VLAN49,VLAN59) 136.1.59.101 150.1.1.1
static (VLAN49,VLAN59) tcp interface 80 136.1.49.1 80
static (VLAN59,VLAN49) tcp interface 2023 136.1.59.2 23
!
access-list VLAN49_INBOUND permit tcp any host 136.1.49.15 eq 2023
access-group VLAN49_INBOUND in interface VLAN49
```

ASA2:

https://t.me/learningnets

```
object network R2_LO1_REAL
 host 150.1.22.22
 nat (VLAN29,any) static 136.1.19.101
!
object network R2_VLAN29_REAL
 host 136.1.29.2
 nat (VLAN29,VLAN19) static interface service tcp 80 80
!
object-group network R1_VLAN19_REAL
 network-object host 136.1.19.1
!
object service R1_TCP_23
 service tcp destination eq 23
!
object service ASA2_TCP_2023
 service tcp destination eq 2023
!
nat (VLAN29,VLAN19) source static any any destination static interface R1_VLAN19_REAL service ASA2_TCP_2023 R1_TCP_2
!
access-list VLAN29_INBOUND permit tcp any host 136.1.19.1 eq 23
access-group VLAN29_INBOUND in interface VLAN29
```

Verification

Check the NAT configuration on ASA1, and notice that it is interface oriented; verify the flags for static and portmap NAT:

```
ASA1# show nat

NAT policies on Interface VLAN49:
  match ip VLAN49 host 150.1.1.1 VLAN59 any
    static translation to 136.1.59.101
    translate_hits = 0, untranslate_hits = 1
  match tcp VLAN49 host 136.1.49.1 eq 80 VLAN59 any
    static translation to 136.1.59.15/80
    translate_hits = 0, untranslate_hits = 0

NAT policies on Interface VLAN59:
  match tcp VLAN59 host 136.1.59.2 eq 23 VLAN49 any
    static translation to 136.1.49.15/2023
    translate_hits = 0, untranslate_hits = 0
!
!ASA1# show xlate detail
```

https://t.me/learningnets

```
3 in use, 3 most used
Flags: D - DNS, d - dump, I - identity, i - dynamic, n - no random,
       r - portmap, s - static
NAT from VLAN49:150.1.1.1 to VLAN59:136.1.59.101 flags s
TCP PAT from VLAN49:136.1.49.1/80 to VLAN59:136.1.59.15/80 flags sr
TCP PAT from VLAN59:136.1.59.2/23 to VLAN49:136.1.49.15/2023 flags sr
```

Telnet from R2 toward 136.1.59.101 to verify the static NAT:

```
R2#telnet 136.1.59.101
Trying 136.1.59.101 ... Open


User Access Verification

Password: R1>
!
!ASA1# show conn address 136.1.59.2
1 in use, 14 most used
TCP VLAN49 136.1.59.101(150.1.1.1):23 VLAN59 136.1.59.2:36418, idle 0:02:02, bytes 266, flags UIO
```

Telnet from R2 toward ASA1 on port 80 to verify the static PAT:

```
R2#telnet 136.1.59.15 80
Trying 136.1.59.15, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Date: Sun, 10 Nov 2013 07:27:43 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request[Connection to 136.1.59.15 closed by foreign host]
!
!ASA1# show conn address 136.1.59.2
1 in use, 14 most used
TCP VLAN49 136.1.59.15(136.1.49.1):80 VLAN59 136.1.59.2:44690, idle 0:00:37, bytes 0, flags U
```

Telnet from R1 toward ASA1 on port 2023 to verify the static PAT:

```
R1#telnet 136.1.49.15 2023
Trying 136.1.49.15, 2023 ... Open
```

https://t.me/learningnets

```
User Access Verification

Password: R2>
!
!ASA1# show conn address 136.1.49.1 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:136.1.49.1/50180 (136.1.49.1/50180) VLAN59:136.1.59.2/23 (136.1.49.15/2023), flags UIOB, idle 1m52s, upti
```

Check the NAT configuration on ASA2, and notice that it is not interface oriented; verify the flags for object-nat and twice-nat:

```
ASA2# show nat
Manual NAT Policies (Section 1)
1 (VLAN29) to (VLAN19) source static any any   destination static interface R1_VLAN19_REAL service ASA2_TCP_2023 R1_
    translate_hits = 2, untranslate_hits = 2

Auto NAT Policies (Section 2)
1 (VLAN29) to (VLAN19) source static R2_VLAN29_REAL interface   service tcp www www
    translate_hits = 0, untranslate_hits = 1
2 (VLAN29) to (any) source static R2_LO1_REAL 136.1.19.101
    translate_hits = 0, untranslate_hits = 2
!
!ASA2# show xlate

4 in use, 6 most used
Flags: D - DNS, i - dynamic, r - portmap, s - static, I - identity, T - twice
       e - extended
TCP PAT from VLAN29:136.1.29.2 80-80 to VLAN19:136.1.19.16 80-80
    flags sr idle 0:13:46 timeout 0:00:00
NAT from VLAN29:150.1.22.22 to any:136.1.19.101
    flags s idle 0:17:41 timeout 0:00:00
```

https://t.me/learningnets

```
TCP PAT from VLAN29:0.0.0.0/0 0 to VLAN19:0.0.0.0/0 0
    flags srIT idle 0:00:20 timeout 0:00:00
TCP PAT from VLAN19:136.1.19.1 23-23 to VLAN29:136.1.29.16 2023-2023
    flags srT idle 0:00:20 timeout 0:00:00
```

Telnet from both R1 and **TEST-SRV-A** toward 136.1.19.101 to verify the static NAT:

```
R1#telnet 136.1.19.101
Trying 136.1.19.101 ... Open

User Access Verification

Password: R2>
!
!ASA2# show conn
2 in use, 53 most used
TCP VLAN29 136.1.19.101(150.1.22.22):23 VLAN19 136.1.19.1:60161, idle 0:00:02, bytes 102, flags UIO
TCP VLAN29 136.1.19.101(150.1.22.22):23 VLAN26 172.16.10.100:59912, idle 0:00:44, bytes 101, flags UIO
```

Telnet from R1 to ASA2's VLAN 19 port 80 to verify the static PAT:

```
R1#telnet 136.1.19.16 80
Trying 136.1.19.16, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Date: Sun, 10 Nov 2013 11:31:57 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request[Connection to 136.1.19.16 closed by foreign host]
!
!ASA2# show conn address 136.1.19.1
1 in use, 53 most used
TCP VLAN29 136.1.19.16(136.1.29.2):80 VLAN19 136.1.19.1:23330, idle 0:00:04, bytes 0, flags U
```

Telnet from R2 to ASA2's VLAN 29 port 2023 to verify the static PAT:

```
R2#telnet 136.1.29.16 2023
Trying 136.1.29.16, 2023 ... Open

User Access Verification
```

https://t.me/learningnets

```
Password: R1>


!
!ASA2# show conn address 136.1.29.2 long
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/19267 (136.1.29.2/19267) VLAN19:136.1.19.1/23 (136.1.29.16/2023), flags UIOB, idle 49s, uptime
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Static Policy NAT and PAT

Task

- Configure static NAT on ASA1:
  - Traffic sourced from R2's Loopback0 and destined to R1's Loopback0 should be translated to 136.1.49.101.
  - Traffic sourced from R2's Loopback0 and destined to R1's VLAN 49 should be translated to 136.1.49.102.
  - Telnet sessions sourced from R1's Loopback0 and destined to ASA1's VLAN 49 should be redirected to R2's Loopback0.
- Configure static NAT on ASA2:
  - Traffic sourced from R1's Loopback1 and destined to R2's Loopback1 should be translated to 136.1.29.101.
  - Traffic sourced from R1's Loopback1 and destined to R2's VLAN 29 should be translated to 136.1.29.102.
  - Telnet sessions sourced from R2's Loopback1 and destined to ASA2's VLAN 29 should be redirected to R1's Loopback1.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.

https://t.me/learningnets

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Static NAT creates a one-to-one mapping between two IP addresses, and static PAT restricts this mapping for specific TCP/UDP services. This can be further enhanced by making the mapping (static NAT or static PAT) available/valid only to specific hosts/subnets. Considering an ASA with two interfaces, **inside** and **outside** , there are two such cases:

- When host/subnet A from **inside** exists on the **outside** interface, it will be statically NAT'ed to host/subnet B, but only if the traffic destination is host/subnet C.
- Host/subnet A from **inside** is reachable on the **outside** interface as host/subnet B, but only if the traffic source is host/subnet D.

Old NAT Syntax

This is implemented on ASA code before 8.3(1) through a regular static command where an access-list is specified to define the restrictions. For example, let's make the inside HTTP server 10.0.0.1 appear as 172.16.1.1 only to the outside hosts on the subnet 192.168.1.0/24. We first create a special ACL that defines the policy: `access-list POLICY permit tcp host 10.0.0.1 eq 80 192.168.1.0 255.255.255.0` . In this access-list, the source matches the inside server characteristics whereas destination matches the outside hosts. We then create the static rule: `static (inside,outside) tcp 172.16.1.1 80 access-list POLICY` . Policy static NAT/PAT takes precedence over regular static NAT/PAT.

New NAT Syntax

Because we need to somehow define the restrictions/rules for which the NAT is valid, object-NAT is not an option because it is too restrictive in configuration options. For defining any kind of policy NAT, only twice-NAT can be used. There is no need to define any access-lists, because the restrictions are edited inline within the NAT rule. For example, let's make the inside HTTP server 10.0.0.1 appear as 172.16.1.1 only to the outside hosts on the subnet 192.168.1.0/24. We need to

https://t.me/learningnets

create objects or object-groups for all hosts/subnets used in the NAT rule first, and then define the rule:

```
object network SOURCE_REAL
 host 10.0.0.1
!
object network SOURCE_MAPPED
 host 172.16.1.1
!
object network DESTINATION
 subnet 192.168.1.0 255.255.255.0
!
nat (inside,outside) source static SOURCE_REAL SOURCE_MAPPED destination static DESTINATION DESTINATION
```

With the above rule, we basically statically NAT the source only if the destination is 192.168.1.0/24 and preserve the destination as is (don't NAT it).

Configuration

ASA1:

```
access-list R1_LO0_NAT permit ip host 150.1.2.2 host 150.1.1.1
access-list R1_VLAN49_NAT permit ip host 150.1.2.2 host 136.1.49.1
static (VLAN59,VLAN49) 136.1.49.101 access-list R1_LO0_NAT
static (VLAN59,VLAN49) 136.1.49.102 access-list R1_VLAN49_NAT
!
access-list R2_LO0_NAT permit tcp host 150.1.2.2 eq 23 host 150.1.1.1
static (VLAN59,VLAN49) tcp interface 23 access-list R2_LO0_NAT
!
access-list VLAN49_INBOUND permit tcp host 150.1.1.1 host 136.1.49.15 eq 23
access-group VLAN49_INBOUND in interface VLAN49
```

ASA2:

```
object network R1_L1_REAL
 host 150.1.11.11
!
object network R1_L1_R2_L1_MAPPED
 host 136.1.29.101
!
object network R2_L1_REAL
 host 150.1.22.22
!
nat (VLAN19,VLAN29) source static R1_L1_REAL R1_L1_R2_L1_MAPPED destination static R2_L1_REAL R2_L1_REAL
```

https://t.me/learningnets

```
!
!
object-group network R1_L1_R2_VLAN29_MAPPED
 network-object host 136.1.29.102
!
object-group network R2_VLAN29_REAL
 network-object host 136.1.29.2
!
nat (VLAN19,VLAN29) source static R1_L1_REAL R1_L1_R2_VLAN29_MAPPED destination static R2_VLAN29_REAL R2_VLAN29_REAL
!
!
object service TELNET
 service tcp destination eq 23
!
nat (VLAN29,VLAN19) source static R2_L1_REAL R2_L1_REAL destination static interface R1_L1_REAL service TELNET TELNE
!
access-list VLAN29_INBOUND permit tcp host 150.1.22.22 host 150.1.11.11 eq 23
access-group VLAN29_INBOUND in interface VLAN29
```

Verification

Initiate traffic from R2's Loopback0 toward R1's Loopback0, and verify NAT:

https://t.me/learningnets

```
R2#telnet 150.1.1.1 /source-interface loopback0

Trying 150.1.1.1 ... Open

User Access Verification

Password: R1>

!
!ASA1# show conn address 150.1.2.2 long

1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:150.1.1.1/23 (150.1.1.1/23) VLAN59:150.1.2.2/42785 (136.1.49.101/42785), flags UIO, idle 20s, uptime 25s,
```

Initiate traffic from R2's Loopback0 toward R1's VLAN 49, and verify NAT:

```
R2#telnet 136.1.49.1 /source-interface loopback0
Trying 136.1.49.1 ... Open

User Access Verification

Password: R1>
!
!ASA1# show conn address 150.1.2.2 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
```

https://t.me/learningnets

```
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:136.1.49.1/23 (136.1.49.1/23) VLAN59:150.1.2.2/63949 (136.1.49.102/63949), flags UIO, idle 3s, uptime 5s,
```

Initiate telnet traffic from R1's Loopback0 toward ASA1's VLAN 49, and verify NAT:

```
R1#telnet 136.1.49.15  /source-interface loopback0
Trying 136.1.49.15 ... Open


User Access Verification


Password: R2>
!
!ASA1# show conn address 150.1.1.1 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:150.1.1.1/53957 (150.1.1.1/53957) VLAN59:150.1.2.2/23 (136.1.49.15/23), flags UIOB, idle 17s, uptime 19s,
```

Initiate traffic from R1's Loopback1 to R2's Loopback1, and verify NAT:

```
R1#telnet 150.1.22.22 /source-interface loopback1
Trying 150.1.22.22 ... Open


User Access Verification


Password: R2>
!
!ASA2# show conn address 150.1.11.11 long
```

https://t.me/learningnets

```
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:150.1.22.22/23 (150.1.22.22/23) VLAN19:150.1.11.11/23596 (136.1.29.101/23596), flags UIO, idle 30s, uptim
```

Initiate traffic from R1's Loopback1 to R2's VLAN 29, and verify NAT:

```
R1#telnet 136.1.29.2 /source-interface loopback1
Trying 136.1.29.2 ... Open


User Access Verification


Password: R2>
!
!ASA2#show conn address 150.1.11.11 long
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/23 (136.1.29.2/23) VLAN19:150.1.11.11/30959 (136.1.29.102/30959), flags UIO, idle 21s, uptime
```

Initiate traffic from R2's Loopback1 to ASA2's VLAN 29, and verify NAT:

https://t.me/learningnets

```
R2#telnet 136.1.29.16 /source-interface loopback1

Trying 136.1.29.16 ... Open

User Access Verification

Password: R1>
!
!ASA2# show conn address 150.1.22.22 long
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:150.1.22.22/65245 (150.1.22.22/65245) VLAN19:150.1.11.11/23 (136.1.29.16/23), flags UIOB, idle 39s, uptim
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Dynamic NAT and PAT

Task

- Configure dynamic NAT on ASA1:
  - Traffic from VLAN 59 should be translated to an address pool of 136.1.49.120-136.1.49.130; use the last IP in the range as a PAT backup.
  - Traffic from R2's Loopback0 should be translated into ASA's interface IP address.
- Configure dynamic NAT on ASA2:
  - Traffic from VLAN 19 should be translated to an address pool of 136.1.29.120-136.1.29.130; use ASA's interface as a PAT backup.
  - Traffic from R1's Loopback1 should be translated into a PAT pool of 136.1.29.140-136.1.29.145; ensure equal use of IP addresses from the pool.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

https://t.me/learningnets

In the most basic form, the ASA firewall supports two types of dynamic address translations (applied to the source IP addresses, of course):

- NAT (network address translation), where every inside host is dynamically allocated a new outside IP from a configured pool.
- PAT (port address translation), where hosts matching the PAT rule have their addresses translated to a single IP address, with the TCP/UDP ports being overloaded and rewritten as needed.

Old NAT Syntax

To configure dynamic NAT or PAT rules, you must identify the real source that will be translated using a `nat` command and associate it with the mapped source to which translation is performed using a `global` command:

- The syntax for nat rule is `nat <nameif> <nr> <subnet> <mask>` ; the **nameif** identifies the interface on which the packets are received before being translated, and the **nr** is the identifier that binds together a nat rule with a global rule.
- The syntax for global rule is `global <nameif> <nr> [<Start_IP_Address> <End_IP_Address>|interface] netmask <mask>` ; the **nameif** identifies the interface packets exiting the firewall, and the **nr** is the identifier that needs to match between the nat and global rules.

Depending on the global rule configuration, traffic is subject to dynamic NAT or dynamic PAT:

- If a single IP address or the `interface` keyword is configured in the global rule, all sources will be port address translated (PAT), or overloaded.
- If a range of IP addresses is configured in the global rule, each source will be dynamically allocated an IP address from the range.
- If both of the above are configured for the same sequence number of the global rule, each source will be dynamically allocated an IP address from the range until the pool is exhausted at which point all sources will be port address translated (PAT).

For a packet to be dynamically translated, the following conditions must be satisfied:

- It must be received by the ASA on an interface with a nat rule applied and match the nat rule condition.
- It must be routed by the ASA through an interface with a global rule applied.
- The sequence number from the matched nat rule must be the same as the sequence number for the global rule; if the packet matches on a NAT rule but does not match

https://t.me/learningnets

on a global rule, the packet is dropped.

New NAT Syntax

With the new NAT syntax, the **`global`** command has been removed and everything is configured using the **`nat`** command through object-nat or twice-nat. Although both object-nat and twice-nat can be used to configure dynamic NAT or dynamic PAT, from a practical point of view it makes more sense to use object-nat. To configure dynamic object-nat, follow these steps:

1. Create a network object using the command **`object network <name> .`**
2. Configure a single entry within the object, defining the real IP address, real subnet, or IP address range using one of these commands: **`host <IP Address>`** , **`subnet <network> <mask>`** , **`range <start_ip> <end_ip> .`**
3. Configure dynamic object-nat using the object-level command **`nat ([<real_interface>|any] [<mapped_interface>|any]) dynamic [<mapped_IP>|interface|pat-pool <object>]`** . The **`any`** keyword can be used for source and/or destination interface, which makes the NAT rule valid regardless on which interface traffic is coming from or going to; **`any`** keyword cannot be used along with the **`interface`** keyword.

Because the dynamic NAT scope has been already discussed, let's see some examples of mapping the old NAT to the new NAT:

```
! Old NAT Syntax
!
nat (inside) 100 180.10.0.0 255.255.255.0
global (outside) 100 190.1.1.1
!
!New Object NAT Syntax
!
object network MAPPED_SOURCE
 host 190.1.1.1
!
object network REAL_SOURCE
 subnet 180.10.0.0 255.255.255.0
 nat (inside,outside) dynamic pat-pool MAPPED_SOURCE
!
!New Twice NAT Syntax
!
object network REAL_SOURCE
 subnet 180.10.0.0 255.255.255.0
!
```

https://t.me/learningnets

```
object-group network MAPPED_SOURCE
 network-object host 190.1.1.1
!
nat (inside,outside) source dynamic REAL_SOURCE pat-pool MAPPED_SOURCE
```

Configuration

ASA1:

```
nat (VLAN59) 10 136.1.59.0 255.255.255.0
global (VLAN49) 10 136.1.49.120-136.1.49.129
global (VLAN49) 10 136.1.49.130
!
nat (VLAN59) 20 150.1.2.2 255.255.255.255
global (VLAN49) 20 interface
```

ASA2:

```
object network DYNAMIC_POOL
 range 136.1.29.120 136.1.29.130
!
object network VLAN19_REAL_TO_VLAN26
 subnet 136.1.19.0 255.255.255.0
 nat (VLAN19,VLAN26) dynamic DYNAMIC_POOL interface
!
object network PAT_POOL
 range 136.1.29.140 136.1.29.145
!
object-group network R1_LO0_REAL
 network-object host 150.1.11.11
!
nat (VLAN19,any) source dynamic R1_LO0_REAL pat-pool PAT_POOL round-robin
```

Verification

Verify the NAT configuration on ASA1:

```
ASA1# show nat


NAT policies on Interface VLAN59:
  match ip VLAN59 host 150.1.2.2 VLAN49 any
    dynamic translation to pool 20 (136.1.49.15 [Interface PAT])
```

https://t.me/learningnets

```
    translate_hits = 2, untranslate_hits = 0
  match ip VLAN59 host 150.1.2.2 VLAN59 any
    dynamic translation to pool 20 (No matching global)
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 136.1.59.0 255.255.255.0 VLAN49 any
    dynamic translation to pool 10 (136.1.49.120 - 136.1.49.129)
    translate_hits = 3, untranslate_hits = 0
  match ip VLAN59 136.1.59.0 255.255.255.0 VLAN59 any
    dynamic translation to pool 10 (No matching global)
    translate_hits = 0, untranslate_hits = 0
```

Initiate traffic from VLAN 59 subnet, and verify NAT:

```
R2#telnet 150.1.1.1
Trying 150.1.1.1 ... Open



User Access Verification

Password: R1>
!
!ASA1# show conn address 136.1.59.2 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:150.1.1.1/23 (150.1.1.1/23) VLAN59:136.1.59.2/15924 (136.1.49.127/15924), flags UIO, idle 21s, uptime 24s
```

Initiate traffic from R2's Loopback0 subnet, and verify NAT:

```
R2#telnet 150.1.1.1 /source-interface loopback0
Trying 150.1.1.1 ... Open


User Access Verification
```

https://t.me/learningnets

```
Password: R1>

!
!ASA1# show conn address 150.1.2.2 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:150.1.1.1/23 (150.1.1.1/23) VLAN59:150.1.2.2/25677 (136.1.49.15/26211), flags UIO, idle 27s, uptime 29s,
```

Verify the NAT configuration on ASA2:

```
ASA2# show nat

Manual NAT Policies (Section 1)
1 (VLAN19) to (any) source dynamic R1_LO0_REAL pat-pool PAT_POOL round-robin
    translate_hits = 0, untranslate_hits = 0

Auto NAT Policies (Section 2)
1 (VLAN19) to (VLAN26) source dynamic VLAN19_REAL_TO_VLAN26 DYNAMIC_POOL interface
    translate_hits = 0, untranslate_hits = 0
2 (VLAN19) to (VLAN29) source dynamic VLAN19_REAL_TO_VLAN29 DYNAMIC_POOL interface
    translate_hits = 0, untranslate_hits = 0
```

Initiate traffic from VLAN 19 subnet, and verify NAT:

```
R1#telnet 136.1.29.2
Trying 136.1.29.2 ... Open

User Access Verification

Password: R2>
!
!ASA2# show conn address 136.1.19.1 long
```

https://t.me/learningnets

```
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/23 (136.1.29.2/23) VLAN19:136.1.19.1/15454 (136.1.29.124/15454), flags UIO, idle 46s, uptime 4
!
!R1#telnet 172.16.10.100 80
Trying 172.16.10.100, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Content-Type: text/html; charset=UTF-8
Server: Microsoft-HTTPAPI/2.0
Date: Sun, 10 Nov 2013 20:12:45 GMT
Connection: close
Content-Length: 326

<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN""http://www.w3.org/TR/html4/strict.dtd">
<HTML><HEAD><TITLE>Bad Request</TITLE>
<META HTTP-EQUIV="Content-Type" Content="text/html; charset=UTF-8"></HEAD>
<BODY><h2>Bad Request - Invalid Verb</h2>
<hr><p>HTTP Error 400. The request verb is invalid.</p>
</BODY></HTML>[Connection to 172.16.10.100 closed by foreign host]
!
!ASA2# show conn address 136.1.19.1 long
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
```

https://t.me/learningnets

```
       X - inspected by service module
TCP VLAN26:172.16.10.100/80 (172.16.10.100/80) VLAN19:136.1.19.1/31531 (136.1.29.128/31531), flags UO, idle 6s, upti
```

Initiate traffic from R1's Loopback1 subnet:

```
R1#telnet 136.1.29.2 /source-interface loopback1
Trying 136.1.29.2 ... Open

User Access Verification

Password: R2>
!
!R1#telnet 172.16.10.100 /source-interface loopback1
Trying 172.16.10.100 ... Open

Welcome to Microsoft Telnet Service
login:
!
!ASA2# show conn address 150.1.11.11 long
2 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/23 (136.1.29.2/23) VLAN19:150.1.11.11/36395 (136.1.29.142/36395), flags UIO, idle 15s, uptime
TCP VLAN26:172.16.10.100/23 (172.16.10.100/23) VLAN19:150.1.11.11/61502 (136.1.29.142/61502), flags U, idle 4s, upti
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Dynamic Policy NAT and PAT

Task

- Configure dynamic NAT on ASA1:
  - Telnet traffic from VLAN 59 should be translated to 136.1.49.155.
  - ICMP traffic from VLAN 59 should be translated to 136.1.49.165.
  - All other traffic should be translated into ASA's interface IP address.
- Configure dynamic NAT on ASA2:
  - Telnet traffic from VLAN 19 should be translated to 136.1.29.155.
  - HTTP traffic from VLAN 19 should be translated to 136.1.29.165.
  - All traffic from R1's Loopback1 should be translated into ASA'2 interface IP address.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

This is similar to the policy static NAT and PAT; in this case the scope is to dynamically NAT/PAT the traffic source based on the traffic pattern. For example,

https://t.me/learningnets

traffic from source A is dynamically NAT'ed to B for telnet traffic toward destination C, whereas traffic from source A is dynamically NAT'ed to D for HTTP traffic toward destination E. This is implemented differently, based on the ASA code version.

Old NAT Syntax

We still use regular **`nat`** and **`global`** commands, but the **`nat`** command is changed—we no longer reference a subnet/host but an access-list matching the traffic policy (the ACL can be standard or extended). So from a configuration perspective, there are three steps involved:

1. Identify the traffic within an ACL; for example, **`access-list POLICY1 permit http 10.10.10.0 255.255.255.0 172.16.10.0 255.255.25.0 .`**
2. Configure the nat rule for traffic identified by the ACL if coming on the inside interface, **`nat (inside) 20 access-list POLICY1 .`**
3. Configure the global rule if traffic is exiting the outside interface, the same as regular dynamic NAT/PAT, **`global (outside) 20 101.101.101.101 .`**

Dynamic policy NAT/PAT takes precedence over regular dynamic NAT/PAT. The configuration order is important; the rules for the most specific traffic should be configured first, or you won't match the required policy. ASA cannot automatically order nat/global rules for dynamic policy NAT/PAT based on the sequence numbers.

New NAT Syntax

Because we need to somehow define the restrictions/rules for which the NAT is valid, object-NAT is not an option because it is too restrictive in configuration options. For defining any kind of policy NAT, only twice-NAT can be used. There is no need to define any access-lists, because the restrictions are edited inline within the NAT rule. For example, let’s make the HTTP traffic from inside host 10.0.0.1 appear as 172.16.1.1 only if traffic is destined to outside hosts on the subnet 192.168.1.0/24. We must create objects or object-groups for all hosts/subnets used in the NAT rule first, and then define the rule:

```
object network SOURCE_REAL
 host 10.0.0.1
!
object network SOURCE_MAPPED
 host 172.16.1.1
!
object network DESTINATION
 subnet 192.168.1.0 255.255.255.0
!
object service HTTP
 service tcp destinatiomn eq 80
```

https://t.me/learningnets

```
!
nat (inside,outside) source dynamic SOURCE_REAL SOURCE_MAPPED destination static DESTINATION DESTINATION service HTT
```

As a restriction, at this moment twice-NAT does not allow policy dynamic NAT/PAT for ICMP traffic; if you attach a service defining ICMP traffic on the rule, you will receive an error message similar to **`ERROR: real service object includes protocol that doesnt match TCP or UDP. .`**

Configuration

ASA1:

```
access-list TELNET permit tcp 136.1.59.0 255.255.255.0 any eq 23
access-list ICMP permit icmp 136.1.59.0 255.255.255.0 any
access-list ALL permit ip any any
!
nat (VLAN59) 10 access-list TELNET
nat (VLAN59) 20 access-list ICMP
nat (VLAN59) 30 access-list ALL
!
global (VLAN49) 10 136.1.49.155
global (VLAN49) 20 136.1.49.165
global (VLAN49) 30 interface
```

ASA2:

```
object network VLAN19_REAL
 subnet 136.1.19.0 255.255.255.0
!
object network R1_LO1_REAL
 host 150.1.11.11
!
object network VLAN19_TELNET_MAPPED
 host 136.1.29.155
!
object network VLAN19_HTTP_MAPPED
 host 136.1.29.165
!
object network ANY_DESTINATION
 subnet 0.0.0.0 0.0.0.0
!
object service TELNET
 service tcp destination eq 23
!
```

https://t.me/learningnets

```
object service HTTP
 service tcp destination eq 80
!
nat (VLAN19,any) source dynamic VLAN19_REAL VLAN19_TELNET_MAPPED destination static ANY_DESTINATION ANY_DESTINATION
nat (VLAN19,any) source dynamic VLAN19_REAL VLAN19_HTTP_MAPPED destination static ANY_DESTINATION ANY_DESTINATION se
nat (VLAN19,VLAN26) source dynamic R1_LO1_REAL interface
nat (VLAN19,VLAN29) source dynamic R1_LO1_REAL interface
```

Verification

Verify the NAT configuration on ASA1:

```
ASA1# show nat


NAT policies on Interface VLAN59:
  match tcp VLAN59 136.1.59.0 255.255.255.0 VLAN49 any eq 23
    dynamic translation to pool 10 (136.1.49.155)
    translate_hits = 0, untranslate_hits = 0
  match tcp VLAN59 136.1.59.0 255.255.255.0 VLAN59 any eq 23
    dynamic translation to pool 10 (No matching global)
    translate_hits = 0, untranslate_hits = 0
  match icmp VLAN59 136.1.59.0 255.255.255.0 VLAN49 any
    dynamic translation to pool 20 (136.1.49.165)
    translate_hits = 0, untranslate_hits = 0
  match icmp VLAN59 136.1.59.0 255.255.255.0 VLAN59 any
    dynamic translation to pool 20 (No matching global)
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 any VLAN49 any
    dynamic translation to pool 30 (136.1.49.15 [Interface PAT])
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 any VLAN59 any
    dynamic translation to pool 30 (No matching global)
    translate_hits = 0, untranslate_hits = 0
```

Initiate telnet traffic from VLAN 59, and verify NAT:

```
R2#telnet 136.1.49.1
Trying 136.1.49.1 ... Open


User Access Verification


Password: R1>
```

https://t.me/learningnets

```
!
!ASA1# show conn address 136.1.59.2 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:136.1.49.1/23 (136.1.49.1/23) VLAN59:136.1.59.2/46663 (136.1.49.155/30621), flags UIO, idle 9s, uptime 11
```

Initiate ICMP traffic from VLAN 59, and verify NAT:

https://t.me/learningnets

```
R2#ping 136.1.49.1 repeat 100

Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 136.1.49.1, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 1/1/4 ms
!
!ASA1# show conn address 136.1.59.2 long

1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
ICMP VLAN49:136.1.49.1/0 (136.1.49.1/0) VLAN59:136.1.59.2/17 (136.1.49.165/4558), idle 0s, uptime 0s, timeout 2s, by
```

Initiate other traffic from VLAN 59, like HTTP, and verify NAT:

```
R2#telnet 136.1.49.1 80
Trying 136.1.49.1, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Date: Sun, 10 Nov 2013 21:38:26 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request[Connection to 136.1.49.1 closed by foreign host]
!
!ASA1# show conn address 136.1.59.2 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
```

https://t.me/learningnets

```
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:136.1.49.1/80 (136.1.49.1/80) VLAN59:136.1.59.2/51619 (136.1.49.15/29450), flags U, idle 4s, uptime 4s, t
```

Verify the NAT configuration on ASA2:

```
ASA2# show nat

Manual NAT Policies (Section 1)
1 (VLAN19) to (any) source dynamic VLAN19_REAL VLAN19_TELNET_MAPPED   destination static ANY_DESTINATION ANY_DESTINA
    translate_hits = 3, untranslate_hits = 3
2 (VLAN19) to (any) source dynamic VLAN19_REAL VLAN19_HTTP_MAPPED   destination static ANY_DESTINATION ANY_DESTINATI
    translate_hits = 3, untranslate_hits = 3
3 (VLAN19) to (VLAN26) source dynamic R1_LO1_REAL interface
    translate_hits = 0, untranslate_hits = 0
4 (VLAN19) to (VLAN29) source dynamic R1_LO1_REAL interface
    translate_hits = 1, untranslate_hits = 0
```

Initiate telnet traffic from VLAN 19, verify NAT:

```
R1#telnet 136.1.29.2
Trying 136.1.29.2 ... Open


User Access Verification

Password: R2>
!
!R1#telnet 172.16.10.100
Trying 172.16.10.100 ... Open

Welcome to Microsoft Telnet Service
login:
!
!ASA2# show conn address 136.1.19.1 long
1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
```

https://t.me/learningnets

```
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/23 (136.1.29.2/23) VLAN19:136.1.19.1/30880 (136.1.29.155/30880), flags UIO, idle 16s, uptime 1
TCP VLAN26:172.16.10.100/23 (172.16.10.100/23) VLAN19:136.1.19.1/55966 (136.1.29.155/55966), flags UO, idle 5s, upti
```

Initiate HTTP traffic from VLAN 19, and verify NAT:

```
R1#telnet 136.1.29.2 80
Trying 136.1.29.2, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Date: Sun, 10 Nov 2013 21:46:38 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request[Connection to 136.1.29.2 closed by foreign host]
!
!R1#telnet 172.16.10.100 80
Trying 172.16.10.100, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Content-Type: text/html; charset=us-ascii
Server: Microsoft-HTTPAPI/2.0
Date: Sun, 10 Nov 2013 21:54:21 GMT
Connection: close
Content-Length: 326

<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN""http://www.w3.org/TR/html4/strict.dtd">
<HTML><HEAD><TITLE>Bad Request</TITLE>
<META HTTP-EQUIV="Content-Type" Content="text/html; charset=us-ascii"></HEAD>
<BODY><h2>Bad Request - Invalid Verb</h2>
<hr><p>HTTP Error 400. The request verb is invalid.</p>
</BODY></HTML>[Connection to 172.16.10.100 closed by foreign host]
!
!ASA2# show conn address 136.1.19.1 long
```

https://t.me/learningnets

```
3 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/80 (136.1.29.2/80) VLAN19:136.1.19.1/39475 (136.1.29.165/39475), flags U, idle 10s, uptime 27s
TCP VLAN26:172.16.10.100/80 (172.16.10.100/80) VLAN19:136.1.19.1/27428 (136.1.29.165/27428), flags U, idle 3s, uptim
```

Initiate traffic from R1's Loopback1, and verify NAT:

https://t.me/learningnets

```
R1#telnet 136.1.29.2 /source-interface loopback1

Trying 136.1.29.2 ... Open



User Access Verification

Password: R2>

!
!ASA2# show conn address 150.1.11.11 long

1 in use, 53 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN29:136.1.29.2/23 (136.1.29.2/23) VLAN19:150.1.11.11/57329 (136.1.29.16/57329), flags UIO, idle 11s, uptime 1
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Outside Dynamic NAT

Task

- Configure dynamic NAT on ASA1:
  - Traffic from VLAN 49 should be translated to 136.1.59.166.
  - Traffic from R1's Loopback0 toward R2's Loopback0 should be translated to 136.1.59.177.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Outside Dynamic NAT is a feature that you would rarely need in the real world. Technically, the end result is the same as in regular dynamic NAT or dynamic policy NAT, but it applies translations to the source IP addresses as traffic is originated from the lower security-level interface to the higher security-level interface (or more commonly to the outside hosts accessing the inside destinations). You might occasionally need this in situations where your inside servers have no information about the outside hosts behind your firewall. For example, they may have their default gateway set to some other host, or you may not have any default gateway

https://t.me/learningnets

set at all. Of course, you may use static NAT rules to simplify routing in this situation, but they might not cover the whole source address space if it is too big (such as hundreds of hosts).

The only configuration difference from the regular dynamic NAT or dynamic policy NAT is that the **nat** rule must have the special keyword **outside** appended to the command. This limitation only applies to ASA code version prior to 8.3(1), so for the old nat style. For traffic to be allowed as it passes through the ASA with NAT configured:

- If traffic matches a NAT rule one-way, it will be translated, and when return traffic is received by the ASA, it is untranslated.
- If traffic does not match a NAT rule one-way, when return traffic is received by the ASA it does not need to match on a NAT rule.

These being said, if asymmetric NAT rules are matched (as a session is initiated it does not match a NAT rule, but return traffic will match a NAT rule), ASA will actually drop all traffic and not even allow the session to be initiated; the reason code for this, visible with the command `show asp drop` , would be **NAT reverse path failed (nat-rpf-failed)**. The most common scenario for the problem is when dynamic NAT/PAT is configured (from inside to outside, or from outside to inside).

Let's take ASA1 with VLAN49 interface having a security-level of 10, and VLAN59 interface with a security-level of 20. If we want to configure dynamic NAT for VLAN59 subnet, using ASA1's interface, the configuration would be:

```
nat (VLAN59) 10 136.1.59.0 255.255.255.0
global (VLAN49) 10 interface
```

All traffic from VLAN 59 will be dynamically PAT'ed to the ASA1 VLAN49 IP address as it exits the respective interface. The implicit side-effect is that the VLAN 59 subnet will no longer be reachable from hosts/subnets behind ASA1 VLAN49 interface; so in our case, traffic sourced from the VLAN49 subnet or R1's Loopback0 will no longer be functional for destinations in VLAN 59. This is because traffic is asymmetric; as it goes from R1 toward R2 it does not match a NAT rule, but return traffic will match the configured dynamic NAT rule, so ASA actually restricts the traffic. To fix this, we need to configure static NAT rules for any hosts in VLAN 59 that we want/need to be reachable from R1 (VLAN 49 or Loopback0).

However, based on the above configuration, traffic from R1 (VLAN 49 or Loopback0) toward R2's Loopback0 will be functional, because it does not match any NAT rules (as it is initiated, or for return).

https://t.me/learningnets

The same problem would occur for outside dynamic NAT, so with our task requirements and the below configuration, R2 would no longer be able to access VLAN 49 resources unless static NAT were configured.

Configuration

```
ASA1:

nat (VLAN49) 10 136.1.49.0 255.255.255.0 outside
global (VLAN59) 10 136.1.59.166
!
access-list R1_L0_R2_L0 permit ip host 150.1.1.1 host 150.1.2.2
nat (VLAN49) 20 access-list R1_L0_R2_L0 outside
global (VLAN59) 20 136.1.59.177
!
access-list VLAN49_INBOUND permit ip 136.1.49.0 255.255.255.0 any
access-list VLAN49_INBOUND permit ip host 150.1.1.1 host 150.1.2.2
access-group VLAN49_INBOUND in interface VLAN49
```

Verification

Verify NAT configuration on ASA1:

```
ASA1# show nat


NAT policies on Interface VLAN49:
  match ip VLAN49 host 150.1.1.1 VLAN49 host 150.1.2.2
    dynamic translation to pool 20 (No matching global)
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN49 host 150.1.1.1 VLAN59 host 150.1.2.2
    dynamic translation to pool 20 (136.1.59.177)
    translate_hits = 1, untranslate_hits = 0
  match ip VLAN49 136.1.49.0 255.255.255.0 VLAN49 any
    dynamic translation to pool 10 (No matching global)
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN49 136.1.49.0 255.255.255.0 VLAN59 any
    dynamic translation to pool 10 (136.1.59.166)
    translate_hits = 1, untranslate_hits = 0
```

https://t.me/learningnets

Initiate traffic from VLAN 49, and verify NAT:

```
R1#telnet 136.1.59.2
  Trying 136.1.59.2 ... Open


  User Access Verification

  Password:  R2>
  !
  !ASA1# show conn address 136.1.49.1 long
  1 in use, 14 most used
  Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
         B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
         D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
         G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
         i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
         k - Skinny media, M - SMTP data, m - SIP media, n - GUP
         O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
         q - SQL*Net data, R - outside acknowledged FIN,
         R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
         s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
         V - VPN orphan, W - WAAS,
         X - inspected by service module
TCP VLAN49:136.1.49.1/30917 (136.1.59.166/56652) VLAN59:136.1.59.2/23 (136.1.59.2/23), flags UIOB, idle 12s, uptime
```

Generate traffic from R1's Loopback0 to R2's Loopback0, and verify NAT:

```
R1#telnet 150.1.2.2 /source-interface loopback0
Trying 150.1.2.2 ... Open


User Access Verification


Password: R2>
!
!ASA1# show conn address 150.1.1.1 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
```

https://t.me/learningnets

```
        q - SQL*Net data, R - outside acknowledged FIN,
        R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
        s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
        V - VPN orphan, W - WAAS,
        X - inspected by service module
TCP VLAN49:150.1.1.1/17573 (136.1.59.177/12824) VLAN59:150.1.2.2/23 (150.1.2.2/23), flags UIOB, idle 20s, uptime 24s
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA NAT Control

Task

- Configure ASA1 so that any traffic passing through requires a NAT translation; however:
  - R2's Loopback0 should be reachable by R1 through 150.1.2.2.
  - Traffic from VLAN 59 destined to R1's Loopback0 should not require a translation.
  - All other traffic from VLAN 59 should be translated with the original IP address preserved.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Until version 7.x of the ASA code, configuring NAT translation rules was necessary to make the PIX appliance pass traffic through. Even if you had inside network advertised via IGP to the outside hosts, you had to configure some sort of NAT to make everything work. Since version 7.x, the default behavior does not require a NAT translation entry prior to permitting a session through the firewall (this is now

https://t.me/learningnets

done based on the security-levels and access-lists). In ASA code prior to 8.3(1), you can revert to the old behavior by issuing the global command `nat-control` ; this command has been completely removed with the new NAT syntax, starting with code 8.3(1), and there is no replacement for it because it is considered a legacy feature. With nat-control enabled, any traffic requires matching on a NAT rule, but there are a few mechanisms for bypassing this requirement:

NAT Exempt

You might want certain traffic to be allowed through the firewall (such as traffic from inside hosts toward VPN clients for VPN sessions terminated on the outside interface of the ASA) without being NAT'ed. For this reason there is something called NAT Exempt, where you basically identify the traffic to be exempted from NAT within an access-lists (so traffic will be allowed untranslated through the firewall) and configure a special nat rule for it. The required steps are:

1. Identify the traffic in an ACL; however, only IP statements are allowed, and you can't specify TCP/UDP ports or other IP protocols within an ACL used for NAT Exempt, so `permit ip host 1.1.1.1 host 2.2.2.2` is valid, whereas `permit tcp host 1.1.1.1 host 2.2.2.2` is not valid.
2. Configure a NAT rule using the special NAT identifier of zero and applying the configured ACL (just like in policy dynamic NAT); so `nat (inside) 0 access-list TEST` is a correct NAT Exempt statement, whereas `nat (inside) 11 access-list TEST` is actually a dynamic policy NAT statement.

NAT Exempt does not require any corresponding `global` commands, which means the identified traffic will be exempt from NAT regardless of the exit interface. So traffic matching NAT Exempt rules will not create any xlate/translations on the ASA, it will just be allowed through.

Static Identity NAT

You might want a host from the inside interface to be visible on the outside interface by its real/original IP address, or a host from the outside to be visible on the inside by its real/original IP address. This can be achieved by using something called Static Identity NAT, which is nothing more than a regular static command. However, real and mapped addresses are the same (for example, `static (inside,outside) 10.10.10.10 10.10.10.10`); as a static NAT, it allows the session to be initiated from either side, and it also creates a permament xlate/translation.

https://t.me/learningnets

Dynamic Identity NAT

Because nat-control is enabled, you might want traffic from certain sources to be allowed through the firewall without having the IP address changed/translated, but you don't want to use static Identity NAT because one of the following:

- There are too many sources to create a static entry for each.
- You want any new sources to be allowed without having to add new static commands.
- You don't want traffic to be possibly initiated from either side.

Dynamic Identity NAT behaves like a regular dynamic NAT, except that when translation is performed, the real/original IP address is preserved; configuration requires only a regular nat command but using the special NAT identifier of zero (just as NAT Exempt uses, with the exception that traffic is not matched by an ACL, but is edited inline like in any dynamic NAT rule). No global commands are required, which means that traffic from matched sources will be subject to an identity translation (IP source is preserved) regardless of the exit interface. So `nat (inside) 0 10.10.10.0 255.255.255.0` is a correct Dynamic Identity NAT statement, whereas `nat (inside) 11 10.10.10.0 255.255.255.0` is actually a dynamic NAT statement.

> Note that all three above discussed scenarios can be used regardless of whether nat-control is enabled, but using dynamic identity NAT makes no sense unless nat-control is active.

There is one exception: Traffic between interfaces with the same security-level is allowed to pass with nat-control enabled, but after a nat rule is configured all traffic requires nat.

Configuration

ASA1:

```
nat-control
static (VLAN59,VLAN49) 150.1.2.2 150.1.2.2
!
access-list NAT_EXEMPT permit ip 136.1.59.0 255.255.255.0 host 150.1.1.1
nat (VLAN59) 0 access-list NAT_EXEMPT
!
nat (VLAN59) 0 0.0.0.0 0.0.0.0
!
access-list VLAN49_INBOUND permit ip any host 150.1.2.2
```

https://t.me/learningnets

```
access-group VLAN49_INBOUND in interface VLAN49
```

Verification

Verify NAT configuration on ASA1:

```
ASA1# show nat

NAT policies on Interface VLAN59:
  match ip VLAN59 136.1.59.0 255.255.255.0 VLAN49 host 150.1.1.1
    NAT exempt
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 136.1.59.0 255.255.255.0 VLAN59 host 150.1.1.1
    NAT exempt
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 host 150.1.2.2 VLAN49 any
    static translation to 150.1.2.2
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 any VLAN49 any
    identity NAT translation, pool 0
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 any VLAN59 any
    identity NAT translation, pool 0
    translate_hits = 0, untranslate_hits = 0
  match ip VLAN59 any VLAN49 any
    no translation group, implicit deny
    policy_hits = 13
```

Initiate traffic from R1 toward R2's Loopback0 address; evidence that static identity NAT is matched is if the traffic is functional. Alternatively, you could use a packet-tracer utility to verify it:

```
R1#telnet 150.1.2.2
Trying 150.1.2.2 ... Open

User Access Verification

Password: R2>
!
!ASA1# show conn address 136.1.49.1 long
1 in use, 14 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
```

https://t.me/learningnets

```
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN49:136.1.49.1/59698 (136.1.49.1/59698) VLAN59:150.1.2.2/23 (150.1.2.2/23), flags UIOB, idle 1m1s, uptime 1m4
!
!ASA1# packet-tracer input VLAN49 tcp 136.1.49.1 20000 150.1.2.2 23

Phase: 1
Type: UN-NAT
Subtype: static
Result: ALLOW
Config:
static (VLAN59,VLAN49) 150.1.2.2 150.1.2.2 netmask 255.255.255.255
nat-control
  match ip VLAN59 host 150.1.2.2 VLAN49 any
    static translation to 150.1.2.2
    translate_hits = 0, untranslate_hits = 6
Additional Information:
NAT divert to egress interface VLAN59
Untranslate 150.1.2.2/0 to 150.1.2.2/0 using netmask 255.255.255.255

Phase: 2
Type: ACCESS-LIST
Subtype: log
Result: ALLOW
Config:
access-group VLAN49_INBOUND in interface VLAN49
access-list VLAN49_INBOUND extended permit ip any host 150.1.2.2
Additional Information:

Phase: 3
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:
Phase: 4
Type: NAT
Subtype: rpf-checkResult: ALLOW
```

https://t.me/learningnets

```
Config: static (VLAN59,VLAN49) 150.1.2.2 150.1.2.2 netmask 255.255.255.255
nat-control
 match ip VLAN59 host 150.1.2.2 VLAN49 any
static translation to 150.1.2.2
 translate_hits = 0, untranslate_hits = 6
Additional Information:

Phase: 5
Type: NAT
Subtype: host-limits
Result: ALLOW
Config:
static (VLAN59,VLAN49) 150.1.2.2 150.1.2.2 netmask 255.255.255.255
nat-control
  match ip VLAN59 host 150.1.2.2 VLAN49 any
    static translation to 150.1.2.2
    translate_hits = 0, untranslate_hits = 6
Additional Information:

Phase: 6
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 7
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Config:
Additional Information:
New flow created with id 200, packet dispatched to next module

Result:
input-interface: VLAN49
input-status: up
input-line-status: up
output-interface: VLAN59
output-status: up
output-line-status: upAction: allow
```

Initiate traffic from VLAN 59 toward R1's Loopback0, and verify that traffic is

https://t.me/learningnets

exempted from NAT and that there are no XLATES created but the one for configured static identity NAT. To make sure that NAT Exempt is matched and not dynamic identity NAT, look at the hits counter when using the packet-tracer utility:

```
R2#telnet 150.1.1.1
Trying 150.1.1.1 ... Open

User Access Verification

Password: R1>
!
!ASA1# show xlate
1 in use, 3 most used Global 150.1.2.2 Local 150.1.2.2
!
!ASA1# packet-tracer input VLAN59 tcp 136.1.59.2 30000 150.1.1.1 23

Phase: 1
Type: ROUTE-LOOKUP
Subtype: input
Result: ALLOW
Config:
Additional Information:
in   150.1.1.1       255.255.255.255 VLAN49

Phase: 2
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:
Phase: 3
Type: NAT-EXEMPT
Subtype: Result: ALLOW
Config: nat-control
match ip VLAN59 136.1.59.0 255.255.255.0 VLAN49 host 150.1.1.1
NAT exempt
translate_hits = 2, untranslate_hits = 0
Additional Information:

Phase: 4
Type: NAT
Subtype:
Result: ALLOW
Config:
nat (VLAN59) 0 0.0.0.0 0.0.0.0
nat-control
```

https://t.me/learningnets

```
  match ip VLAN59 any VLAN49 any
    identity NAT translation, pool 0
    translate_hits = 0, untranslate_hits = 0
Additional Information:

Phase: 5
Type: NAT
Subtype: host-limits
Result: ALLOW
Config:
nat (VLAN59) 0 0.0.0.0 0.0.0.0
nat-control
  match ip VLAN59 any VLAN49 any
    identity NAT translation, pool 0
    translate_hits = 0, untranslate_hits = 0
Additional Information:

Phase: 6
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 7
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Config:
Additional Information:
New flow created with id 202, packet dispatched to next module

Result:
input-interface: VLAN59
input-status: up
input-line-status: up
output-interface: VLAN49
output-status: up
output-line-status: upAction: allow
```

Initiate traffic from VLAN 59 to VLAN 49, and verify the dynamic identity NAT:

```
R2#telnet 136.1.49.1
```

https://t.me/learningnets

```
Trying 136.1.49.1 ... Open


User Access Verification

Password: R1>
!
!ASA1# show xlate
2 in use, 3 most used
Global 150.1.2.2 Local 150.1.2.2Global 136.1.59.2 Local 136.1.59.2
!
!ASA1# packet-tracer input VLAN59 tcp 136.1.59.2 40000 136.1.49.1 23

Phase: 1
Type: ROUTE-LOOKUP
Subtype: input
Result: ALLOW
Config:
Additional Information:
in   136.1.49.0      255.255.255.0   VLAN49

Phase: 2
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:
Phase: 3
Type: NAT
Subtype: Result: ALLOW
Config:nat (VLAN59) 0 0.0.0.0 0.0.0.0
nat-control
  match ip VLAN59 any VLAN49 any
identity NAT translation, pool 0
translate_hits = 2, untranslate_hits = 0
Additional Information:Dynamic translate 136.1.59.2/0 to 136.1.59.2/0 using netmask 255.255.255.255

Phase: 4
Type: NAT
Subtype: host-limits
Result: ALLOW
Config:
nat (VLAN59) 0 0.0.0.0 0.0.0.0
nat-control
  match ip VLAN59 any VLAN49 any
    identity NAT translation, pool 0
```

https://t.me/learningnets

```
    translate_hits = 2, untranslate_hits = 0
Additional Information:

Phase: 5
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 6
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Config:
Additional Information:
New flow created with id 204, packet dispatched to next module

Result:
input-interface: VLAN59
input-status: up
input-line-status: up
output-interface: VLAN49
output-status: up
output-line-status: upAction: allow
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Legacy DNS Doctoring

Task

- Configure R1 to act as DNS server. Create a new host entry for the name **www.ine.com** with the IP address 150.1.1.1.
- Hosts VLAN 59 using R1 as their DNS server should see the name **www.ine.com** resolved to the IP address of 136.1.59.1.
- Use the `alias` command on ASA1 to accomplish this.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The legacy `alias` command was designed to work like a half-functioning static command. When you configure the `alias (<nameif>) <mapped_ip> <real_ip> <netmask>` command, ASA will determine whether the `mapped_ip` belongs to the subnet of the **nameif** interface. If it does, the firewall will install an ARP entry to pull traffic to **mapped_ip** through itself. As soon as the packet is received inbound on the **nameif** interface, destined to **mapped_ip**, the firewall will rewrite the destination IP address with the **real_ip** and forward it, effectively implementing the aliasing function. The

https://t.me/learningnets

source IP might be modified as well but only if there are additional NAT rules configured and matched; note that the netmask option allows you to create a whole block of aliases—for example, by setting the netmask 255.255.255.224, you will get about 30 aliases installed by the ASA.

The `alias` command only works one way, so the traffic initiated from **real_ip** will not be source-rewritten to **mapped_ip**. Of course, these days you would implement this redirection using the reverse static NAT command, which works symmetrically. The other important feature of the legacy alias command is automatic DNS rewrite. When a DNS A-record response is forwarded by the firewall outbound on the interface specified in the `alias` command and the IP address in the DNS packet matches the **real_ip**, the firewall will change it to **mapped_ip**. This is called DNS Doctoring becuase the ASA inspects DNS traffic and changes data in DNS payload on the fly; it was the most common use for the `alias` command.

The `alias` command creates a proxy-ARP entry in the firewall by default for the **mapped_ip**, so if you only use it for DNS doctoring, disable Proxy-ARP by using the command `sysopt noproxyarp` to avoid occasional IP address conflict messages. Also, if you're wondering about a way to globally disable DNS rewrites, there are two options:

- Legacy method, using the global `sysopt nodnsalias [inboud|outbound]` command, where inbound means direction from a lower security level to a higher security level interface, and vice-versa for outbound.
- Newest method, using the `no nat-rewrite` command within a DNS inspection policy-map:

```
policy-map type inspect dns DNS_INSPECT
 parameters
  no nat-rewrite
```

> The **`alias`** command is no longer available starting with ASA code 8.3(1).

Configuration

ASA1:

```
alias (VLAN59) 136.1.59.1 150.1.1.1 255.255.255.255
```

R1:

https://t.me/learningnets

```
ip dns server
ip host www.ine.com 150.1.1.1
```

Verification

Configure R2 to use R1 as its DNS server, and test DNS resolution before applying the configuration on ASA1:

```
R2:
ip name-server 136.1.49.1
ip domain-lookup

R2#ping www.ine.com
Translating "www.ine.com"...domain server (136.1.49.1) [OK]

Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 150.1.1.1
, timeout is 2 seconds: !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Apply the configuration on ASA1 and test DNS resolution again:

```
R2#ping www.ine.com
Translating "www.ine.com"...domain server (136.1.49.1) [OK]

Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 136.1.59.1
, timeout is 2 seconds: .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA DNS Doctoring

Task

- Use the `static` command on ASA1 so that hosts VLAN 59 using R1 as their DNS server should see the name **www.ine.com** resolved to the IP address of 136.1.59.1.
  - Configure R1 to act as DNS server. Create a new host entry for the name **www.ine.com** with the IP address 150.1.1.1.
- Configure ASA2 so that hosts VLAN 19 using R2 as their DNS server should see the name **www.inelab.com** resolved to the IP address of 136.1.19.22.
  - Configure R2 to act as DNS server. Create a new host entry for the name **www.inelab.com** with the IP address 150.1.22.22.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

https://t.me/learningnets

Because the `static` command replaced the legacy `alias` command, it was natural to incorporate the DNS rewrite feature into this command. When you specify the dns keyword with the command `static [<local-iface>,<global-iface>] <global-ip> <local-ip> dns` , ASA will inspect the following DNS packets:

- Going from **global-iface** to **local-iface** and having **global-ip** inside the DNS A-record contents; the **global-ip** will be changed to **local-ip**.
- Going from **local-iface** to **global-iface** and having **local-ip**> inside the DNS A-record; the **local-ip** will be replaced with **global-ip**.

Because the `static` command was removed starting with code 8.3(1), the DNS rewrite function was migrated to both object-nat and twice-nat; basically, you must configure a regular static object or twice nat and specify the `dns` keyword.

Configuration

ASA1:

```
static (VLAN49,VLAN59) 136.1.59.1 150.1.1.1 dns
```

R1:

```
ip dns server
ip host www.ine.com 150.1.1.1
```

ASA2:

```
object network R2_LO1_MAPPED
 host 136.1.19.22
!
object network R2_LO1_REAL
 host 150.1.22.22
 nat (VLAN29,VLAN19) static R2_LO1_MAPPED dns
```

R2:

```
ip dns server
ip host www.inelab.com 150.1.22.22
```

Verification

https://t.me/learningnets

Configure R2 to use R1 as its DNS server, and test DNS resolution before applying the configuration on ASA1:

```
R2:
ip name-server 136.1.49.1
ip domain-lookup

R2#ping www.ine.com
Translating "www.ine.com"...domain server (136.1.49.1) [OK]

Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 150.1.1.1
, timeout is 2 seconds: !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Apply the configuration on ASA1, and test DNS resolution again:

```
R2#ping www.ine.com
Translating "www.ine.com"...domain server (136.1.49.1) [OK]

Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 136.1.59.1
, timeout is 2 seconds: .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Configure R1 to use R2 as its DNS server, and test DNS resolution before applying the configuration on ASA2:

```
R1:
ip name-server 136.1.29.2
ip domain-lookup

R1#ping www.inelab.com
Translating "www.inelab.com"...domain server (136.1.29.2) [OK]

Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 150.1.22.22
, timeout is 2 seconds: !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Apply the configuration on ASA2, and test DNS resolution again:

```
R1#ping www.inelab.com
```

https://t.me/learningnets

```
Translating "www.inelab.com"...domain server (136.1.29.2) [OK]

Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 136.1.19.22
, timeout is 2 seconds: .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Transparent Firewall

Task

- Configure ASA1 in transparent mode between VLAN 49 and VLAN 59:
  - Use subnet 136.1.59.0/24 and assign R1 the IP address of 136.1.59.1.
  - Ensure that multicast traffic is allowed bidirectionally without configuring access-lists.
  - The BGP session between R1 and R2 should be in the UP state.
- Configure ASA2 in transparent mode between VLAN 26 and VLAN 29:
  - Use subnet 172.16.10.0/24 and assign R2 the IP address of 172.16.10.2.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Transparent firewall mode turns the ASA into a layer 2 bridging device. In this mode, the firewall acts as a bump on the wire, inspecting traffic transparently as it flows through because hosts on the network do not see the firewall as a routing hop. When configuring the transparent firewall, you do not assign any IP addresses to the interfaces. However, every interface must still have a nameif and security-level

https://t.me/learningnets

configured.

For ASA code before 8.3(1), you must globally assign an IPv4 and/or IPv6 address to the firewall unit itself, or it won't allow any IPv4/IPv6 traffic through. Also, the firewall supports a maximum of two active interfaces (usually the inside and the outside).

Like any bridge, the firewall builds a MAC address table associated with the interfaces. You can review this table using the command `show mac-address-table .` You may even statically associate MAC addresses with the interfaces using the global command `mac-address-table static <nameif> <mac_address>` if you want to bind a host to one zone/interface of the firewall. The firewall will forward frames based on the destination MAC addresses and the bridging table. However, if the MAC address is unknown, it will not flood the frame out as a regular switch/bridge; instead it will drop the first frame and:

- If the destination IP address from the packet is on the directly connected subnet of the ASA, it will generate an ARP request to create the entry in its MAC address table.
- If the destination IP address from the packet is not on the directly connected subnet of the ASA, it will generate a ping toward the destination to create the entry in its MAC address table.

In both cases, if the unicast destination MAC address is not learned, all traffic for that MAC address is dropped.

All IPv4/IPv6 unicast IP traffic is permitted from higher to lower security interface by default, just like in routed mode; in the opposite direction, it must be allowed through access-lists. Non-IP traffic types (IPX, Apple-Talk, MPLS, and CDP) could be permitted using Ethertype access-lists; by default these are dropped by the firewall. Also, the firewall cannot permit protocols with an Ethertype value lower than 0x600, the exception being for STP and IS-IS. Some critical layer 2 traffic is by default permitted to flow bidirectionally between interfaces (lower to higher and higher to lower):

- ARP requests and responses, required to permit IPv4 address resolution. You may still control ARP packets flow using ARP inspection.
- Spanning Tree BPDUs are permitted to allow for loop-less topology calculation by STP algorithm.

However, the firewall only permits IEEE BPDUs and untagged PVST BPDUs. Therefore, if you want to ensure that STP traffic is not blocked by the firewall, follow these guidelines:

- Firewall interfaces are connected to access ports as Cisco switches send IEEE STP

https://t.me/learningnets

BPDUs out of access ports.
- Firewall interfaces are connected to dot1q trunk links with only one allowed VLAN, which is native on the trunks. In this case, only untagged SSTP BPDUs and IEEE STP BPDUs are sent by Cisco switches.

As for multicast and broadcast traffic, the firewall permits some UDP multicast (such as RIP) to flow from inside to outside without an explicit permission. Other types of multicast traffic (such as OSPF, PIM, and IGMP) should be explicitly permitted even on the inside interface. If firewall interfaces have the same security-level and the command `same-security-traffic permit inter-interface` is configured, it will allow all multicast traffic bidirectionally as long as the multicast MAC address is in the range of 0100.5E00.0000-0100.5EFE.FFFF for IPv4 and 3333.0000.0000-3333.FFFF.FFFF for IPv6; also, broadcast traffic will be allowed if the destination MAC address is FFFF.FFFF.FFFF.

For ASA code after 8.3(1), transparent mode configuration has been changed, but functionality remains the same:

- Interfaces are members of a bridge-group, configured with the interface-level command `bridge-group <nr> .`
- ASA supports a maximum of 8 bridge-groups and a maximum of 4 active interfaces per bridge-group.
- No global IPv4/IPv6 addresses are configured; instead, the IPv4/IPv6 address is configured on the logical BVI interface, just like with bridging on routers.

Because the BGP session between R1 and R2 is authenticated, ASA1 must be configured appropriately (the requirements are the same for both routed and transparent mode):

- By default, ASA strips TCP options from the header; because the BGP authentication hash is embedded into TCP option 19, ASA will remove it.
- By default, ASA randomizes the ISN; because BGP uses the ISN when creating the MD5 hash, authentication will fail.

Configuration

```
R1:
interface GigabitEthernet0/1
 ip addres 136.1.59.1 255.255.255.0


R2:
```

https://t.me/learningnets

```
interface GigabitEthernet0/0
 ip address 172.16.10.2 255.255.255.0
```

ASA1:

```
firewall transparent
!
interface Ethernet0/0
 nameif VLAN49
 no shutdown
!
interface Ethernet0/2
 nameif VLAN59
 no shutdown
!
same-security-traffic permit inter-interface
ip address 136.1.59.15 255.255.255.0
!
class-map BGP
 match port tcp eq 179
!
tcp-map BGP
 tcp-options range 19 19 allow
!
policy-map global_policy
 class BGP
  set connection random-sequence-number disable
  set connection advanced-options BGP
```

ASA2:

```
firewall transparent
!
interface Ethernet0/1
 no shutdown
!
interface Ethernet0/1.26
 vlan 26
 nameif VLAN26
 bridge-group 1
!
interface Ethernet0/1.29
 vlan 29
 nameif VLAN29
 bridge-group 1
```

https://t.me/learningnets

```
!
same-security-traffic permit inter-interface
!
interface BVI1
 ip address 172.16.10.16 255.255.255.0
```

Verification

Verify the firewall mode, IP address, and MAC table on ASA1:

```
ciscoasa# show firewall
Firewall mode: Transparent
!
!ciscoasa# show ip address
Management System IP Address:
        ip address 136.1.59.15 255.255.255.0Management Current IP Address:
ip address 136.1.59.15 255.255.255.0
!
!ciscoasa# show mac-address-table

interface                   mac  address          type      Age(min)
-------------------------------------------------------------------
VLAN59                     d867.d9e0.bbc1          dynamic    3
VLAN59                     001b.5408.f38f          dynamic    5
VLAN49                     001b.5408.f38e          dynamic    3
VLAN49                     d48c.b5da.d549          dynamic    3
```

Verify that the BGP session between R1 and R2 is UP and authenticated:

```
R1#show ip bgp summary | b Neighbor
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
136.1.59.2      4           12      34       35        1    0    0 00:27:54        0
!
!R1#show ip bgp neighbors 136.1.59.2 | section Option
Option Flags: nagle, path mtu capable,md5
```

Verify that ASA1 forwards BPDUs:

```
SW2#show spanning-tree interface fastEthernet 1/0/12 detail
 Port 14 (FastEthernet1/0/12) of VLAN0049 is designated forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.14.
   Designated root has priority 49, address 001b.d58b.2180
```

https://t.me/learningnets

```
   Designated bridge has priority 32817, address 001b.5408.f380
   Designated port id is 128.14, designated path cost 19
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default BPDU: sent 2927, received 3
!
!SW2#show spanning-tree interface fastEthernet 1/0/13 detail
 Port 15 (FastEthernet1/0/13) of VLAN0059 is root forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.15.
   Designated root has priority 49, address 001b.d58b.2180
   Designated bridge has priority 32817, address 001b.5408.f380
   Designated port id is 128.14, designated path cost 19
   Timers: message age 2, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default BPDU: sent 4, received 2884
```

Verify the firewall mode, IP address, and MAC table on ASA2:

https://t.me/learningnets

```
ciscoasa# show firewall
Firewall mode: Transparent

!
!ciscoasa# show bridge-group

Static mac-address entries:  0 (in use), 65535 (max)
Dynamic mac-address entries: 2 (in use), 65535 (max)
 Bridge Group: 1
 Interfaces:
 Ethernet0/1.26
 Ethernet0/1.29


 Management System IP Address:  172.16.10.16 255.255.255.0
Management Current IP Address: 172.16.10.16 255.255.255.0

 Management IPv6 Global Unicast Address(es):
    N/A
 Static mac-address entries:  0
 Dynamic mac-address entries: 2
!
!ciscoasa# show mac-address-table

interface                   mac  address          type      Age(min)   bridge-group
--------------------------------------------------------------------------------
VLAN29                     d867.d9e0.bbc0          dynamic      5         1
VLAN26                     0050.5698.62b1          dynamic      5         1
```

Verify that ASA2 forwards traffic:

```
R2#ping 172.16.10.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.10.100, timeout is 2 seconds:.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Verify that ASA2 forwards BPDUs:

```
SW1#show spanning-tree vlan 26 interface fastEthernet 1/0/15 detail
 Port 17 (FastEthernet1/0/15) of VLAN0026 is designated forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.17.
   Designated root has priority 26, address 001b.d58b.2180
```

https://t.me/learningnets

```
   Designated bridge has priority 26, address 001b.d58b.2180
   Designated port id is 128.17, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default BPDU: sent 813, received 2
!
!SW1#show spanning-tree vlan 29 interface fastEthernet 1/0/15 detail
 Port 17 (FastEthernet1/0/15) of VLAN0029 is root forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.17.
   Designated root has priority 26, address 001b.d58b.2180
   Designated bridge has priority 26, address 001b.d58b.2180
   Designated port id is 128.17, designated path cost 0
   Timers: message age 2, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default BPDU: sent 2, received 813
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA ARP Inspection

Task

- Configure ASA1 in transparent mode between VLAN 49 and VLAN 59:
  - Use subnet 136.1.59.0/24 and assign R1 the IP address of 136.1.59.1.
  - ASA should enforce consistency in ARP requests and responses.
  - ASA should allow ARP traffic only for statically configured mappings.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

ARP traffic is permitted to flow across the transparent firewall to facilitate in IPv4 address resolution. However, certain types of network attacks utilize ARP and might compromise your network security. The most common attack is known as ARP spoofing, where a malicious user sends crafted ARP packets to poison another host's ARP cache and pretend it is operating under a different IP address. This allows an attacker to attract traffic not intended for it, by taking over the default gateway's IP, for example.

https://t.me/learningnets

ARP inspection prevents this type of attack by looking inside the ARP replies/responses and ensuring that the IP to MAC address bindings from the ARP payload match the pre-configured table. To activate this feature, you must first populate the ARP static bindings table using the global command **`arp [nameif] <ip_address> <mac_address>`**, which binds the IP to the MAC on the particular interface. To enable ARP inspection for ARP packets entering a particular interface, use the global command **`arp-inspection [nameif] enable [flood|no-flood]`**. The flood option permits ARP packets that contains IP addresses not matching the static table to be flooded out on other interfaces. The no-flood option will permit only valid ARP packets for IP addresses in the inspection table. Unlike ARP Inspection in the Catalyst switches, the ASA firewall cannot use DHCP snooping to dynamically populate the ARP bindings table.

Configuration

Identify MAC addresses of R1 and R2; these will differ based on your rack assignment:

```
R1#show interfaces gigabitEthernet0/1 | i address
  Hardware is CN Gigabit Ethernet, address is d48c.b5da.d549 (bia d48c.b5da.d549)
  Internet address is 136.1.59.1/24
!
!R2#show interfaces gigabitEthernet0/1 | i address
  Hardware is CN Gigabit Ethernet, address is d867.d9e0.bbc1 (bia d867.d9e0.bbc1)
  Internet address is 136.1.59.2/24
```

R1:

```
interface GigabitEthernet0/1
 ip addres 136.1.59.1 255.255.255.0
```

ASA1:

```
firewall transparent
!
interface Ethernet0/0
 nameif VLAN49
 no shutdown
!
interface Ethernet0/2
 nameif VLAN59
 no shutdown
!
```

https://t.me/learningnets

```
same-security-traffic permit inter-interface
ip address 136.1.59.15 255.255.255.0
!
arp VLAN49 136.1.59.1 d48c.b5da.d549
arp VLAN59 136.1.59.2 d867.d9e0.bbc1
!
arp-inspection VLAN49 enable no-flood
arp-inspection VLAN59 enable no-flood
```

Verification

Verify the ARP inspection configuration:

```
ciscoasa# show arp-inspection
interface                arp-inspection         miss
-----------------------------------------------------
VLAN49                   enabled                no-flood
VLAN59                   enabled                no-flood
!
!ciscoasa# show arp

        VLAN49 136.1.59.1 d48c.b5da.d549 -
        VLAN59 136.1.59.2 d867.d9e0.bbc1 -
```

Verify that there is IP connectivity between R1 and R2:

```
R1#ping 136.1.59.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.59.2, timeout is 2 seconds:.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Change the MAC address on R1 or R2, and verify that traffic no longer works, because ARP is no longer functional:

```
R1:
interface gigabitEthernet 0/1
 mac-address aaaa.aaaa.aaaa


R1#ping 136.1.59.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.59.2, timeout is 2 seconds:.....
```

https://t.me/learningnets

```
Success rate is 0 percent (0/5)
!
!R1#show ip arp 136.1.59.2
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  136.1.59.2              0  Incomplete
      ARPA
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA EtherType Access-Lists

Task

- Configure ASA1 in transparent mode between VLAN 49 and VLAN 59:
  - Use subnet 136.1.59.0/24 and assign R1 the IP address of 136.1.59.1.
  - Block spanning-tree BPDUs going across the firewall.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

As we know, the firewall by default allows the ARP packets and STP BPDUs (IEEE and PVST+ untagged). As for other non-IP packets, the firewall will not permit 802.2 LLC and SNAP frames because it only recognizes Ethernet II frames with the EtherType field higher than or equal to 0x600. Thus, the firewall will not pass any LLC encapsulated traffic, such as VTP, UDLD, or IPX. (STP BPDUs are also LLC encapsulated, but the firewall still permits them as an exception; the same is true for CDP, but CDP is not allowed by default.)

As for the valid Ethernet II frames, you may control them using the ethertype access-

https://t.me/learningnets

lists. The format for the access-list is **`access-list <name> ethertype permit <Ethertype>`** with reserved names for MPLS, STP, and IPX Ethernet II traffic. The special keyword **bpdu** allows you to filter the reception of STP BPDUs (which are in LLC format, not Ethernet II). You may use this to deny STP protocol across the transparent firewall. When the ethertype access-list has been configured, you may apply it using the regular command **`access-group`** . Note that IPv4, IPv6, and Ethertype access-list may apply to the same interface at the same time.

Configuration

```
R1:
interface GigabitEthernet0/1
 ip addres 136.1.59.1 255.255.255.0


ASA1:

firewall transparent
!
interface Ethernet0/0
 nameif VLAN49
 no shutdown
!
interface Ethernet0/2
 nameif VLAN59
 no shutdown
!
same-security-traffic permit intra-interface
ip address 136.1.59.15 255.255.255.0
!
access-list LAYER2 ethertype deny bpdu
access-list LAYER2 ethertype permit any
!
access-group LAYER2 in interface VLAN49
access-group LAYER2 in interface VLAN59
```

Verification

Clear the STP counters and verify that no BPDUs are received by SW1 from ASA1 interfaces:

```
SW2#clear spanning-tree counters
```

https://t.me/learningnets

```
!
!SW2#show spanning-tree interface fastEthernet 1/0/12 detail
 Port 14 (FastEthernet1/0/12) of VLAN0049 is designated forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.14.
   Designated root has priority 49, address 001b.d58b.2180
   Designated bridge has priority 32817, address 001b.5408.f380
   Designated port id is 128.14, designated path cost 19
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default BPDU: sent 77, received 0
!
!SW2#show spanning-tree interface fastEthernet 1/0/13 detail
 Port 15 (FastEthernet1/0/13) of VLAN0059 is designated forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.15.
   Designated root has priority 59, address 001b.d58b.2180
   Designated bridge has priority 32827, address 001b.5408.f380
   Designated port id is 128.15, designated path cost 19
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   The port is in the portfast mode
   Link type is point-to-point by default BPDU: sent 78, received 0
```

Verify that BPDUs are matched but dropped:

```
ciscoasa# show access-list
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096)
            alert-interval 300
access-list LAYER2; 2 elements access-list LAYER2 ethertype deny bpdu (hitcount=429)
access-list LAYER2 ethertype permit any (hitcount=25)
!
!ciscoasa# show asp drop frame l2_acl
FP L2 rule drop (l2_acl)                                                43



Last clearing: 10:58:19 UTC Nov 17 2013 by enable_15
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Transparent NAT

Task

- Configure ASA1 in transparent mode between VLAN 49 and VLAN 59:
  - Use subnet 136.1.59.0/24 and assign R1 the IP address of 136.1.59.1.
  - Traffic from R2 toward R1 should be dynamically NAT'ed into 136.1.109.150.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

One of the new features introduced in version 8.0 was NAT support in transparent mode firewall. You may now apply NAT translations to traffic crossing the firewall, using the regular NAT rules (it works for both old and new NAT). Remember, though, that you cannot use interface PAT, because the physical interfaces have no IP addresses assigned. Additionally:

- If the translated address is on the same subnet as the ASA, ASA will proxy-arp for it, just like in routed mode.
- If the translated addresses is not on the same subnet as the ASA, you will need

https://t.me/learningnets

additional static routes configured on the upstream routers to route traffic back through the firewall, just like in routed mode.

- If the original source address being translated is not on the same subnet as the ASA, you will need additional static route configured on the ASA for the original source; this is because ASA will do a route-lookup for the real/original address before doing the untranslation, so packet switching is not based on the destination IP address.

When translating IP addresses on the same segment as the ASA, ARP inspection no longer applies. Therefore, the ARP responses might uncover the real IP addresses of the hosts in the segment, possibly causing confusion.

Configuration

```
R1:
interface GigabitEthernet0/1
 ip addres 136.1.59.1 255.255.255.0
!
ip route 136.1.109.150 255.255.255.255 136.1.59.2


ASA1:

firewall transparent
!
interface Ethernet0/0
 nameif VLAN49
 no shutdown
!
interface Ethernet0/2
 nameif VLAN59
 no shutdown
!
same-security-traffic permit inter-interface
ip address 136.1.59.15 255.255.255.0
!
nat (VLAN59) 10 136.1.59.2 255.255.255.255
global (VLAN49) 10 136.1.109.150
```

Verification

https://t.me/learningnets

Initiate traffic from R2 and verify the translation on ASA1:

```
R2#ping 136.1.59.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.59.1, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!ciscoasa# show xlate
1 in use, 3 most usedPAT Global 136.1.109.150(28005) Local 136.1.59.2 ICMP id 14
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Firewall Contexts

Task

- Configure ASA3 to run in multi-context mode as follows:

Context	Interface1	Interface1 Nameif	Interface1 IP	Interface2	Interface2 Nameif	I
SW1-R3	Gi0/0.34	VLAN34	136.1.34.17/24	Gi0/1	VLAN93	13
SW2-R3	Gi0/0.34	VLAN34	136.1.34.18/24	Gi0/3	VLAN94	13

- Physical interfaces should not be visible in the **SW1-R3** context.
- Security-level for the interfaces should map the VLAN number.
- Ensure automatic MAC address assignment for the shared interface.
- Telnet from SW1/SW2 Loopback0 interfaces toward R3's VLAN34 IP address should be functional.
- Create a dedicated admin context.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config`

https://t.me/learningnets

instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

ASA firewall supports software virtualization by means of so-called firewall contexts. Every context has its own set of routing, filtering/inspection, address translation rules, and assigned IPS sensors. For ASA code versions in the current CCIE Security Blueprint, there are some restrictions for ASA running in multiple-context mode:

- Only static routing is supported.
- Features such as QoS and VPN are not supported
- All contexts must be in either routing or transparent firewall mode, you cannot mix modes in different contexts.

Contexts are generally useful when you have different sets of security policies applied to different traffic flows. This might be the case when the firewall protects multiple customers or departments in the same organization. It is common to see other virtualization technologies, such as VLANs or VRFs, used alongside with the firewall contexts; however, the firewall contexts have significant differences from the VRFs seen in the IOS routers, as we will discuss later. To turn the firewall to the multiple contexts mode, you should enter the global command `mode multiple` when logged in via the console port (you may do this remotely, converting the existing running configuration into the so-called **admin** context, but you risk losing connection to the box); this will force the mode change and reload the appliance. If you connect to the appliance on the console port, you are logging in to the **system** context; the sole purpose of this context is to define other contexts and allocate resources to them. You should do the following while logged in to the system context:

- Configure physical interfaces; you must un-shutdown the interfaces that you want to allocate to the contexts. If you are creating sub-interfaces using VLANs, you should do it under the system context as well.
- Define the admin context; this is a special context that allows logging in to the firewall remotely (via ssh, telnet, or https). This context should be configured first, because the firewall won't let you create any other contexts

https://t.me/learningnets

before designating the admin context using the global command `admin-context <name>`. When you convert from the single-context mode, the running configuration of the single mode is converted and saved as the admin context, to allow remote firewall administration. To access the system context remotely, you should log in to the admin context using any configured remote access method and issue the command `changeto system`. Every configured context should have a configuration URL defined using the context-level command `config-url <path>` to store its configuration; without this command the context configuration is incomplete. After the context has been defined, you may switch to the context configuration using the command `changeto context <name>`.

- Define additional contexts if needed using the command `context <name>`, and allocate physical interfaces to the contexts using the context-level command `allocate-interface <physical-interface> [<logical-name>]`. Here, *physical-interface* is the physical interface or sub-interface name and *logical-name* is the name that the context sees for this interface (this is not the interface nameif, as this is configured within the context). Using this command, you can hide the real physical interfaces from the context administrators (for example, hide VLAN numbers) to provide an additional level of isolation from the physical configuration.
- Change to the context configuration and proceed to configure it just like a regular ASA: Assign interface names, security levels, and IP addresses. However, this time you must generally has different IP and MAC addresses on this interface (it is possible to share the IP address as well, though). You may use non-overlapping subnets or simply different IPs on the same subnet. Remember that by default both contexts will inherit the same MAC address from the shared physical interface. This might result in the firewall not being able to classify the incoming traffic properly. Use the command `mac-address auto` in the system context to automatically generate a MAC address for every new virtual interface, or change the MAC at the context level from the interface configuration mode.

Use the command `write memory all` in the system context to save all contexts configuration on the persistent storage. You may also save configuration for a context individually when logged in under the particular context using the command `write memory`. Note that firewall contexts do not share IP routing tables, so if you want to establish communication between the routing contexts, you must do one of

https://t.me/learningnets

the following:

- Configure each context with a set of static routes for the subnets connected or located behind the other context.
- Use an external router that has full knowledge of the subnets behind each of the contexts to provide connectivity.

Configuration

SW1:

```
interface FastEthernet1/0/16
 switchport mode access
 switchport access vlan 94
 no shutdown
!
interface FastEthernet1/0/17
 switchport mode access
 switchport access vlan 93
 no shutdown
```

SW2:

```
interface FastEthernet1/0/16
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
```

ASA3 System Context:

```
mode multiple
!
mac-address auto
!
interface GigabitEthernet0/0
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
!
interface GigabitEthernet0/1
 no shutdown
!
interface GigabitEthernet0/3
```

https://t.me/learningnets

```
 no shutdown
!
admin-context admin
context admin
 config-url disk0:/admin.cfg
!
context SW1-R3
 allocate-interface GigabitEthernet0/0.34 VLAN34
 allocate-interface GigabitEthernet0/1 VLAN93
 config-url disk0:/SW1-R3.cfg
!
context SW2-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/3
 config-url disk0:/SW2-R3.cfg
!
end
```

ASA3 SW1-R3 Context:

```
changeto context SW1-R3
!
interface VLAN34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.17 255.255.255.0
!
interface VLAN93
 nameif VLAN93
 security-level 93
 ip address 136.1.93.17 255.255.255.0
!
route VLAN93 150.1.9.9 255.255.255.255 136.1.93.9
```

ASA3 SW2-R3 Context:

```
changeto context SW2-R3
!
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.18 255.255.255.0
!
interface GigabitEthernet0/3
 nameif VLAN94
```

https://t.me/learningnets

```
 security-level 94
 ip address 136.1.94.18 255.255.255.0
!
route VLAN94 150.1.10.10 255.255.255.255 136.1.94.10
```

Verification

Verify context configuration:

```
ciscoasa# show context
Context Name      Class      Interfaces           URL *admin
            default                          disk0:/admin.cfg SW1-R3
           default    GigabitEthernet0/0.34, disk0:/SW1-R3.cfg
                              GigabitEthernet0/1   SW2-R3
           default    GigabitEthernet0/0.34, disk0:/SW2-R3.cfg
                              GigabitEthernet0/3

Total active Security Contexts: 3
```

Verify that in **SW1-R3** context, physical interfaces are not visible:

```
ciscoasa/SW1-R3# show int ip brief
Interface                  IP-Address      OK? Method Status                Protocol vlan93
                     136.1.93.17     YES manual up                    up   vlan34
                     136.1.34.17     YES manual up                    up
!
!ciscoasa/SW2-R3# show int ip brief
Interface                  IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0.34
      136.1.34.18     YES manual up                    up   GigabitEthernet0/3
         136.1.94.18     YES manual up                    up
```

Verify nameifs and security-levels:

https://t.me/learningnets

```
ciscoasa/SW1-R3# show nameif

Interface                Name                    Security vlan93 VLAN93                  93
vlan34 VLAN34                   34

!
!ciscoasa/SW2-R3# show nameif

Interface                Name                    Security GigabitEthernet0/0.34
VLAN34                   34
GigabitEthernet0/3 VLAN94                  94
```

Initiate traffic from SW1's Loopback0:

```
SW1#telnet 136.1.34.3 /source-interface loopback0
Trying 136.1.34.3 ... Open


User Access Verification

Password: R3>show users
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:01:48
*514 vty 0                 idle                 00:00:00 150.1.9.9

  Interface    User               Mode         Idle     Peer Address
!
!ciscoasa/SW1-R3# show conn
1 in use, 1 most used
TCP VLAN34 136.1.34.3:23 VLAN93 150.1.9.9:40678, idle 0:00:18, bytes 407, flags UIO
```

Initiate traffic from SW2's Loopback0:

```
SW2#telnet 136.1.34.3 /source-interface loopback0
Trying 136.1.34.3 ... Open


User Access Verification

Password: R3>show users
```

https://t.me/learningnets

```
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:02:22
*514 vty 0                 idle                 00:00:00 150.1.10.10

  Interface    User               Mode         Idle     Peer Address
!
!ciscoasa/SW2-R3# show conn
1 in use, 2 most used
TCP VLAN34 136.1.34.3:23 VLAN94 150.1.10.10:12331, idle 0:00:12, bytes 426, flags UIO
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Contexts Traffic Classification

Task

- Configure ASA3 to run in multi-context mode as follows:

Context	Interface1	Interface1 Nameif	Interface1 IP	Interface2	Interface2 Nameif	I
SW1-R3	Gi0/0.34	VLAN34	136.1.34.17/24	Gi0/1	VLAN93	13
SW2-R3	Gi0/0.34	VLAN34	136.1.34.18/24	Gi0/3	VLAN94	13

- Security-level for the interfaces should map the VLAN number.
- SW1 and SW2 should reach R3's VLAN 34 on the IP addresses of 136.1.93.3 and 136.1.94.3.
- Do not change the default MAC addresses.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by

https://t.me/learningnets

the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Let’s discuss how the firewall classifies incoming traffic and assigns it to the correct context. If the interface where traffic has been received is uniquely allocated to the context, there are no requirements, but if the interface is shared, additional rules are needed:

- First, ASA tries to map traffic to the correct context based on the destination MAC address; configure different MAC addresses for the shared interface using the system context `mac-address auto` global level command or the specific context's `mac-address <value>` interface-level command.
- If the MAC address is the same in both contexts for the same interface, the firewall attempts to use NAT configuration in every context to resolve the conflicts. This may happen if you intentionally assign the same IP address to both contexts or did not assign different MAC addresses to the shared interfaces. The firewall attempts to match the destination IP address and TCP/UDP port information in the packet with the active translation slots in every context. The context with the matching translation slot is selected as the target context. This type of classification allows sharing the same IP subnet or even IP address on the shared interface. You are not required to have unique MAC addresses in each context, because the translation slots are used for traffic classification.
- If all contexts on the shared interface use the same IP address and MAC, you cannot access the contexts on the shared interface and therefore all traffic is affected. This is why it is generally recommended to use separate IP addresses (MAC could be the same) on the shared interfaces.

If you look at the solution for our scenario, you will see that only NAT-based classification is used because MAC address change is not allowed. Because VLAN34 is the shared interface, return traffic from VLAN 34 toward VLAN 93 and VLAN 94 must be correctly classified, so dynamic NAT was configured for initial traffic.

https://t.me/learningnets

ASA does not support interface sharing for transparent firewall mode.

Configuration

SW1:

```
interface FastEthernet1/0/16
 switchport mode access
 switchport access vlan 94
 no shutdown
!
interface FastEthernet1/0/17
 switchport mode access
 switchport access vlan 93
 no shutdown
```

SW2:

```
interface range FastEthernet1/0/16 , FastEthernet1/0/18
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
```

ASA3 System Context:

```
mode multiple
!
interface GigabitEthernet0/0
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
!
interface GigabitEthernet0/1
 no shutdown
!
interface GigabitEthernet0/3
 no shutdown
!
admin-context SW1-R3
!
context SW1-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/1
 config-url disk0:/SW1-R3.cfg
```

https://t.me/learningnets

```
!
context SW2-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/3
 config-url disk0:/SW2-R3.cfg
!
end
```

ASA3 SW1-R3 Context:

```
changeto context SW1-R3
!
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.17 255.255.255.0
!
interface GigabitEthernet0/1
 nameif VLAN93
 security-level 93
 ip address 136.1.93.17 255.255.255.0
!
object network VLAN93
 subnet 136.1.93.0 255.255.255.0
 nat (VLAN93,VLAN34) dynamic interface
!
object network R3
 host 136.1.34.3
 nat (VLAN34,VLAN93) static 136.1.93.3
```

ASA3 SW2-R3 Context:

```
changeto context SW2-R3
!
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.18 255.255.255.0
!
interface GigabitEthernet0/3
 nameif VLAN94
 security-level 94
 ip address 136.1.94.18 255.255.255.0
!
object network VLAN94
 subnet 136.1.94.0 255.255.255.0
 nat (VLAN94,VLAN34) dynamic interface
```

https://t.me/learningnets

```
!
object network R3
 host 136.1.34.3
 nat (VLAN34,VLAN94) static 136.1.94.3
```

Verification

Verify NAT configuration on both contexts:

```
ciscoasa/SW1-R3# show nat

Auto NAT Policies (Section 2)
1 (VLAN34) to (VLAN93) source static R3 136.1.93.3
    translate_hits = 0, untranslate_hits = 3
2 (VLAN93) to (VLAN34) source dynamic VLAN93 interface
    translate_hits = 5, untranslate_hits = 7
!
!ciscoasa/SW2-R3# show nat


Auto NAT Policies (Section 2)
1 (VLAN34) to (VLAN94) source static R3 136.1.94.3
    translate_hits = 0, untranslate_hits = 0
2 (VLAN94) to (VLAN34) source dynamic VLAN94 interface
    translate_hits = 3, untranslate_hits = 3
```

Verify that traffic between SW1 and R3 is functional:

```
SW1#telnet 136.1.93.3
Trying 136.1.93.3 ... Open


User Access Verification

Password: R3>show users
    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:19:22
*514 vty 0                 idle                 00:00:00 136.1.34.17

  Interface    User               Mode         Idle     Peer Address
!
!ciscoasa/SW1-R3# show conn long
1 in use, 1 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
```

https://t.me/learningnets

```
        B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
        D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
        G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
        i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
        k - Skinny media, M - SMTP data, m - SIP media, n - GUP
        O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
        q - SQL*Net data, R - outside acknowledged FIN,
        R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
        s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
        V - VPN orphan, W - WAAS,
        X - inspected by service module
TCP VLAN34:136.1.34.3/23 (136.1.93.3/23) VLAN93:136.1.93.9/15551 (136.1.34.17/21644), flags UIO, idle 39s, uptime 44
```

Verify that traffic between SW2 and R3 is functional:

https://t.me/learningnets

```
SW2#telnet 136.1.94.3

Trying 136.1.94.3 ... Open


User Access Verification

Password: R3>show users

    Line       User       Host(s)              Idle       Location
   0 con 0                 idle                 00:21:50
*514 vty 0                 idle                 00:00:00 136.1.34.18


  Interface    User               Mode         Idle     Peer Address
!
!ciscoasa/SW2-R3# show conn long

1 in use, 1 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed, C - CTIQBE media,
       D - DNS, d - dump, E - outside back connection, F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, M - SMTP data, m - SIP media, n - GUP
       O - outbound data, P - inside back connection, p - Phone-proxy TFTP connection,
       q - SQL*Net data, R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up,
       V - VPN orphan, W - WAAS,
       X - inspected by service module
TCP VLAN34:136.1.34.3/23 (136.1.94.3/23) VLAN94:136.1.94.10/39747 (136.1.34.18/35861), flags UIO, idle 15s, uptime 2
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Contexts Resource Management

Task

- Configure ASA3 to run in multi-context mode as follows:

Context	Interface1	Interface1 Nameif	Interface1 IP	Interface2	Interface2 Nameif	I
SW1-R3	Gi0/0.34	VLAN34	136.1.34.17/24	Gi0/1	VLAN93	1
SW2-R3	Gi0/0.34	VLAN34	136.1.34.18/24	Gi0/3	VLAN94	1

- Security-level for the interfaces should map the VLAN number.
- Context **SW1-R3** should be allowed to have no more than 1000 host and xlate entries, and no more than 10000 connections.
- Context **SW2-R3** should be limited to no more than 500 host and xlate entries, and no more than 5000 connections.
- Context **admin** should use the default resource limits.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of

https://t.me/learningnets

reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The firewall has limited resources, shared between the contexts. The resources include concurrent connections, inspections, translation slots, management sessions (telnet, ssh, and https), number of inside hosts, and so on; you can find the full list in the firewall configuration guide. To avoid resource contention and exhaustion, the firewall allows limiting per-context resources using the resource class concept. Every class specifies the amount of resources available to a context and classes are assigned to the contexts to enforce the limits. By default, all contexts are assigned class **default**. Note that contexts do not share the particular class resources, they only inherit the resource limits set by a class.

When you create a new class, it inherits all limits from the **default** resource class. When you re-define any particular limit in the new class, you automatically override the default setting for this limit. You may also configure the default class settings and all classes will inherit these values, unless overridden within the specific class.

The appliance never reserves any resources for classes. It simply uses them to compute the resource limits and satisfies any request that is within the limit for a given class. For example, suppose the system supports up to 1000 connection maximum, and you create a new class with the limit of 500 connections. You assign this class to 3 contexts. At the peak of their usage, every context may request up to 500 connections, exceeding the total limit of 1000. Thus it is up to the administrator to properly set limits and prevent resource starvation. You may set resource limits in absolute values (such as number of connections or hosts) or in percentages of the maximum resource available. The syntax is:

```
class <name>
  limit-resource <resource> [<value>|{1-100%}]
```

Some resources, like connections, inspects, and syslogs, support rate limiting, using the class command `limit-resource rate [{conns|inspects|syslogs}|{1-100%}]`. In our scenario, we limit only certain resources for two classes. All other parameters remain unlimited, because they are inherited from the default class.

https://t.me/learningnets

Configuration

SW1:

```
interface FastEthernet1/0/16
 switchport mode access
 switchport access vlan 94
 no shutdown
!
interface FastEthernet1/0/17
 switchport mode access
 switchport access vlan 93
 no shutdown
```

SW2:

```
interface range FastEthernet1/0/16 , FastEthernet1/0/18
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
```

ASA3 System Context:

```
mode multiple
!
interface GigabitEthernet0/0
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
!
interface GigabitEthernet0/1
 no shutdown
!
interface GigabitEthernet0/3
 no shutdown
!
class SW1-R3
 limit-resource hosts 1000
 limit-resource xlates 1000
 limit-resource conns 10000
!
class SW2-R3
 limit-resource hosts 500
 limit-resource xlates 500
```

https://t.me/learningnets

```
 limit-resource conns 5000
!
admin-context admin
context admin
 config-url disk0:/admin.cfg
 member default
!
context SW1-R3
 allocate-interface GigabitEthernet0/0.34 VLAN34
 allocate-interface GigabitEthernet0/1 VLAN93
 config-url disk0:/SW1-R3.cfg
 member SW1-R3
!
context SW2-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/3
 config-url disk0:/SW2-R3.cfg
 member SW2-R3
!
end
```

ASA3 SW1-R3 Context:

```
changeto context SW1-R3
!
interface VLAN34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.17 255.255.255.0
!
interface VLAN93
 nameif VLAN93
 security-level 93
 ip address 136.1.93.17 255.255.255.0
```

ASA3 SW2-R3 Context:

```
changeto context SW2-R3
!
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.18 255.255.255.0
!
interface GigabitEthernet0/3
 nameif VLAN94
 security-level 94
```

https://t.me/learningnets

```
 ip address 136.1.94.18 255.255.255.0
```

Verification

Verify the available resources and their types:

```
ciscoasa# show resource types

Rate limited resource types:
  Conns           Connections/sec
  Inspects        Inspects/sec
  Syslogs         Syslogs/sec

Absolute limit types:
  Conns           Connections
  Hosts           Hosts
  ASDM            ASDM Connections
  SSH             SSH Sessions
  Telnet          Telnet Sessions
  Xlates          XLATE Objects
  All             All Resources
```

Check resource allocation per context:

```
ciscoasa# show resource allocation detail

Resource Origin:
        A       Value was derived from the resource 'all'
        C       Value set in the definition of this class
        D       Value set in default class
Resource         Class          Mmbrs  Origin      Limit      Total    Total %
Conns [rate]     default          all      CA  unlimited
                 SW1-R3             1      DA  unlimited
                 SW2-R3             1      DA  unlimited
                 All Contexts:      3

Inspects [rate]  default          all      CA  unlimited
                 SW1-R3             1      DA  unlimited
                 SW2-R3             1      DA  unlimited
                 All Contexts:      3

Syslogs [rate]   default          all      CA  unlimited
                 SW1-R3             1      DA  unlimited
                 SW2-R3             1      DA  unlimited
```

https://t.me/learningnets

```
                All Contexts:      3
Conns           default            all      CA  unlimited
                SW1-R3             1        C       10000      10000      3.98%
SW2-R3          1        C         5000       5000       1.99%
All Contexts:   3                             15000      5.97%
Hosts           default            all      CA  unlimited
                SW1-R3             1        C        1000       1000       0.00%
SW2-R3          1        C         500        500        0.00%
All Contexts:   3                             1500       0.00%

IPSec           default            all      CA  unlimited
                SW1-R3               1      DA  unlimited
                SW2-R3               1      DA  unlimited
                All Contexts:      3

Mac-addresses   default            all      CA  unlimited
                SW1-R3               1      DA  unlimited
                SW2-R3               1      DA  unlimited
                All Contexts:      3

ASDM            default            all       C          5                2.50%
                SW1-R3               1       D          5          5     2.50%
                SW2-R3               1       D          5          5     2.50%
                All Contexts:      3                              15     7.50%

SSH             default            all       C          5                5.00%
                SW1-R3               1       D          5          5     5.00%
                SW2-R3               1       D          5          5     5.00%
                All Contexts:      3                              15    15.00%

Telnet          default            all       C          5                5.00%
                SW1-R3               1       D          5          5     5.00%
                SW2-R3               1       D          5          5     5.00%
                All Contexts:      3                              15    15.00%
Xlates          default            all      CA  unlimited
                SW1-R3             1        C        1000       1000       0.00%
SW2-R3          1        C         500        500        0.00%
All Contexts:   3                             1500       0.00%
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Active-Standby Failover

Task

- Configure ASA3 and ASA4 for active-standby failover; the failover pair should use ASA3's MAC addresses:

Device	Interface1	Interface1 Nameif	Interface1 IP	Interface2	Interface2 Nameif	In
ASA3	Gi0/0.34	VLAN34	136.1.34.17/24	Gi0/1	VLAN93	13
ASA4	Gi0/0.34	VLAN34	136.1.34.18/24	Gi0/1	VLAN93	13

- Configure stateful failover for all sessions using Gi0/2 as the failover link with the name **FAILOVER** and the IP subnet 136.1.100.0/24.
- The units should monitor each other using the minimum poll times.
- Security-level should map the VLAN number and ASA hostname should not be visible at the command line.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure`

https://t.me/learningnets

command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Failover allows two firewall units to operate in hot standby mode. For two units to operate as a failover pair, their hardware and software configurations must be identical (flash disk and minor software version differences are allowed for zero downtime upgrade of a failover pair). One firewall unit is designated as primary and another as secondary, and by default the primary unit receives the active role and the secondary receives the standby role. At any moment of time, only one unit is active and forwards traffic, and the other remains in standby mode. When the active unit fails or interface failure criteria is met, the standby takes the role of the active unit, also inheriting its IP/MAC addresses. Failover is available in both transparent firewall and routed modes. The firewall supports two types of failover: stateful and stateless. During the stateless failover, the states of the currently active sessions, NAT translations, IPsec session, and other data are not copied between the active and standby units; after failover, users must re-initiate their connections. Stateful failover preserves all connection states during the failover and makes the switchover process almost seamless (there are, however, some exceptions—for example, clientless SSL VPN sessions cannot be replicated because of their functional design). The configurations of both units are kept synchronized all the time, because the commands from the active unit are always replicated to the standby but not vice-versa, so all configuration changes should be performed on the active unit; otherwise, because configuration is no longer synchronized, failover will no longer be functional.

Failover configuration always involves a failover interface, which is used to exchange information between the two units (for example, configuration updates) and health monitoring. In a stateful failover scenario, the same link can be used to constantly replicate session states. You may designate a dedicated physical interface for failover or allocate a sub-interface on any active data interface, but it is strongly recommended that you use a dedicated physical interface for stateful failover, of the same type and performance with data traffic interfaces (if at least one data traffic interface is GigabitEthernet, do not use a FastEthernet interface for session state replication, but also a GigabitEthernet).

All failover commands are global. To configure a failover link, use `failover lan interface <name> <physical-interface>` . Additionally, you must assign IP addresses to the failover interface to allow information exchange using the command `failover interface ip <name> <primary-ip> <mask> standby <secondary-ip>` . The last two mandatory commands to be entered are `failover lan unit [primary|secondary]` to designate the

https://t.me/learningnets

primary/secondary unit and `failover` to activate the failover configuration. These four commands should be configured on both units, assigning their respective roles. After this, the firewalls will detect each other, and the secondary will replicate the configuration from the primary unit. Note that the default failover mode is stateless; if you want to enable stateful failover mode, enter the command `failover link <name> <physical-interface>` .

When configuring the IP addresses on the interfaces, remember that the primary and secondary units should have different IP addresses. Use the command `ip address <IP_address> <mask> standby <IP_address>` to configure the IP addresses for the primary and secondary units at the same time that you configure the active unit. Failover occurs under three general conditions:

- The active unit detects system health issues (such as hardware, software, or power failure); active and standby roles are swapped when this is detected, based on the default or configured unit polling timers.
- The standby unit detects loss of contact with the active unit across the failover interface; both units constantly send keepalive message to each other across the failover link. If the standby unit loses three consecutive keepalives, it will try to restore contact with the active unit. The standby unit will broadcast ARP requests out of all interfaces, asking for the IP address of the active unit. If it receives the ARP response on the failover link, nothing changes. If the response is only received on a non-failover link, the standby unit marks the failover link as non-functional, but it does not fail over. Manual intervention is required to fix the problem, and this is normal behavior to avoid a dual active scenario. If no response is received on any interface, the standby unit fails over.
- The active unit detects loss of the monitored interfaces above the configured threshold. By default, when interface monitoring is enabled, every single physical interface failure would force the active unit to give its role to the standby; if the unit detects loss of carrier on any interface, it immediately declares the interface to be down. To account for more complex cases, interface monitoring is performed by sending and receiving keepalave packets to the standby unit. If the active unit does not receive any hello packets for the duration of half of the hold-interval, it will attempt counting packets on the monitored interface to see if any traffic enters the interface. If this does not succeed, the unit will start sending ARP requests for known destinations to provoke some responses and determine whether

https://t.me/learningnets

this generates traffic; if this fails, the unit will initiate failover.

By default, the firewall unit monitors all physical interfaces with IP addresses assigned, but not sub-interfaces. You may set up unit and interface polling timers using the command **`failover polltime [unit|interface] msec <hello> holdtime msec <hold>`** . If needed, you may increase the threshold for the number of failed monitored interfaces using the command **`failover interface-policy <N>%`** ; when this command is configured, the active unit will only initiate failover if the number or percentage of failed interfaces is equal to or above configured values. With default settings, any single interface failure will trigger failover.

From the configuration perspective, everything is configured on the primary unit; on the secondary units, only the four mandatory failover commands should be configured and failover interface be enabled.

Configuration

SW1:

```
interface range FastEthernet1/0/17 , FastEthernet1/0/19
 switchport mode access
 switchport access vlan 93
 no shutdown
```

SW2:

```
interface range FastEthernet1/0/16 , FastEthernet1/0/18
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
!
interface range FastEthernet1/0/17 , FastEthernet1/0/19
 switchport mode access
 switchport access vlan 100
 no shutdown
```

ASA3:

```
interface GigabitEthernet0/0
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
 nameif VLAN34
 security-level 34
```

https://t.me/learningnets

```
 ip address 136.1.34.17 255.255.255.0 standby 136.1.34.18
!
interface GigabitEthernet0/1
 nameif VLAN93
 security-level 93
 ip address 136.1.93.17 255.255.255.0 standby 136.1.93.18
 no shutdown
!
interface GigabitEthernet0/2
 no shutdown
!
failover lan unit primary
failover lan interface FAILOVER GigabitEthernet0/2
failover link failover GigabitEthernet0/2
failover interface ip FAILOVER 136.1.100.17 255.255.255.0 standby 136.1.100.18
failover polltime unit msec 200 holdtime msec 800
failover polltime interface msec 500 holdtime 5
failover replication http
failover
!
monitor-interface VLAN34
prompt state
```

ASA4:

```
interface GigabitEthernet0/2
 no shutdown
!
failover lan unit secondary
failover lan interface FAILOVER GigabitEthernet0/2
failover link failover GigabitEthernet0/2
failover interface ip FAILOVER 136.1.100.17 255.255.255.0 standby 136.1.100.18
failover
```

Verification

Verify failover interface configuration and IP address:

```
act# show failover interface
        interface FAILOVER GigabitEthernet0/2
                System IP Address: 136.1.100.17 255.255.255.0 My IP Address    : 136.1.100.17
                Other IP Address : 136.1.100.18
!
!act# failover exec mate show failover interface
```

https://t.me/learningnets

```
        interface FAILOVER GigabitEthernet0/2
                System IP Address: 136.1.100.17 255.255.255.0 My IP Address    : 136.1.100.18

                Other IP Address : 136.1.100.17
```

Verify that both interfaces are monitored and in normal state:

```
act# show monitor-interface

        This host: Primary - Active
                Interface VLAN34 (136.1.34.17): Normal (Monitored)
                Interface VLAN93 (136.1.93.17): Normal (Monitored)
        Other host: Secondary - Standby Ready
                Interface VLAN34 (136.1.34.18): Normal (Monitored)
                Interface VLAN93 (136.1.93.18): Normal (Monitored)
```

Verify failover state on both units as well as failover configuration:

```
act# show failover
Failover On
Failover unit Primary
Failover LAN Interface: FAILOVER GigabitEthernet0/2 (up)
Unit Poll frequency 200 milliseconds, holdtime 800 milliseconds
Interface Poll frequency 500 milliseconds, holdtime 5 seconds
Interface Policy 1
Monitored Interfaces 2 of 114 maximum
failover replication http
Version: Ours 8.6(1)10, Mate 8.6(1)10
Last Failover at: 16:22:18 UTC Nov 18 2013 This host: Primary - Active
                Active time: 4560 (sec)
                slot 0: ASA5515 hw/sw rev (1.0/8.6(1)10) status (Up Sys)
                  Interface VLAN34 (136.1.34.17): Normal (Monitored)
                  Interface VLAN93 (136.1.93.17): Normal (Monitored)
                slot 1: IPS5515 hw/sw rev (N/A/) status (Unresponsive/Up)
Other host: Secondary - Standby Ready

                Active time: 0 (sec)
                slot 0: ASA5515 hw/sw rev (1.0/8.6(1)10) status (Up Sys)
                  Interface VLAN34 (136.1.34.18): Normal (Monitored)
                  Interface VLAN93 (136.1.93.18): Normal (Monitored)
                slot 1: IPS5515 hw/sw rev (N/A/) status (Unresponsive/Up)

Stateful Failover Logical Update Statistics
        Link : FAILOVER GigabitEthernet0/2 (up)
        Stateful Obj    xmit       xerr       rcv        rerr
```

https://t.me/learningnets

```
General         600        0          599        0
sys cmd         599        0          599        0
up time         0          0          0          0
RPC services    0          0          0          0
TCP conn        0          0          0          0
UDP conn        0          0          0          0
ARP tbl         0          0          0          0
Xlate_Timeout   0          0          0          0
IPv6 ND tbl     0          0          0          0
VPN IKEv1 SA    0          0          0          0
VPN IKEv1 P2    0          0          0          0
VPN IKEv2 SA    0          0          0          0
VPN IKEv2 P2    0          0          0          0
VPN CTCP upd    0          0          0          0
VPN SDI upd     0          0          0          0
VPN DHCP upd    0          0          0          0
SIP Session     0          0          0          0
Route Session   0          0          0          0
User-Identity   1          0          0          0

Logical Update Queue Information
                Cur     Max     Total
Recv Q:         0       3       599
Xmit Q:         0       30      5867
```

Initiate a session from SW1 toward R3, and verify the session state on both active and standby units:

https://t.me/learningnets

```
SW1#telnet 136.1.34.3

Trying 136.1.34.3 ... Open


User Access Verification

Password: R3>

!
!act# show conn

9 in use, 11 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:15288, idle 0:03:55, bytes 111, flags UIO

!
!stby# show conn

9 in use, 11 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:15288, idle 0:03:44, bytes 111, flags UIO
```

Bring down one interface of ASA3, and verify that failover occurs and the session is not interrupted:

```
SW2:
interface fastEthernet1/0/16
 shutdown

stby# show failover
Failover On
Failover unit Primary
Failover LAN Interface: FAILOVER GigabitEthernet0/2 (up)
Unit Poll frequency 200 milliseconds, holdtime 800 milliseconds
Interface Poll frequency 500 milliseconds, holdtime 5 seconds
Interface Policy 1
Monitored Interfaces 2 of 114 maximum
failover replication http
Version: Ours 8.6(1)10, Mate 8.6(1)10
Last Failover at: 17:51:10 UTC Nov 18 2013 This host: Primary - Failed
                Active time: 5331 (sec)
                slot 0: ASA5515 hw/sw rev (1.0/8.6(1)10) status (Up Sys)
                  Interface VLAN34 (136.1.34.18): No Link (Waiting)
                  Interface VLAN93 (136.1.93.18): Normal (Monitored)
```

https://t.me/learningnets

```
                slot 1: IPS5515 hw/sw rev (N/A/) status (Unresponsive/Up)
Other host: Secondary - Active
                Active time: 73 (sec)
                slot 0: ASA5515 hw/sw rev (1.0/8.6(1)10) status (Up Sys)
                  Interface VLAN34 (136.1.34.17): Normal (Waiting)
                  Interface VLAN93 (136.1.93.17): Normal (Monitored)
                slot 1: IPS5515 hw/sw rev (N/A/) status (Unresponsive/Up)

Stateful Failover Logical Update Statistics
        Link : FAILOVER GigabitEthernet0/2 (up)
        Stateful Obj    xmit       xerr       rcv        rerr
        General         718        0          712        0
        sys cmd         712        0          712        0
        up time         0          0          0          0
        RPC services    0          0          0          0
        TCP conn        3          0          0          0
        UDP conn        0          0          0          0
        ARP tbl         2          0          0          0
        Xlate_Timeout   0          0          0          0
        IPv6 ND tbl     0          0          0          0
        VPN IKEv1 SA    0          0          0          0
        VPN IKEv1 P2    0          0          0          0
        VPN IKEv2 SA    0          0          0          0
        VPN IKEv2 P2    0          0          0          0
        VPN CTCP upd    0          0          0          0
        VPN SDI upd     0          0          0          0
        VPN DHCP upd    0          0          0          0
        SIP Session     0          0          0          0
        Route Session   0          0          0          0
        User-Identity   1          0          0          0

        Logical Update Queue Information
                        Cur     Max     Total
        Recv Q:         0       3       803
        Xmit Q:         0       30      6881'
!
!act# show conn
8 in use, 12 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:15288, idle 0:00:24, bytes 132, flags UIO
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Active-Active Failover

Task

- Configure ASA3 and ASA4 for active-active failover: use any IP address for standby unit:

Context	Interface1	Interface1 Nameif	Interface1 IP	Interface2	Interface2 Nameif	I
SW1-R3	Gi0/0.34	VLAN34	136.1.34.17/24	Gi0/1	VLAN93	13
SW2-R3	Gi0/0.34	VLAN34	136.1.34.18/24	Gi0/3	VLAN94	13

- Ensure that ASA3 always forwards traffic for **SW1-R3** and ASA4 for **SW2-R3** contexts.
- Configure authenticated stateful failover using Gi0/2 as the failover link with the name **FAILOVER** and the IP subnet 136.1.100.0/24.
- The ASA hostname should define failover role, context name, and state.
- Security-level for all interfaces maps the VLAN number.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config`

https://t.me/learningnets

instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

When configuring failover, it is mandatory to set both firewalls in either single or multiple context mode simultaneously with multiple contexts mode supporting a unique failover function known as Active/Active failover. With Active/Active failover, the primary unit is active for the first group of security contexts and standby for the second group, whereas the secondary unit is active for the second group of security contexts and standby for first the group; only two failover groups are supported because there are only two ASAs within a failover pair, and the admin context is always a member of group one. Thus, there are no longer purely active and purely standby units; both units forward traffic at the same time. Still, there are primary and secondary boxes, with respect to failover configuration, because the primary unit replicates all system context configurations to the secondary unit; within each user-configured context, configuration is replicated from the active unit to the standby unit.

To implement Active/Active failover, you configure one firewall unit as primary for the first group and secondary for the second group. After this, you assign firewall contexts to the groups. Every firewall will become active for the failover group where it is primary and standby for the group where it is secondary; if one of the units fails, its contexts will migrate to the working unit. Use the command **`failover group [1|2]`** to enter the firewall group configuration. Under the group configuration, you can issue the command **`primary`** or **`secondary`** to specify the firewall role for this group. Another important command here is **`preempt`** , which instructs each firewall to take over its assigned contexts when it becomes available again; by default, if a unit fails and the other unit takes over its contexts, the contexts are not returned back automatically when the original unit comes back. By entering the **`preempt`** command within the group, you enable the lower priority preemption behavior. Finally, to assign the context to a failover group, use the command **`join-failover-group [1|2]`** at the context level, when logged in to the system context. It is important to note that failover is group-wide; if a failover event occurs, the whole group fails over to the other unit, meaning all contexts assigned to the single group.

The last thing that differs for multiple-context mode is interface monitoring. By default, all physical interfaces mapped to a context and configured with the IP addresses are monitored. The monitored interface statistics are aggregated per-group; as soon as the number of failed interfaces exceeds the per-group threshold,

https://t.me/learningnets

the whole group fails over to the standby unit. The command to monitor an interface remains the same, `monitor-interface <nameif>` , but this time you apply it when logged under the particular context. Because of the group-wide failover behavior, the poll timers are now specific per group. To configure the timers, you must access the system context and enter the `failover group [1|2]` configuration mode; the commands remain the same and you can also set the interface-policy threshold, which applies to all failed interfaces within all contexts mapped to this group.

For Active/Active failover, ASA automatically generates virtual MAC addresses for all interface members of failover groups; for this reason, in the shared interface scenario, ASA can correctly assign traffic to each context and there is no need to manually generate MAC addresses or configure NAT.

Configuration

SW1:

```
interface range FastEthernet1/0/17 , FastEthernet1/0/19
 switchport mode access
 switchport access vlan 93
 no shutdown
!
interface range FastEthernet1/0/16 , FastEthernet1/0/18
 switchport mode access
 switchport access vlan 94
 no shutdown
```

SW2:

```
interface range FastEthernet1/0/16 , FastEthernet1/0/18
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
!
interface range FastEthernet1/0/17 , FastEthernet1/0/19
 switchport mode access
 switchport access vlan 100
 no shutdown
```

ASA3 System Context:

```
mode multiple
!
interface GigabitEthernet0/0
 no shutdown
!
```

https://t.me/learningnets

```
interface GigabitEthernet0/0.34
 vlan 34
!
interface GigabitEthernet0/1
 no shutdown
!
interface GigabitEthernet0/2
 no shutdown
!
interface GigabitEthernet0/3
 no shutdown
!
failover group 1
 primary
 preempt
!
failover group 2
 secondary
 preempt
!
failover lan unit primary
failover lan interface FAILOVER GigabitEthernet0/2
failover link failover GigabitEthernet0/2
failover interface ip FAILOVER 136.1.100.17 255.255.255.0 standby 136.1.100.18
failover key cisco
failover
!
admin-context SW1-R3
!
context SW1-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/1
 config-url disk0:/SW1-R3.cfg
 join-failover-group 1
!
context SW2-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/3
 config-url disk0:/SW2-R3.cfg
 join-failover-group 2
!
prompt priority context state
```

ASA3 SW1-R3 Context:

```
changeto context SW1-R3
!
```

https://t.me/learningnets

```
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.17 255.255.255.0 standby 136.1.34.178
!
interface GigabitEthernet0/1
 nameif VLAN93
 security-level 93
 ip address 136.1.93.17 255.255.255.0 standby 136.1.93.178
!
monitor-interface VLAN34
```

ASA3 SW2-R3 Context:

```
changeto context SW2-R3
!
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 ip address 136.1.34.18 255.255.255.0 standby 136.1.34.187
!
interface GigabitEthernet0/3
 nameif VLAN94
 security-level 94
 ip address 136.1.94.18 255.255.255.0 standby 136.1.94.187
!
monitor-interface VLAN34
```

ASA4:

```
interface GigabitEthernet0/2
 no shutdown
!
failover lan unit secondary
failover lan interface FAILOVER GigabitEthernet0/2
failover link failover GigabitEthernet0/2
failover interface ip FAILOVER 136.1.100.17 255.255.255.0 standby 136.1.100.18
failover
```

Verification

Verify failover from system context:

```
pri/act# show failover group 1

  Last Failover at: 21:45:36 UTC Nov 18 2013
```

https://t.me/learningnets

```
 This host:    Primary
State:          Active
                Active time:    30542 (sec)

                SW1-R3 Interface VLAN34 (136.1.34.17): Normal (Monitored)
                SW1-R3 Interface VLAN93 (136.1.93.17): Normal (Monitored)

 Other host:   Secondary
State:          Standby Ready
                Active time:    0 (sec)

                SW1-R3 Interface VLAN34 (136.1.34.178): Normal (Monitored)
                SW1-R3 Interface VLAN93 (136.1.93.118): Normal (Monitored)

  Stateful Failover Logical Update Statistics
        Status: Configured.
        RPC services    0          0          0          0
        TCP conn        9          0          0          0
        UDP conn        0          0          0          0
        ARP tbl         6          0          0          0
        Xlate_Timeout   0          0          0          0
        IPv6 ND tbl     0          0          0          0
        SIP Session     0          0          0          0
        Route Session   0          0          0          0
        User-Identity   4          0          0          0
!
!pri/act# show failover group 2

  Last Failover at: 21:47:26 UTC Nov 18 2013
 This host:    Primary
State:          Standby Ready
                Active time:    109 (sec)

                SW2-R3 Interface VLAN34 (136.1.34.187): Normal (Monitored)
                SW2-R3 Interface VLAN94 (136.1.94.187): Normal (Monitored)

 Other host:   Secondary
State:          Active

                Active time:    30463 (sec)

                SW2-R3 Interface VLAN34 (136.1.34.18): Normal (Monitored)
                SW2-R3 Interface VLAN94 (136.1.94.18): Normal (Monitored)

  Stateful Failover Logical Update Statistics
        Status: Configured.
```

https://t.me/learningnets

```
        RPC services    0          0          0          0
        TCP conn        0          0          7          0
        UDP conn        0          0          0          0
        ARP tbl         0          0          3          0
        Xlate_Timeout   0          0          0          0
        IPv6 ND tbl     0          0          0          0
        SIP Session     0          0          0          0
        Route Session   0          0          0          0
        User-Identity   1          0          0          0
```

Verify failover from user configured context:

```
pri/SW1-R3/act# show failover
Failover On
Last Failover at: 21:45:36 UTC Nov 18 2013 This context: Active
                Active time: 30286 (sec)
                  Interface VLAN34 (136.1.34.17): Normal (Monitored)
                  Interface VLAN93 (136.1.93.17): Normal (Monitored) Peer context: Standby Ready
                Active time: 0 (sec)
                  Interface VLAN34 (136.1.34.178): Normal (Monitored)
                  Interface VLAN93 (136.1.93.118): Normal (Monitored)

Stateful Failover Logical Update Statistics
        Status: Configured.
        Stateful Obj    xmit       xerr       rcv        rerr
        RPC services    0          0          0          0
        TCP conn        9          0          0          0
        UDP conn        0          0          0          0
        ARP tbl         6          0          0          0
        Xlate_Timeout   0          0          0          0
        IPv6 ND tbl     0          0          0          0
        SIP Session     0          0          0          0
        Route Session   0          0          0          0
        User-Identity   2          0          0          0
!
!sec/SW2-R3/act# show failover
Failover On
Last Failover at: 21:47:24 UTC Nov 18 2013 This context: Active
                Active time: 30198 (sec)
                  Interface VLAN34 (136.1.34.18): Normal (Monitored)
                  Interface VLAN94 (136.1.94.18): Normal (Monitored) Peer context: Standby Ready

                Active time: 109 (sec)
                  Interface VLAN34 (136.1.34.187): Normal (Monitored)
                  Interface VLAN94 (136.1.94.187): Normal (Monitored)
```

https://t.me/learningnets

```
Stateful Failover Logical Update Statistics
        Status: Configured.
        Stateful Obj    xmit       xerr       rcv        rerr
        RPC services    0          0          0          0
        TCP conn        7          0          0          0
        UDP conn        0          0          0          0
        ARP tbl         4          0          0          0
        Xlate_Timeout   0          0          0          0
        IPv6 ND tbl     0          0          0          0
        SIP Session     0          0          0          0
        Route Session   0          0          0          0
        User-Identity   0          1          1          0
```

Check that virtual MAC addresses are automatically assigned to both ASAs; for example, look at VLAN 34:

```
pri/act# show interface gigabitEthernet 0/0 | i MAC
        MAC address a493.4caa.f508, MTU not set
!
!sec/stby# show interface gigabitEthernet 0/0 | i MAC
        MAC address a493.4caa.f78a, MTU not set
!
!R3#show ip arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  136.1.34.3              -   001e.f779.4770  ARPA   FastEthernet0/0
Internet  136.1.34.17            29   1200.0022.0200  ARPA   FastEthernet0/0
Internet  136.1.34.18            29   1200.0022.0300  ARPA   FastEthernet0/0
```

Initiate traffic from SW1 toward R3:

```
SW1#telnet 136.1.34.3
Trying 136.1.34.3 ... Open



User Access Verification


Password: R3>
!
!pri/SW1-R3/act# show conn
5 in use, 5 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:15280, idle 0:00:20, bytes 111, flags UIO
!
!pri/SW1-R3/act# failover exec mate show conn
```

https://t.me/learningnets

```
5 in use, 5 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:15280, idle 0:00:21, bytes 111, flags UIO
```

Initiate traffic from SW2 toward R3:

```
SW2#telnet 136.1.34.3
Trying 136.1.34.3 ... Open
c


User Access Verification


Password: R3>
!
!sec/SW2-R3/act# show conn
5 in use, 5 most used
TCP VLAN34 136.1.34.3:23 VLAN94 136.1.94.10:36324, idle 0:00:12, bytes 111, flags UIO
!
!sec/SW2-R3/act# failover exec mate show conn
5 in use, 5 most used
TCP VLAN34 136.1.34.3:23 VLAN94 136.1.94.10:36324, idle 0:00:10, bytes 111, flags UIO
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Traffic Policing

Task

- To mitigate ICMP flooding attacks, configure ASA1 to limit ICMP traffic to 64 Kbps on its VLAN 59 interface.
- Ensure that both ingress and egress traffic flows conform to this restriction.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Traffic policing is a QoS feature that enforces average traffic rate by dropping exceeding packets. The main purpose of policing is to ensure that traffic flows meet the contracted rates. Another useful purpose is to limit the rate of potentially dangerous traffic, such as ICMP flood, to prevent the potential or actual DoS attack. Policing applies to traffic matched by a specific L3/L4 class-map. You specify policing under a policy-map using the syntax:

https://t.me/learningnets

```
policy-map POLICY
  class TRAFFIC
   police [inbound|outbound] cir [burst]
```

Policing applies either ingress or egress on the interface. Here, CIR (Committed Access Rate) is the average rate enforced in Bps (bits per second). The Burst value is measured in bytes and defines the maximum amount of traffic that could be sent/received over the interface in a given unit of time. This value is not mandatory, and you may omit it; the firewall picks up the optimal value automatically. The larger the burst, the more averaging the policer does when metering the traffic, allowing for sudden spikes of traffic rate. You can use the value Tc=Burst/CIR as the reference averaging interval; over this interval the average rate is always CIR.

You may specify actions for conforming or exceeding traffic using the command syntax **`police cir [burst] conform-action [transmit|drop] exceed-action [transmit|drop] .`** Although the ASA cannot mark or remark traffic using the DSCP field, you can use these actions in two scenarios:

- Dropping all packets matching the specific class, such as **`police 64000 8000 conform-action drop`** ; this could be viewed as alternative to access-list for dropping traffic; although this option is available, it will not work, in that traffic will actually still be allowed but policed to 64Kbps.
- Checking whether the traffic rate conforms to the contract without affecting traffic flow (that is, without dropping any packets), such as **`police 128000 64000 conform-action transmit exceed-action transmit`** ; although this option is available, it will not work, in that traffic will actually still be policed to 64 Kbps.

Policing could be configured in a policy-map that applies either globally or per-interface. When the policy applies globally, it affects all traffic ingress and/or egress on all interfaces, depending on the direction of the feature, whereas if applied per interface it only affects ingress and/or egress traffic on that interface. Remember that the policer drops packets exceeding the average traffic rate which might severely affect TCP performance, because this protocol reacts to packet drops by slowing sending the rate. To avoid this effect, try setting the burst to a larger value, based on the Tc of 1.5-3 seconds of CIR (that is, Burst = CIR*3/8); however, the optimal value is usually picked up empirically.

Traffic destined to the ASA is not affected by the policing configuration, and this is because ASA internal architecture makes a clear difference between data-plane,

https://t.me/learningnets

control-plane, and management-plane traffic and supports specific actions for each; policing only affects data-plane traffic, therefore traffic transiting the ASA. Additionally, input policing only matches on return traffic, it does not match as traffic is initiated.

Configuration

```
ASA1:

access-list ICMP permit icmp any any
!
class-map ICMP
 match access-list ICMP
!
policy-map ICMP
 class ICMP
  police input 64000
  police output 64000
!
service-policy ICMP interface VLAN59
```

Verification

Check policy-map statistics before generating any ICMP traffic:

```
ASA1# show service-policy interface VLAN59

Interface VLAN59:
  Service-policy: ICMP
    Class-map: ICMP Input police Interface VLAN59:
cir 64000 bps, bc 2000 bytes
conformed 0 packets, 0 bytes; actions:  transmit
        exceeded 0 packets, 0 bytes; actions:  drop
        conformed 0 bps, exceed 0 bps Output police Interface VLAN59:
cir 64000 bps, bc 2000 bytes
conformed 0 packets, 0 bytes; actions:  transmit

        exceeded 0 packets, 0 bytes; actions:  drop
        conformed 0 bps, exceed 0 bps
```

https://t.me/learningnets

Generate traffic from R2 to exceed the policing rate:

```
R2#ping 136.1.49.1 size 1500 repeat 100
Type escape sequence to abort.
Sending 100, 1500-byte ICMP Echos to 136.1.49.1, timeout is 2 seconds:
!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!!!!!.
!!!!!!.!!!!!!.!!!!!!.!!!!!!.!!Success rate is 86 percent (86/100), round-trip min/avg/max = 1/1/4 ms
!
!ASA1#show service-policy interface VLAN59

Interface VLAN59:
  Service-policy: ICMP
    Class-map: ICMP
      Input police Interface VLAN59:
        cir 64000 bps, bc 2000 bytes
        conformed 0 packets, 0 bytes; actions:  transmit
        exceeded 0 packets, 0 bytes; actions:  drop
        conformed 0 bps, exceed 0 bps
      Output police Interface VLAN59:
        cir 64000 bps, bc 2000 bytes conformed 86 packets, 130204 bytes; actions:  transmit
exceeded 14 packets, 21196 bytes; actions:  drop

        conformed 0 bps, exceed 0 bps
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Low Latency Queuing

Task

- Configure ASA1 with priority queue service to VoIP traffic going to VLAN 49 users.
- Classify the VoIP packets based on RTP port range 16384 32767.
- Set priority queue depth to 5 packets on VLAN 49 interface.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Queuing is used to buffer outgoing packets on firewall interfaces when the links are too congested to accept the traffic. There are two levels of queuing: hardware queue (tx-ring or transmit ring) and software queue, managed by the firewall. The tx-ring is emptied by the hardware, and when this queue is filled up, the packets will go to the software queue. The only supported discipline for the tx-ring is FIFO (first in first out), because this queue is usually serviced pretty fast, almost at the line rate.

The software queue may hold packets when the physical interface is severely congested. By default, this queue is serviced in the same FIFO manner as the tx-ring. However, certain types of IP traffic, most notably VoIP bearer packets, cannot

https://t.me/learningnets

tolerate long delays caused by queuing in software. This is why the firewall supports the so-called priority or low-latency queuing (LLQ). Priority software ensures that certain classes of traffic assigned to this queue are always serviced first and deposited into tx-ring ahead of any other packets.

You may enable priority queue per-interface using the command `priority-queue [nameif]` in global configuration mode. Priority queuing can be enabled only at the physical interface level and affects all associated sub-interfaces. This will put you in the priority-queue configuration mode, where you can further set tx-ring size (hardware queue size in packets) or queue-limit (priority queue depth, in packets). You may want to lessen the priority-queue depth to limit the effect of priority queue on other traffic—that is, to prevent bandwidth starvation for non-priority traffic; however, this might cause excessive packet drops in the priority queue. In addition to the LLQ, the interface will still have Best Effort (BE) queue assigned that services non-priority traffic (after the priority queue has been emptied).

To assign actual traffic to the priority queue, you must configure and apply a policy-map to the respective interface. The command syntax is:

```
policy-map POLICY_NAME
 class CLASS_NAME
   priority
```

Traffic matching the L3/L4 class-map is serviced using the expedite (priority) queue on the interface; all other traffic will use the Best Effort queue. The policy-map could be applied per-interface or globally. In the latter case, it will affect every interface enabled for priority-queue service. It is common to enable priority-queuing for sensitive but bandwidth-limited traffic such as VoIP; however, avoid putting bursty traffic, such as file transfers, in the priority queue, because this might starve other traffic classes. You cannot police and prioritize the same flow of traffic. When assigning VoIP traffic to priority queue, you might want to do the following:

- Police non-prioritized traffic classes to make sure they don’t fill the tx-ring and affect VoIP quality.
- Tune tx-ring size (which is defaulted to 128) to make it shorter. This will force the software priority queue to be engaged faster and might reduce the VoIP latency. However, setting the tx-ring too small might affect the overall performance because of constant calls for the LLQ scheduler.

Note that the solution below uses a special `match rtp` command to classify VoIP bearer traffic. When using Cisco VoIP solutions, it is common to see VoIP traffic using RTP port range 16384 32767. However, the `match rtp` command takes two

https://t.me/learningnets

parameters: initial port, and the step—that is, the number of ports to step up to cover the RTP port range; so the command `match rtp 16384 16383` actually covers the default VoIP port range.

Configuration

```
ASA1:

class-map LLQ
 match rtp 16384 16383
!
policy-map LLQ
 class LLQ
  priority
!
priority-queue VLAN49
 queue-limit 5
!
service-policy LLQ interface VLAN49
```

Verification

Check priority queuing configuration:

```
ASA1# show priority-queue config
Priority-Queue Config interface VLAN49
                current         default         range
queue-limit     5               2048            0 - 2048

tx-ring-limit   512             512             3 - 512

Priority-Queue Config interface VLAN59
                current         default         range
queue-limit     0               2048            0 - 2048
tx-ring-limit   -1              512             3 - 512
```

Check priority queue statistics, where BE is the Best Effort Queue:

```
ASA1# show priority-queue statistics
Priority-Queue Statistics interface VLAN49
Queue Type          = BE
```

https://t.me/learningnets

```
Tail Drops          = 0
Reset Drops         = 0
Packets Transmit    = 0
Packets Enqueued    = 0
Current Q Length    = 0
Max Q Length        = 0
Queue Type          = LLQ
Tail Drops          = 0
Reset Drops         = 0
Packets Transmit    = 0
Packets Enqueued    = 0
Current Q Length    = 0
Max Q Length        = 0
!
!ASA1# show service-policy interface VlAN49

Interface VLAN49:
  Service-policy: LLQ
    Class-map: LLQ
      Priority: Interface VLAN49: aggregate drop 0, aggregate transmit 0

    Class-map: class-default
        Default Queueing
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Traffic Shaping

Task

- The VLAN 49 interface of ASA1 connects to an ISP that provides only 512 Kbps of guaranteed traffic rate (CIR).
- Configure the firewall to conform to this requirement, provided that the ISP sets the measurement interval to 100ms.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

https://t.me/learningnets

Traffic shaping is a QoS procedure that formats traffic flow according to the traffic rate specified. This commonly happens when the physical line rate (Access Information Rate, or AIR, such as 100 Mbps) is higher than the contracted average transmission rate (for example, 1 Mbps). The shaper meters the traffic flow rate and delays packets that exceed the contracted speed (CIR, or Committed Information Rate). The delayed packets are buffered and processed every timer interval Tc (committed measurement interval in milliseconds). Every interval Tc, the shaper will emit no more than CIR*Tc bits at physical line rate AIR.

Because AIR is commonly higher than CIR, the firewall will only emit traffic for the duration of CIR/AIR*Tc ms and then remain idle until the end of the current interval. The traffic that is still in queue will have to wait for the beginning of the next interval to be serviced. The CIR/AIR*Tc value might be viewed as the coefficient that defines the shaper effectiveness. Thus, by lowering the Tc you may increase the overall interval utilization but at same time increase the load on the firewalls CPU, because the queue scheduler is required to fire more often. So what is the optimal value of Tc? Commonly, this value is specified in the traffic contract as the measurement interval. You should set your shaper Tc value no larger than the contracted Tc to avoid excessive drops on the provider's edge. However, some types of traffic, such as VoIP bearer, might require the Tc value to be set as low as possible to minimize delays introduced by the shaper. This is acceptable, as long as the firewall CPU is not highly over-utilized.

Now think of the following situation: What if during the current Tc there is no traffic to send, but in the next Tc the queue accumulates more than Bc bits of traffic? Using just the regular rules, the shaper will underutilize the current interval and send only Bc bits in the next interval. The average rate per two intervals will be Bc/(2*Tc) = CIR/2. To minimize this unfairness, the shaper implements a special credit counter that is incremented every time the current interval is underutilized. For example, if the current interval did not send any traffic, the credit counter is incremented by Bc. However, there is a limit to this counter, called Be (Burst Excessive), the maximum amount of extra credit allowed to the flow. Now, every Tc the shaper is allowed to send up to Bc+Credit bits, that is up to Bc+Be if the credit allows. This procedure allows the underutilization during any interval to allow sending extra traffic during the next interval. The ASA firewall automatically sets Be=Bc to allow for more fair shaping.

When configuring traffic shaping in the ASA firewall, consider the following:

- You may only apply shaping at the interface level (not globally) and only under class-default; no other classes can be defined in the policy that performs traffic-shaping.
- Shaping applies to both transit and firewall-originated traffic at the same time.
- The interface-level priority queue does not work along with shaping. However, you

https://t.me/learningnets

may enable a hierarchical queue inside the shaper, as demonstrated in a separate following task.

- Shaping accounts for full packet size, including IPsec encapsulation and layer 2 overheads.
- The shaping queue size is 64 packets, and the service discipline is FIFO by default (it can be changed with hierarchical queueing).

The syntax for the shape command is:

```
policy-map POLICY_NAME
  class class-default
    shape average <rate> <burst>
```

Here, Tc is not set explicitly, but rather calculated by the shaper using the formula Tc=Burst/Rate. Note that Burst is set in bits, not bytes as with the policing.

Configuration

```
ASA1:
policy-map SHAPER
 class class-default
  shape average 512000 51200
!
service-policy SHAPER interface VLAN49
```

Verification

Verify the configuration:

https://t.me/learningnets

```
ASA1# show service-policy shape


Interface VLAN49:
  Service-policy: SHAPER
    Class-map: class-default
 shape (average) cir 512000, bc 51200

      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0
```

Initiate traffic from R2 toward R1, and check the average latency:

```
R2#ping 136.1.49.1 size 1500 repeat 100
Type escape sequence to abort.
Sending 100, 1500-byte ICMP Echos to 136.1.49.1, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!Success rate is 100 percent (100/100), round-trip min/avg/max = 1/21/24 ms
!
!ASA1# show service-policy interface VLAN49

Interface VLAN49:
  Service-policy: SHAPER
    Class-map: class-default

      shape (average) cir 512000, bc 51200
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0 (pkts output/bytes output) 100/150000
```

Remove the shaping configuration, initiate traffic from R2 toward R1 again, and check the average latency:

```
ASA1:
no service-policy SHAPER interface VLAN49
!
!R2#ping 136.1.49.1 size 1500 repeat 100
Type escape sequence to abort.
Sending 100, 1500-byte ICMP Echos to 136.1.49.1, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

https://t.me/learningnets

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!! Success rate is 100 percent (100/100), round-trip min/avg/max = 1/2/4 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Hierarchical Queuing

Task

- Configure priority queuing for shaped VoIP bearer and GETVPN signaling traffic on ASA1:
  - GETVPN signaling is defined as GDOI exchange on the default port.
  - VoIP bearer traffic is marked with the DSCP value of EF.
- All other traffic should receive best-effort service.
- Adjust traffic-shaping interval to provide minimum delay for VoIP traffic.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

When you apply shaping per-interface, you create another software queue, the one that the shaper uses for the delayed traffic. Essentially, this queue represents the new virtual link with the speed equal to the contracted rate (the CIR). It would be beneficial to segregate traffic flows within this virtual link and provide differentiated services. As you remember, enabling shaping on the interface will disallow the use of any interface-level priority queue. However, you are allowed to create the priority

https://t.me/learningnets

queue within the shaper's delayed traffic buffer. To accomplish this, you must nest another service policy (child policy) under the shaped class-default of the parent policy-map.

Under the child policy, you may assign L3/L4 classes and enable priority-queue where needed. The result is that the shaper's queue is split in two queues: LLQ (priority) and BE (best effort). Notice that you cannot apply anything (such as policing) except the `priority` command under the child policy. For example:

```
class-map VOICE
 match dscp ef
!
policy-map CHILD_POLICY
  class VOICE
    priority
!
policy-map PARENT_POLICY
  class class-default
    shape average 256000 2560
    service-policy CHILD_POLICY
```

When implementing priority queuing for VoIP traffic, you may want to set Bc=CIR*10ms to minimize traffic delays.

Configuration

ASA1:

```
class-map VPN
 match port udp eq 848
!
class-map VOICE
 match dscp ef
!
policy-map CHILD_POLICY
 class VPN
  priority
 class VOICE
  priority
!
policy-map SHAPER
 class class-default
  shape average 512000 5120
```

https://t.me/learningnets

```
  service-policy CHILD_POLICY
!
service-policy SHAPER interface VLAN49
```

Verification

Verify configuration:

```
ASA1# show service-policy interface VLAN49

Interface VLAN49:
  Service-policy: SHAPER
    Class-map: class-default
 shape (average) cir 512000, bc 5120

      (pkts output/bytes output) 0/0
      (total drops/no-buffer drops) 0/0

      Service-policy: CHILD_POLICY Class-map: VPN
 priority

          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0
 Class-map: VOICE
 priority


          Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0

        Class-map: class-default

          Default Queueing
          queue limit 64 packets
          (queue depth/total drops/no-buffer drops) 0/0/0
          (pkts output/bytes output) 0/0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA URL Filtering

Task

- Configure ASA2 to filter ActiveX and JavaScript from all HTTP requests on port 80.
  - Use UDP to communicate with the websense server located at 172.16.10.100.
  - Filter HTTP traffic from 136.1.19.0/24 network on ports 80 and 8080, except from host 136.1.19.100; block proxy-requests going on port 8080.
  - Filter FTP traffic from network 136.1.19.0/24 and deny interactive FTP connections.
  - In case of websense server failure, HTTP/FTP requests should be allowed.
- For best performance, configure ASA to remember decisions made by the filtering server, regardless of the user requesting access.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The firewall may work in tandem with external N2H2 or Websense URL filtering server to perform application-level access-control. When a user accesses the web,

https://t.me/learningnets

the firewall sends the issued HTTP or HTTPS request to the original destination and the filtering service at the same time. The response from the original server is only allowed if the filtering service allows it. The firewall is also capable of filtering the requested FTP files, by redirecting the FTP GET command contents to the filtering service. Thus, the content filtering engine may perform screening based on the object names.

You start configuring URL filtering with the command `url-server`, which defines the filtering server IP address and parameters. After this, you may use commands `filter ftp`, `filter url`, and `filter https` to define the traffic that should be inspected. The generic syntax is as follows: `filter [ftp|url|https] <port> <src-net> <src-mask> <dst-net> <dst-mask> [allow] [proxy-block]`. By using zero instead of address and mask you effectively simulate the **any** keyword. You can also inspect URL/FTP requests on any port, even non-default; for example, you may filter all HTTP requests sent via a Squid proxy server listening on port 8080 using the command `filter url 8080 0 0 0 0 proxy-block`. The **proxy-block** keyword denies any proxy HTTP requests. Now pay attention to the **allow** keyword; when you configure a filtering rule with this keyword, the firewall will permit all requests if it detects failure of the filtering service. By default, when the filtering service fails, all HTTP requests are denied. The special **interact-block** keyword is used to prevent interactive FTP sessions initiated by users.

Because the firewall sends the original user's HTTP/HTTPS request to the destination server and the filtering service at the same time, it may happen that under heavy load, the response from the server arrives before the filtering decision is made. By default, the firewall would drop the content. However, you may use the command `url-block block <buffer-size>` to set the buffer that stores the server's response while waiting for filtering decision. This may improve performance with busy filtering services. Additionally, the ASA can remember the response of the URL filtering server and cache it (based on the destination and/or source IP address of the session), so when an identical request is received, it will no longer query the server but allow or drop the session directly.

There is also a special command `filter https`, which enforces HTTPs traffic inspection. Of course, the firewall cannot look inside the encrypted sessions. However, it queries the filtering service about the destination IP address used to establish the session and some other open attributes. If the service denies the request, the SSL handshake never completes. Finally, the commands `filter activex` and `filter java` have nothing to do with the URL filtering process, they just specify the sessions (based on source and/or destination) that are not allowed to download ActiveX objects or Java applets; the firewall simply removes those objects from the HTTP data stream if the source/destination pair matches the configuration. You can look for Java/ActiveX components on any port, not just the default HTTP

https://t.me/learningnets

port 80.

Configuration

```
ASA2:

url-server (VLAN26) vendor websense host 172.16.10.100 protocol udp version 4
url-cache dst 128
!
filter activex 80 0 0 0 0
filter java 80 0 0 0 0
!
filter url 80 136.1.19.0 255.255.255.0 0 0 allow
filter url 8080 136.1.19.0 255.255.255.0 0 0 allow proxy-block
filter url except 136.1.19.100 255.255.255.255 0 0
filter ftp 21 136.1.19.0 255.255.255.0 0 0 allow interact-block
```

Verification

Check the URL server statistics, but note that there is no websense server installed:

```
ASA2# show url-server statistics

Global Statistics:
--------------------
URLs total/allowed/denied         0/0/0
URLs allowed by cache/server      0/0
URLs denied by cache/server       0/0
HTTPSs total/allowed/denied       0/0/0
HTTPSs allowed by cache/server    0/0
HTTPSs denied by cache/server     0/0
FTPs total/allowed/denied         0/0/0
FTPs allowed by cache/server      0/0
FTPs denied by cache/server       0/0
Requests dropped                  0
Server timeouts/retries           0/0
Processed rate average 60s/300s   0/0 requests/second
Denied rate average 60s/300s      0/0 requests/second
Dropped rate average 60s/300s     0/0 requests/second

Server Statistics:
--------------------172.16.10.100                      DOWN
```

https://t.me/learningnets

```
Vendor                          websense
 Port                           15868

  Requests total/allowed/denied   0/0/0
  Server timeouts/retries         0/0
  Responses received              0
  Response time average 60s/300s  0/0

URL Packets Sent and Received Stats:
------------------------------------
Message                 Sent    Received
STATUS_REQUEST          215     0
LOOKUP_REQUEST          0       0
LOG_REQUEST             0       NA

Errors:
-------
RFC noncompliant GET method     0
URL buffer update failure       0
```

Verify the that filtering configuration is correct:

```
ASA2# show running-config filter

filter url except 136.1.19.100 255.255.255.255 0.0.0.0 0.0.0.0
filter activex 80 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
filter java 80 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
filter url http 136.1.19.0 255.255.255.0 0.0.0.0 0.0.0.0 allow
filter url 8080 136.1.19.0 255.255.255.0 0.0.0.0 0.0.0.0 allow proxy-block
filter ftp 21 136.1.19.0 255.255.255.0 0.0.0.0 0.0.0.0 allow interact-block
```

To verify that filtering will apply to traffic from VLAN 19, use the **packet-tracer** tool:

```
ASA2# packet-tracer input VLAN19 tcp 136.1.19.20 20000 136.1.29.2 80


Phase: 1
Type: ROUTE-LOOKUP
Subtype: input
Result: ALLOW
Config:
Additional Information:
in   136.1.29.0      255.255.255.0   VLAN29

Phase: 2
```

https://t.me/learningnets

```
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 3 Type: FILTER
Subtype: filter-activex
Result: ALLOW
Config:
Additional Information:

Phase: 4 Type: FILTER
Subtype: filter-java
Result: ALLOW
Config:
Additional Information:

Phase: 5 Type: FILTER
Subtype: filter-url

Result: ALLOW
Config:
Additional Information:

Phase: 6
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Config:
Additional Information:

Phase: 7
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Config:
Additional Information:
New flow created with id 9866, packet dispatched to next module

Result:
input-interface: VLAN19
input-status: up
input-line-status: up
output-interface: VLAN29
output-status: up
```

https://t.me/learningnets

```
output-line-status: up
Action: allow
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Network Time Protocol

Task

- Configure the ASA1 for time synchronization via NTP with R2.
- For added security, authenticate NTP updates using the string **CISCO**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

NTP is used for time synchronization and is very important when implementing distributed logging or certificate-based authentication. The firewall may work only as NTP client, synchronizing with one or multiple servers; most importantly, you can authenticate NTP time updates using MD5 hash. The configuration is straightforward; just remember that you must enable authentication and trust the authentication key in the client device only. The server only needs the key defined, with the correct key index; indexes are carried along in NTP messages so that devices might select the proper keys for hash computations.

https://t.me/learningnets

Configuration

```
ASA1:
ntp authentication-key 2 md5 CISCO
ntp authenticate
ntp server 136.1.59.2 key 2
ntp trusted-key 2

R2:

ntp master
ntp authentication-key 2 md5 CISCO
ntp trusted-key 2
```

Verification

Depending on the clock difference, it may take a while for time to get synchronized between client and server:

https://t.me/learningnets

```
ASA1# show ntp status
Clock is synchronized, stratum 9, reference is 136.1.59.2

nominal freq is 99.9984 Hz, actual freq is 100.0195 Hz, precision is 2**6
reference time is d6387050.651cd587 (11:43:12.394 UTC Thu Nov 21 2013)
clock offset is -36.0324 msec, root delay is 0.76 msec
root dispersion is 58.99 msec, peer dispersion is 22.54 msec
!
!ASA1# show  ntp associations detail

136.1.59.2 configured, authenticated, our_master, sane, valid, stratum 8

ref ID 127.127.1.1, time d6387043.86592642 (11:42:59.524 UTC Thu Nov 21 2013)
our mode client, peer mode server, our poll intvl 1024, peer poll intvl 1024
root delay 0.00 msec, root disp 0.41, reach 377, sync dist 23.331
delay 0.76 msec, offset -36.0324 msec, dispersion 22.54
precision 2**21, version 3
org time d6387050.5bca6763 (11:43:12.358 UTC Thu Nov 21 2013)
rcv time d6387050.651cd587 (11:43:12.394 UTC Thu Nov 21 2013)
xmt time d6387050.64e8d355 (11:43:12.394 UTC Thu Nov 21 2013)
filtdelay =     0.76    0.75    0.76    0.76    0.78    0.78    0.75    0.75
filtoffset =  -36.03  -32.51  -28.95  -25.19  -21.38  -17.51  -13.64  -11.78
filterror =    15.63   31.25   46.88   62.48   78.11   93.73  109.36  117.17
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA System Logging

Task

- Configure ASA2 to generate timestamped system logging messages.
- Collect the debugging messages in the system memory buffer and limit buffer size to 65536 bytes.
- Save the memory buffer contents when it wraps to the FTP server 172.16.10.100 using the username **cisco** and the password **cisco**.
- Send informational and higher priority messages to the syslog server at 172.16.10.100 using the facility value of 23.
- The console port should receive only the alerts and higher messages.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

https://t.me/learningnets

The firewall lets you collect system logging information to aid in monitoring and troubleshooting; logging features of the ASA firewall are sophisticated, and it takes time to get fully used to them. Log messages are generated by various system components and have the following attributes:

- ID number, which uniquely identifies this type of the messages; every message has its own ID, and you can set the device ID (used for remote logging) by using the command `logging device-id .`
- Severity, which varies from 0 (emergencies, most important) to 7 (debugging, least important).
- Class, which defines the functional area of the firewall that originated the message (such as VPN or HA).
- Timestamp, assigned if you enabled timestamps using the command `logging timestamp` in global configuration mode.

Log messages can be saved to various destinations. When you activate logging using the command `logging enable` , the firewall will start sending messages to all configured destinations. Before a message is sent to all configured destinations, it is queued in the internal message queue, with the size defined by the command `logging queue <messages>` . You may configure the following destinations for syslog messages:

- Syslog server, an external host that collects messages sent over the syslog protocol. You configure this destination using the command `logging host <nameif> <ip_address> [tcp/[port]] [udp/[port]]` . Using UDP for syslog messages delivery is the default, but you can change it to reliable TCP; the default UDP port is 514 and the default TCP port is 1470.

  When you configure TCP logging, your firewall might block all connections if the information about the connections is being logged; this happens when the logging host is down and the firewall cannot make a reliable connection to it. To permit connections across the firewall even with the syslog host down, use the permit-hostdown option to the logging host command. To filter the messages sent to the syslog server based on their importance, use the command `logging trap <level>` , where level is between 0 and 7, to specify the cut-off level for the syslog messages. Only message of the level priority or more important are being recorded to this destination.

  To allow the remote system to distinguish the type of the device sending the syslog message, you can change the logging facility value from the

https://t.me/learningnets

default 20. Facility is a UNIX concept, with some of the default facility values being kernel (0), user (1), and mail (2). The default facility value 20 maps to the range of reserved values 16-23 that corresponds to facility names local0-local7, which stands for locally-significant facility names. In our task, we set the syslog facility to local7. Note that the syslog destination supports secure logging over SSL; to enable this, use TCP for syslog connection along with the keyword secure; for example, **`logging host 10.0.0.100 tcp/1470 secure .`**

- Console logging, which is enabled using the command **`logging console <level>`** . Be careful with it if you are managing the appliance using the console port, because a busy device might generate a large number of messages. Consider using message filtering (described below) or rate-limiting the messages in such cases.
- Email logging, which allows you to send certain message to a specific email address. You must specify the severity level for this destination and configure the email settings, such as recipient, sender, and SMTP server. For example (note that you can set separate severities per recipient and specify multiple recipients):

```
logging mail 0
logging from-address ccie@ine.com
logging recipient-address admin@ine.com [severity]
smtp-server 172.16.10.100
```

- ASDM logging, which stores and sends messages to the ASDM management console and is enabled using the command **`logging asdm <severity>`** . ASDM messages are stored in a separate buffer, with the size defined by using the command **`logging asdm-buffer-size <number_of_messages> .`**
- Logging to selected remote Telnet/SSH session (monitor logging); use the command **`logging monitor <severity>`** to activate this destination. Note that remote users must issue the command **`terminal monitor`** to start receiving the logging messages.
- Logging to the memory buffer; this allows you to save syslog messages in the appliance’s memory, avoiding the hurdle of messages on the console screen. Activate this destination using the command **`logging buffered <level>`**

https://t.me/learningnets

; you can tune the buffer size using the command **logging buffer-size <bytes>** . To view the log buffer, use the command **show logging** , and to clear its contents, use the command **clear logging buffer** . Note that the firewall will overwrite the memory buffer when it exhausts the memory allocated to it. You can configure the firewall to save the old contents to the flash memory every time the buffer wraps using the command **logging flash-bufferwrap** . You can also save the buffer to the flash memory at any time using the command **logging savelog <filename>** . The maximum amount of flash memory available to buffer saves is set using the command **logging flash-maximum-allocation <kilobytes>** . Alternatively, the contents of the memory buffer could be saved using a remote FTP server and the command **logging ftp-bufferwrap** . To define the remote FTP server parameters, use the command **logging ftp-server <server> <path> <username> <password> .**

Configuration

ASA2:

```
logging enable
!
logging timestamp
logging buffer-size 65536
logging console alerts
logging buffered debugging
logging trap informational
logging facility 23
logging host VLAN26 172.16.10.100
!
logging ftp-bufferwrap
logging ftp-server 172.16.10.100 / cisco cisco
```

Verification

Check the logging configuration:

```
ASA2# show logging setting
Syslog logging: enabled
Facility: 23
```

https://t.me/learningnets

```
Timestamp logging: enabled

    Standby logging: disabled
    Debug-trace logging: disabled Console logging: level alerts, 0 messages logged
    Monitor logging: disabled Buffer logging: level debugging, 40222 messages logged
Trap logging: level informational, facility 23, 17 messages logged
Logging to VLAN26 172.16.10.100 errors: 2  dropped: 7

    Permit-hostdown logging: disabled
    History logging: disabled
    Device ID: disabled
    Mail logging: disabled
    ASDM logging: disabled
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Log Filtering

Task

- Configure logging on ASA1:
  - Ensure that the message with the ID 103012 is never sent to any destination.
  - Limit the messages visible to the users on the console to VPN-related errors or more severe information.
  - The buffer should only receive messages with IDs in the range 101002-105044.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Filtering system log message is important in situations where you are overwhelmed with the amount of collected data. The ASA firewall provides you with the following options to filter the logging messages:

- Setting the severity for a particular destination.
-

https://t.me/learningnets

Class-based filtering, using the command `logging class <class> <destination> <severity>` ; this configuration would override any previous settings for the same destination.

- Changing the logging level for the selected syslog Message-ID; if you know your message ID, you can fully disable it by using the command `no logging message <message-ID>` or change the logging level using the command `logging message <message-ID> level <lvl> .`
- Creating a logging list and applying it to a destination; the logging list is a re-usable combination of either message class/level or a range of message IDs. It is just a more flexible way of applying the previous two methods. You define the logging list using the command `logging list <name> [level <level> [class <message_class>] | <msg_start_id>[-<msg_end_id>]]` . You may then apply the list to the particular destination using the command `logging <destination> <list-name> .`

Configuration

ASA1:

```
logging enable
!
no logging message 103012
logging list RESTRICTED message 101002-105044
logging buffered RESTRICTED
!
logging class vpdn console errors
logging class vpn console errors
logging class vpnc console errors
logging class vpnfo console errors
logging class vpnlb console errors
logging class webfo console errors
logging class webvpn console errors
```

Verification

Check the logging configuration:

```
ASA1# show logging setting
```

https://t.me/learningnets

```
Syslog logging: enabled
    Facility: 20
    Timestamp logging: disabled
    Standby logging: disabled
    Debug-trace logging: disabled
Console logging:  class vpdn vpn vpnc vpnfo vpnlb webfo webvpn, 71 messages logged
    Monitor logging: disabled Buffer logging: list RESTRICTED, 11470 messages logged
    Trap logging: disabled
    History logging: disabled
    Device ID: disabled
    Mail logging: disabled
    ASDM logging: disabled
!
!ASA1# show logging message
syslog 103012: default-level informational (disabled)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA SNMP Monitoring

Task

- Configure SNMP on ASA2 as follows:
  - Deny SNMP version 1 request using an SNMP map.
  - Send all SNMP traps to host 172.16.10.100 running version 2c.
  - Use a community of **CISCO** and location of **Bucharest,RO**.
  - SSL VPN messages of critical or higher level should be delivered as SNMP traps.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The firewall appliance supports all SNMP versions 1, 2c, and the newer version 3. Additionally, you can only poll the device via SNMP or configure the ASA to send SNMP traps, but you cannot configure the device via SNMP, so it does not support an RW community. To define a host that allows polling and receiving traps from the appliance, use the command `snmp-server host <nameif> <IP> [trap|poll] [community <KEY>] [version <1|2c|3>]` ; this statement permits the configured host to poll the

https://t.me/learningnets

device and receive SNMP traps from it. If you specify one of the keywords **trap** or **poll**, the device is limited to that function only; specify the community string explicitly if it differs from the globally configured community set using the command `snmp-server community <KEY>` . The community value is used for authentication of SNMP messages and is a pre-shared secret on both the managed device and the management station, for SNMP version 1 and 2c.

To enable sending of SNMP traps, issue the command `snmp-server enable-traps [all|<list_of_traps>]` . If you don't specify any arguments to this command, the default set of SNMP traps is enabled. Note the special type of syslog traps, which are used to convey the system log messages when you enable `logging history` in the global configuration mode; this allows logging of the syslog messages via SNMP traps. You can selectively deny certain SNMP version messages using the global command `snmp deny version <1|2c|3>` ; however, in this scenario we use the MPF inspection to block SNMP version 1 packets, by configuring a special SNMP map that denies SNMPv1 packets. Other SNMP commands, such as `snmp-server location` , are intuitive and define various informational parameters.

Although SNMP version 3 has the same scope of previous versions, it add significant security improvements and therfore configuration requirements; security is no longer tied to a community string, and all SNMP packets can be authenticated through MD5 or SHA and encryted through DES, 3DES, or AES. There are three steps to configure SNMP version 3:

1. Configure an SNMP group that basically defines SNMP security level: **noauth** (no authentication, no encrytion), **auth** (authentication, no encryption), **priv** (authentication, encryption); use the command `snmp-server group <name> v3 <noauth|auth|priv> .`
2. Configure a user who will be used in the communication between SNMP client and server, associate it with a previosuly configured group to inherit the security level, and specify strings used for authentication/encryption; use the command `snmp-server user <name> v3 auth <md5|sha> <string> priv <des|3des|aes> <string> .`
3. Configure the SNMP server and associate it with the previously configured user; use the command `snmp-server host <naemeif> <IP_address> version 3 <user> .`

Configuration

https://t.me/learningnets

ASA2:

```
snmp-server community CISCO
snmp-server host VLAN26 172.16.10.100 trap version 2c
snmp-server location Bucharest,RO
snmp-server enable traps all
!
logging class webvpn history critical
logging class webfo history critical
logging class svc history critical
!
snmp-map SNMP_VERSION_1
 deny version 1
!
policy-map global_policy
 class inspection_default
   inspect snmp SNMP_VERSION_1
```

Verification

Check the exchange of SNMP messages:

```
ASA2# show snmp-server statistics
0 SNMP packets input
    0 Bad SNMP version errors
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    0 Number of requested variables
    0 Number of altered variables
    0 Get-request PDUs
    0 Get-next PDUs
    0 Get-bulk PDUs
    0 Set-request PDUs (Not supported)11 SNMP packets output

    0 Too big errors (Maximum packet size 512)
    0 No such name errors
    0 Bad values errors
    0 General errors
    0 Response PDUs
    11 Trap PDUs
```

https://t.me/learningnets

Check that SSL VPN syslog messages are sent as SNMP traps:

```
ASA2# show logging setting

Syslog logging: enabled
    Facility: 20
    Timestamp logging: disabled
    Standby logging: disabled
    Debug-trace logging: disabled
    Console logging: disabled
    Monitor logging: disabled
    Buffer logging: disabled
    Trap logging: disabled History logging:  class webfo webvpn svc, 0 messages logged

    Device ID: disabled
    Mail logging: disabled
    ASDM logging: disabled
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA DHCP Server

Task

- Configure the ASA1 to act as a DHCP server for VLAN 49 users.
  - Use the IP address range 136.1.49.100-136.1.49.254.
  - Assign the domain-name **inelab.com** to the DHCP clients.
  - Lease the IP addresses for 30 minutes.
  - Verify by configuring R1 for DHCP address allocation on its Ethernet interface.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The ASA may function as DHCP Server or DHCP Relay even in transparent mode, but not both at the same time; with either enabled, it can also act like a DHCP client on other interfaces. As a server, the appliance could be configured to allocate IP addresses, usually on its trusted interfaces. You can set the address range for clients using the command `dhcpd address <start_address>-<end_address> <nameif> .` Other default DHCP options, such as DNS/WINS servers, domain name, and lease

https://t.me/learningnets

duration could be set using the respective DHCP sub-commands, all starting with `dhcpd` ; do not forget to enable the DHCP server on the interface using the command `dhcpd enable <nameif>` . Note the interesting command `dhcpd auto_config <nameif>` , where you mentioned the ASA interface configured for IP addressing via DHCP (therefore, the ASA is a DHCP client); then you can import the DNS/WINS/domain settings from the DHCP client to the DHCP server on the ASA.

The following should be noted about DHCP options. The firewall allows you to set any legal option value using the command `dhcpd option` . What you may need is option 3 (default gateway IP address) in transparent firewall mode; it allows you to set the default gateway to something different than the management IP address of the firewall. Another option, useful when you are dealing with Cisco IP Phones, is option 150 (TFTP Server List), which you can set to the IP address of the Cisco Call Manager if needed.

Finally, you may configure the firewall to work as DHCP relay. In this mode, as soon as the firewall receives a DHCP request on the interface configured for DHCP relaying, it sets the giaddr (gateway IP address) field in the DHCP packet to the IP of the relaying interface and forwards the request to the configured DHCP server. The commands to configure DHCP relay are `dhcprelay server <IP> <nameif>` to define the actual DHCP server and `dhcprelay enable <nameif>` to define the relaying interface; the command `dhcprelay setroute <nameif>` rewrites the default gateway option in the DHCP response from the actual server and sets it to the IP address of the respective firewall's interface.

Configuration

```
ASA1:
dhcpd address 136.1.49.100-136.1.49.254 VLAN49
dhcpd domain inelab.com interface VLAN49
dhcpd lease 1800 interface VLAN49
dhcpd enable VLAN49

R1:

interface GigabitEthernet0/1
 ip address dhcp
```

Verification

Check that R1 obtained an IP address via DHCP:

https://t.me/learningnets

```
R1#show ip aliases

Address Type             IP Address      Port
Interface                136.1.19.1
Interface                136.1.27.1 Interface                 136.1.49.100
Interface                150.1.1.1
Interface                150.1.11.11
!
!ASA1# show dhcpd  binding all

IP address       Client Identifier        Lease expiration        Type
136.1.49.100
  0063.6973.636f.2d64.           1742 seconds    Automatic
                 3438.632e.6235.6461.
                 2e64.3534.392d.4769.
                 302f.31
```

Verify DHCP settings:

https://t.me/learningnets

```
ASA1# show dhcpd state

Context  Configured as DHCP ServerInterface VLAN49, Configured for DHCP SERVER

Interface VLAN59, Not Configured for DHCP
!
!ASA1# show dhcpd statistics

DHCP UDP Unreachable Errors: 0
DHCP Other UDP Errors: 0
 Address pools          1
Automatic bindings    1

Expired bindings      0
Malformed messages    0

Message               Received
BOOTREQUEST           0DHCPDISCOVER           1
DHCPREQUEST           1

DHCPDECLINE           0
DHCPRELEASE           0
DHCPINFORM            0

Message               Sent
BOOTREPLY             0DHCPOFFER              1
DHCPACK               1

DHCPNAK               0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA IDENT Handling

Task

- Configure ASA1 to quickly terminate the IDENT lookup sessions going from VLAN 49 for TCP sessions initiated by VLAN 59 users.
- Consider both users translated using identity mappings the outside interface IP address.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

IDENT is the simple protocol that is designed to disclose the owner of the active TCP connection. This is how IDENT works:

- On a user's machine, a special IDENT daemon is running and listening on port 113.
- When a user makes an outbound connection, the remote server might initiate a reverse IDENT connection to the user's IP address and port 113, asking for the information on the TCP connection owner (to record it in the system logs).
- The daemon returns the information, and the server permits the original user's

https://t.me/learningnets

connection.
- If for some reason the IDENT daemon does not respond, it’s up to the server to decide the further connection’s fate; these days, most servers would still permit the connection.

IDENT protocol is clearly unsecure and unreliable. It might work in situations where the user has no admin privileges over the machine he or she works from. In addition to that, it requires that the user is not located behind any sort of NAT or firewall because the translating/filtering device might not properly handle the IDENT lookup. However, some applications (such as sendmail and some FTP dameons) continue to use IDENT lookups with their default configuration, even though they don’t deny the hosts not running the IDENT process. Now imagine that a user behind the ASA firewall connects to the server that does IDENT lookups. Two situations could occur:

- The user's IP is NAT translated using the IP that belongs to the ASA firewall; in this case, the reverse IDENT lookup will terminate at the ASA firewall. By default, the ASA simply discards TCP SYN packets sent to the ports it is not listening on.
- The user’s IP is not NAT translated by the firewall; in this case, the IDENT session will be attempted to the original user’s IP address. Most likely, your firewall does not permit such connections, so the TCP SYN packets will be silently discarded.

In both cases, the original server must wait for reverse IDENT TCP connection to time out before allowing the original connection, which results in delays to the end user. If the firewall sent TCP RST response instead of discarding the packets, the server would know that IDENT lookup fails immediately and would present the user with the service prompt. You may allow such reset behavior by using two commands:

```
service resetinbound
service resetoutside
```

The first comman will send TCP resets for sessions that are denied, but do not terminate on the firewall. The second command will instruct the ASA to send TCP resets for the sessions that are rejected by the firewall itself. Together, these two commands will make the firewall less stealthy but will improve user experience.

https://t.me/learningnets

Configuration

ASA1:

```
service resetinbound
service resetoutside
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Fragmented Traffic

Task

- ASA1 should drop all fragmented packets on its VLAN 59 interface.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Fragmented IP traffic presents a potentially serious security threat:

https://t.me/learningnets

- Fragment reassembly code in many OSes has been vulnerable to numerous bugs for a long time, exposing networks to attacks such as Ping of Death.
- Fragmented traffic puts a higher CPU/Memory burden on the end hosts, because it requires them to create large buffers (64 K to potentially fit any fragment size) for packet reassembly and keep them for a potentially long time, while waiting for all packet fragments.
- It is a common technique to use fragmented traffic to bypass security checks (such as IPS or firewall) by breaking malicious packets into smaller fragments (this might bypass simple firewalls, for example).

Even if it is a good practice to avoid traffic fragmentation at any cost, it is not possible in some situations. Two well-known examples are the use of NFSv2 (Sun's Network File System), which transports data blocks in large UDP packets, and VPN-protected traffic, which might exceed network MTU because of encryption/tunneling overheads.

The ASA firewall implements packet reassembly techniques to perform deep inspection of fragmented traffic. The firewall does not permit a fragmented packet through, until it fully assembles it and verifies its contents. By default, the first fragment is delayed for the duration set by the `fragment timeout <seconds>` command (5 seconds by default) until all fragments are received and assembled. If the timeout expires until all fragments arrive, the packet is discarded. The firewall will accept a maximum of `fragment chain <limit> [nameif]` fragments on the particular interface, with the default being 24 fragments. If the number of fragments exceeds the limit, the packet is discarded, so if you want to disable fragmented packets at all, use the command `fragment chain 1` to instruct the firewall to drop all fragmented packets (more than 1 fragment). Finally, all fragments are stored in a system-wide re-assembly buffer with the size set using the command `fragment size <number>` , where the maximum number of packets held in the buffer is specified.

Configuration

```
ASA1:

fragment chain 1 VLAN59
```

Verification

Generate non-fragmented traffic from R2; it should work:

https://t.me/learningnets

```
R2#ping 136.1.49.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.49.1, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!ASA1# show fragment VLAN59
Interface: VLAN59
    Size: 200, Chain: 1, Timeout: 5, Reassembly: virtual Queue: 0, Assembled: 0, Fail: 0, Overflow: 0
```

Generate fragmented traffic from R2; it should not work:

```
R2#ping 136.1.49.1 size 2000
Type escape sequence to abort.
Sending 5, 2000-byte ICMP Echos to 136.1.49.1, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!ASA1# show fragment VLAN59
Interface: VLAN59
    Size: 200, Chain: 1, Timeout: 5, Reassembly: virtual  Queue: 0, Assembled: 0, Fail: 5, Overflow: 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA ICMP Traffic

Task

- Configure ASA2 so that it does not respond to ICMP Echo.
- Ping, Path MTU discovery and traceroute from the firewall should work.
- All other ICMP messages terminated on firewall interfaces should be discarded.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Back in PIX 6.x code, all ICMP message terminating on the firewall were discarded. This default behavior was changed starting with version 7.x, and now the firewall accepts ICMP messages by default. However, the firewall does not respond to ICMP messages sent to the subnet broadcast address. If you need to filter any specific ICMP message types, you should create at least one explicit ICMP rule; this will automatically block all other ICMP message types, until they are permitted explicitly. You can also deny an ICMP message type explicitly for one subnet, while allowing it for some others. The ICMP rule statement has the following syntax: `icmp [permit|deny] <subnet> <mask> <nameif>` , with keyword any being available.

https://t.me/learningnets

If you want ASA to perform traceroute operation, configure the firewall to accept ICMP time-exceeded and unreachable messages. It is always recommended to allow the firewall to accept unreachable messages as the message type; fragmentation is required but DF bit set is used by the Path MTU (PMTU) discovery process.

Configuration

```
ASA2:

icmp permit any echo-reply VLAN19
icmp permit any echo-reply VLAN26
icmp permit any echo-reply VLAN29
!
icmp permit any time-exceeded VLAN19
icmp permit any unreachable VLAN19
!
icmp permit any time-exceeded VLAN26
icmp permit any unreachable VLAN26
!
icmp permit any time-exceeded VLAN29
icmp permit any unreachable VLAN29
```

Verification

Verify that the firewall blocks ICMP Echo, but ECHO Reply is allowed:

```
R1#ping 136.1.19.16
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.16, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!ASA2# ping 136.1.19.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.1, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Verify that traceroute initiated from the firewall works:

```
ASA2# traceroute 150.1.11.11
```

https://t.me/learningnets

```
Type escape sequence to abort.
Tracing the route to 150.1.11.11
 1  136.1.19.1 0 msec *  0 msec
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA SMTP Inspection

Task

- Configure ASA2 for SMTP traffic inspection from VLAN 19 toward the SMTP server located at 172.165.10.100:
  - Reject emails from senders containing of the strings **cyberspam.org** or **nullroute.com**.
  - Perform SMTP banner obfuscation to prevent the SMTP server identification.
  - Only accept emails addressed to the domain **cisco.com**, but reject emails that have more than 3 invalid recipients.
  - To protect against TCP SYN flooding, limit the number of half-open connections to 50 and the maximum number of connections to 100.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

In this task, we configure a regex to match senders' email addresses; this regex is later used in the STMP inspection policy map. The same inspection policy map is

https://t.me/learningnets

configured for STMP banner obfuscation in addition to allowing email relaying only for the domain **cisco.com**. We create an access-list and layer3/layer4 class-map that matches STMP traffic to the server; a regular policy-map is created, matching the layer3/layer4 class and setting various TCP connection options. Additionally, ESMTP inspection is configured along with the custom inspection policy map.

Configuration

ASA2:

```
regex UNWANTED “(cyberspam.org|nullroute.com)”
!
policy-map type inspect esmtp SMTP_INSPECT
 parameters
   mask-banner
   mail-relay cisco.com action drop-connection
 match invalid-recipients count gt 3
   reset
 match sender-address regex UNWANTED
   reset
!
access-list SMTP_SERVER permit tcp any host 172.16.10.100 eq 25
!
class-map SMTP_SERVER
 match access-list SMTP_SERVER
!
policy-map VLAN19
 class SMTP_SERVER
  set connection conn-max 100
  set connection embryonic-conn-max 50
  inspect esmtp SMTP_INSPECT
!
service-policy VLAN19 interface VLAN19
```

Verification

Note that the ASA rejects any STMP commands sent over telnet connection because it interprets them as incorrectly formatted (not per the RFC standards). Most STMP servers accept SMTP sessions simulated over telnet, but the firewall checks are very strict on this.

```
ASA2# show service-policy interface VLAN19 inspect esmtp
```

https://t.me/learningnets

```
Interface VLAN19:
  Service-policy: VLAN19
    Class-map: SMTP_SERVER
      Inspect: esmtp SMTP_INSPECT, packet 2, drop 0, reset-drop 0 mask-banner, count 0
mail-relay domain cisco.com, drop 0
match invalid-recipients count gt 3
          reset, packet 0 match sender-address regex UNWANTED
          reset, packet 0
!
!ASA2# show service-policy interface VLAN19 set connection

Interface VLAN19:
  Service-policy: VLAN19
    Class-map: SMTP_SERVER Set connection policy: conn-max 100 embryonic-conn-max 50

        current embryonic conns 0, current conns 0, drop 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA HTTP Inspection

Task

- Configure HTTP inspection on ASA2 for users on VLAN 19:
  - ASA should spoof the HTTP server headers to pretend that it is **Apache/2.2.0 (Unix)**.
  - Allow access to only **ine.com**; otherwise deny the connection and log the violation.
  - Reset the TCP connection upon any HTTP protocol violations for extra security.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The purpose of Modular Policy Framework (MPF) is to allow creating flexible traffic inspection rules. Traffic inspection is the core of the stateful firewall, and it has the following main purposes:

- Inspects protocol commands and dynamically modifies firewall rules to permit

https://t.me/learningnets

secondary information channels for protocols requiring it, like H.323 and FTP.
- Performs deep packet inspection to detect protocol violations, anomalies, and attacks.
- Dynamically modifies packet contents to remove potential threats or hide some sensitive information.

MPF replaced the legacy fixup traffic inspection command from PIX with totally new configuration concepts. MPF configuration is similar to Modular QoS CLI (MQC) configuration on IOS routers. Starting with recent IOS releases, you may find many similarities between the IOS zone-based firewall and the ASA MPF configuration. To inspect traffic in the ASA firewall, you should perform the following basic steps:

- Define layer3/layer 4 classification criteria (basically matching intended traffic). For example, if you want to inspect HTTP traffic to a particular server, you may define an access-list to match it and then create a new layer3/layer4 class-map that matches the traffic flow:

```
class-map HTTP_TRAFFIC_CMAP
 match access-list HTTP_TRAFFIC_ACL
```

This is very similar to MQC configuration in IOS routers, as class-maps are used for traffic classification based on layer3/layer4 criteria. As usual, you may create match-any or match-all class-maps, with the first type matching any of the statements in the class-map, and the second requiring all statements to be matched. In addition to using ACLs, you can match traffic based on DSCP/Precedence, port numbers as well, although using the ACLs is more flexible. In addition, the ASA firewall allows using some unique classification criteria, such as matching VPN traffic flows.

- Define a policy-map that associates certain actions with layer3/layer4 class-maps. For example, here is a policy that inspects HTTP traffic for a previously configured class-map:

```
policy-map INSPECT_HTTP_PMAP
 class HTTP_TRAFFIC_CMAP
   inspect http
```

In addition to inspection, another common setting is the command **`set connection`**, which allows changing various TCP session parameters. Most notably, you can change the maximum number of connections and the maximum number of half-open (embryonic) connections for the particular traffic class. This is much more flexible that using the static command, which is no longer available starting with

https://t.me/learningnets

ASA code 8.3(1).

- Apply the policy map to an interface or globally. When you apply the service policy to an interface using the command `service-policy <policy> interface <nameif>`, all ingress or egress traffic on the interface is inspected according to the policy settings. When you apply the policy-map globally using the command `service-policy <policy> global`, it applies to all inbound traffic flows on all interfaces.

Note that there is always a default global policy applied. If you have both the global and interface policies applied, the interface policy takes precedence for the same feature. For example, if the interface-level policy inspects HTTP and the global policy inspects FTP, the effect is combined. However, if both the interface and global policy inspect the same protocol (with potentially different settings), the interface-level policy-map takes precedence.

The default inspection policy uses a special match statement, `match default-inspection-traffic`, which allows matching multiple default protocols at the same time. The respective class has a group of inspect statements associated in the default global policy-map. This is needed to replicate the default behavior of the old fixup commands. The next question is, What if we would like to tune the inspection engine parameters, instead of simply relying on the default values? For example, we would like to tune HTTP and FTP inspection. To accomplish this, the firewall supports the concepts of inspection class-maps and policy-maps. Defined using the syntax `class-map type inspect <protocol>` and `policy-map type inspect <protocol>`, they are designed to be used together. If we want to create a customer HTTP inspection rule, we would use the following syntax:

```
policy-map type inspect http HTTP_INSPECT
  parameters
   protocol-violation action reset
```

Later, this customer inspection policy map is bound to the inspect statement under the layer3/layer4 class in the regular policy-map, allowing for fine-tuning of the HTTP inspection engine:

```
policy-map INSPECT_HTTP_PMAP
  class HTTP_TRAFFIC_CMAP
    inspect http HTTP_INSPECT
```

Every type of inspect policy map (HTTP, ESMTP, FTP, DNS, and so on) has its own set of unique parameters and options, which are fully described in the firewall documentation. Some advanced inspection settings might require protocol-specific

https://t.me/learningnets

classifications. For example, you might want to set various HTTP inspection options only for sites matching certain URLs. To accomplish this, we create a special HTTP inspect class-map that matches the URIs against certain regex, and then we use this class in the HTTP inspect policy-map and associate and action.

```
  regex CISCO "www\.cisco\.com/.*"
  !
  class-map type inspect http URI_CISCO
    match request uri regex CISCO
  !
  policy-map type inspect http HTTP_POLICY
   class URI_CISCO
     log
```

You cannot use regular layer3/layer4 class-maps with the inspect type policy-map. Remember that inspection policy-maps have no effect until they are associated with the inspect command in a normal policy-map.

Configuration

ASA2:

```
access-list HTTP_TRAFFIC_ACL permit tcp 136.1.19.0 255.255.255.0 any eq 80
!
class-map HTTP_TRAFFIC_CMAP
 match access-list HTTP_TRAFFIC_ACL
!
regex ALLOW_URL_REGEX "ine\.com"
!
class-map type inspect http ALLOW_URL_CMAP
 match not request header host regex ALLOW_URL_REGEX
!
policy-map type inspect http ALLOW_URL_POLICY
 parameters
  spoof-server "Apache/2.2.0 (Unix)"
  protocol-violation action reset
 class ALLOW_URL_CMAP
  drop-connection log
!
policy-map VLAN19
 class HTTP_TRAFFIC_CMAP
  inspect http ALLOW_URL_POLICY
!
```

https://t.me/learningnets

```
service-policy VLAN19 interface VLAN19
```

Verification

Verify settings and statistics before any tests:

```
ASA2# show service-policy interface VLAN19 inspect http

Interface VLAN19:
  Service-policy: VLAN19
    Class-map: HTTP_TRAFFIC_CMAP Inspect: http ALLOW_URL_POLICY, packet 0, drop 0, reset-drop 0

               tcp-proxy: bytes in buffer 0, bytes dropped 0
        protocol violations
          reset, packet 0
        server spoofs, packet 0
        class ALLOW_URL_CMAP
          drop-connection log, packet 0
```

Check the HTTP server header spoofing:

```
R1#telnet 172.16.10.100 80
Trying 172.16.10.100, 80 ... Open
HELO
HTTP/1.1 400 Bad Request
Content-Type: text/html; charset=UTF-8 Server:Apache/2.2.0 (Unix)

Date: Thu, 21 Nov 2013 20:50:36 GMT
Connection: close
Content-Length: 326

<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN""http://www.w3.org/TR/html4/strict.dtd">
<HTML><HEAD><TITLE>Bad Request</TITLE>
<META HTTP-EQUIV="Content-Type" Content="text/html; charset=UTF-8"></HEAD>
<BODY><h2>Bad Request - Invalid Verb</h2>
<hr><p>HTTP Error 400. The request verb is invalid.</p>
</BODY></HTML>
[Connection to 172.16.10.100 closed by foreign host]
```

Configure two host mappings on R1 and verify that if non-**ine.com** is accessed, ASA drops the session:

https://t.me/learningnets

```
R1:

ip host www.cisco.com 172.16.10.100
ip host www.ine.com 172.16.10.100


R1#copy http://www.cisco.com flash0:
Address or name of remote host [www.cisco.com]?
Source filename [test2]? TEST
Destination filename [TEST]?
Accessing http://www.cisco.com/TEST...%Error opening http://www.cisco.com/TEST (I/O error)
!
!R1#copy http://www.ine.com flash0:
Address or name of remote host [www.ine.com]?
Source filename [TEST]? TEST
Destination filename [TEST]?
Accessing http://www.ine.com/TEST...%Error opening http://www.ine.com/TEST (No such file or directory)
!
!ASA2# show service-policy interface VLAN19 inspect http

Interface VLAN19:
  Service-policy: VLAN19
    Class-map: HTTP_TRAFFIC_CMAP Inspect: http ALLOW_URL_POLICY, packet 17, drop 1, reset-drop 0

               tcp-proxy: bytes in buffer 0, bytes dropped 0
        protocol violations
          reset, packet 0
        server spoofs, packet 1
        class ALLOW_URL_CMAP
          drop-connection log, packet 1
```

If ASA2 has logging enabled, the following messages will be logged as well:

```
%ASA-5-415008: HTTP - matched Class 23: ALLOW_URL_CMAP in policy-map ALLOW_URL_POLICY,
header matched - Dropping connection from VLAN19:136.1.19.1/53112 to VLAN26: 172.16.10.100/80
%ASA-5-304001: 136.1.19.1 Accessed URL 172.16.10.100:http://www.cisco.com/TEST
%ASA-4-507003: tcp flow from VLAN19:136.1.19.1/53112 to VLAN26:172.16.10.100/80 terminated by inspection engine, rea
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA TCP Inspection

Task

- Configure ASA1 to enforce additional security checks for TCP connections established across the firewall:
  - Check retransmitted TCP packets.
  - Validate TCP checksums.
  - Clear all reserved bits in TCP headers.
- The policy should apply to all telnet connections crossing the firewall appliance:
  - Limit the concurrent number of open Telnet sessions to 3 per user.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The ASA appliance allows setting well-known TCP connection options, such as:

- Total number of open connections per layer3/layer4 class (such as host, group of hosts, subnet, etc.) using the command `set connection conn-max <number>` under the class-map assigned to a policy-map.

https://t.me/learningnets

- Total number of embryonic (aka half-open or incomplete) sessions per layer3/layer4 class to prevent TCP SYN-flooding attacks. This feature supersedes the legacy static command TCP parameters; the syntax is **`set connection embryonic-conn-max <number>`** under the class-map assigned to a policy-map.
- TCP Initial Sequence Number (ISN) randomization to prevent connection hijacking or packet injection attacks. Use the command **`set random-sequence-number [enable|disable]`** to switch this feature for the hosts matched by the particular class-map.

The firewall also allows setting the connection limits on a per-host basis; for example, if you configure the **`set connection per-client max <number>`** command under a class, the limit will apply to every host matched by this class-map and not to the aggregate number of connections. Of course, you can also limit the embryonic connections on a per-client basis using the command **`set connection per-client-embryonic-max <number>`**. ASA also allows you to tune various aspects of TCP connection normalization. The firewall applies several security checks and modifications to the TCP connections to prevent potential attacks. You can modify some TCP normalization settings using the special tcp-map applied under layer3/layer4 class-map by using the command **`set connection advanced-options`**, as follows:

```
tcp-map TCP_MAP
  window-variation drop
!
access-list extended SSH permit tcp any any eq 22
!
class-map SSH
 match access-list SSH
!
policy-map TEST
 class SSH
  set connection advanced-options TCP_MAP
```

You can read the detailed description of all configurable TCP normalization parameters in the firewall configuration guide. Naturally, to better understand TCP normalization, you should know TCP working internals (such as the format of TCP header, TCP flow control, and TCP flags). Note that when you apply TCP inspection/normalization features per interface, it affects both ingress and egress traffic. When the TCP features are applied using the global policy map, it affects ingress traffic on all interfaces.

https://t.me/learningnets

Configuration

ASA1:

```
tcp-map TCP
 check-retransmission
 checksum-verification
 reserved-bits clear
!
class-map TELNET
 match port tcp eq 23
!
policy-map global_policy
 class TELNET
  set connection per-client-max 3
  set connection advanced-options TCP
```

Verification

Verify your configuration:

https://t.me/learningnets

```
ASA1# show service-policy global set connection


Global policy:
  Service-policy: global_policy Class-map: TELNET
 Set connection policy: per-client-max 3

        current conns 0, drop 0
      Set connection advanced-options: TCP
        Retransmission drops: 0                   TCP checksum drops : 0
        Exceeded MSS drops  : 0                   SYN with data drops: 0
        Invalid ACK drops   : 0                   SYN-ACK with data drops: 0
        Out-of-order (OoO) packets : 0            OoO no buffer drops: 0
        OoO buffer timeout drops : 0              SEQ past window drops: 0
        Reserved bit cleared: 0                   Reserved bit drops : 0
        IP TTL modified     : 0                   Urgent flag cleared: 0
        Window varied resets: 0
        TCP-options:
          Selective ACK cleared: 0                Timestamp cleared  : 0
          Window scale cleared : 0
          Other options cleared: 0
          Other options drops: 0
```

Initiate four concurrent telntet sessions from R2 toward R1; only three should be allowed:

```
ASA1# show conn address 136.1.59.2

3 in use, 3 most used
TCP VLAN49 136.1.49.1:23 VLAN59 136.1.59.2:53626, idle 0:00:55, bytes 102, flags UIO
TCP VLAN49 136.1.49.1:23 VLAN59 136.1.59.2:19020, idle 0:01:30, bytes 102, flags UIO
TCP VLAN49 136.1.49.1:23 VLAN59 136.1.59.2:43994, idle 0:01:41, bytes 102, flags UIO
```

Check the ASA drops and policy-map statistics:

```
ASA1# show asp drop

Frame drop: Connection limit reached (conn-limit)                                  3

Last clearing: 07:01:48 UTC Nov 22 2013 by enable_15

Flow drop:
```

https://t.me/learningnets

```
Last clearing: 07:01:48 UTC Nov 22 2013 by enable_15
!
!ASA1#show service-policy global set connection

Global policy:
  Service-policy: global_policy
    Class-map: TELNET Set connection policy: per-client-max 3
current conns 3, drop 4

      Set connection advanced-options: TCP
        Retransmission drops: 0                   TCP checksum drops : 0
        Exceeded MSS drops  : 0                   SYN with data drops: 0
        Invalid ACK drops   : 0                   SYN-ACK with data drops: 0
        Out-of-order (OoO) packets : 0            OoO no buffer drops: 0
        OoO buffer timeout drops : 0              SEQ past window drops: 0
        Reserved bit cleared: 0                   Reserved bit drops : 0
        IP TTL modified     : 0                   Urgent flag cleared: 0
        Window varied resets: 0
        TCP-options:
          Selective ACK cleared: 0                Timestamp cleared  : 0
          Window scale cleared : 0
          Other options cleared: 0
          Other options drops: 0
```

If ASA1 has logging enabled, the following log message should be seen:

```
%ASA-3-201013: Per-client connection limit exceeded 3/3 for input packet from 136.1.59.2/15218 to 136.1.49.1/23
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Management Traffic Inspection

Task

- Configure ASA2 for integration with a RADIUS server found on server 172.16.10.100:
- The firewall should inspect RADIUS accounting packets going to the IETF default RADIUS ports.
  - Validate the RADIUS attribute number 26 and send accounting responses.
  - Apply the inspection rule globally.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

The ASA firewall supports inspection for traffic originated to/from the firewall itself. Currently, the only protocol supported is RADIUS, with the purpose of validating firewall-to-server communications. When you configure management traffic inspection, you should create a special type of class-map (type management) matching the respective management plane traffic. Note that this map is equivalent to the layer3/layer4 class-map, not the special inspection map used with other protocols. Therefore, you assign this map to the normal policy maps applied globally

https://t.me/learningnets

or per-interface.

ASA supports a special type of policy-map of type radius-accounting that is used to set the RADIUS protocol inspection parameters. This is the policy-map that you use as the parameter to inspect the radius-accounting command. Under the radius-accounting inspection policy map, you can set various supported inspection options; however, the most important configuration set here is the IP address and the shared secret of the RADIUS server. When you're done configuring the RADIUS inspection policy map, you may apply it under a class in the global or per-interface policy-map. Note that this class will only match traffic terminated on the firewall itself, and will not affect inspection of the through traffic.

Configuration

ASA2:

```
aaa-server RADIUS  protocol radius
aaa-server RADIUS (VLAN26) host 172.16.10.100 CISCO
!
class-map type management RADIUS
 match port udp eq radius-acct
!
policy-map type inspect radius-accounting RADIUS_INSPECT
 parameters
  send response
  validate-attribute 26
  host 172.16.10.100 key CISCO
!
policy-map global_policy
 class RADIUS
  inspect radius-accounting RADIUS_INSPECT
```

Verification

Verify your configuration:

```
ASA2# show service-policy global inspect radius-accounting

Global policy:
  Service-policy: global_policy
    Class-map: inspection_default
    Class-map: RADIUS
      Inspect: radius-accounting RADIUS_INSPECT, packet 0
```

https://t.me/learningnets

```
        timeout users 1:00:00
        no enable gprs send response
validate-attribute 26, connections dropped 0
host 172.16.10.100 key CISCO
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Threat Detection

Task

- Enable basic threat detection on ASA1:
  - Set additional monitoring intervals for the ACL drop event so that a message is generated every time there are more than 10,000 drops per two hours or more than 1,000 drops per 10 seconds.
  - Enable advanced scanning attack detection and automatic shunning of the attackers.
  - Configure the firewall to shun the attackers for 10 minutes, but never clear any connections originated from the 150.1.1.1 host.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Threat detection is a new ASA v8.x code feature. It allows detection of network attacks by collecting the drop events occurring in the firewall and calculating their rates (speed). The firewall monitors certain events, such as ACL packet drops, invalid packets, embryonic connections exceeding the limit, interface buffer

https://t.me/learningnets

overloading, etc., and fires a system log message when the event rate exceeds configured thresholds. You may find the complete list of monitored system events in the firewall configuration guide or by using the firewall shell's online prompt. Note that all those events are related to firewall drops of IP packets from various causes, and therefore they are only a side-effect of firewall functions, so threat detection does not impose any significant load on the firewall's CPU. The threat Detection feature has three main components:

- Basic Threat Detection
- Advanced Threat Detection
- Scanning Threat Detection

Basic Threat Detection monitors the rate at which packets are dropped and is enabled by default; every event type has default rate thresholds configured. Statistics generated by basic threat detection apply to the entire device and are generally not granular enough to provide information about the source or specific nature of a threat. You may disable basic threat detection by using the command `no threat-detection basic .` Also, you may configure two additional sets of thresholds for every event type by using the command `threat-detection rate [event-type] rate-interval <rate_interval> average-rate <avg_rate> burst-rate <burst_rate> .` The firewall counts events occurring within time intervals using two sliding time windows. The first one is defined by the rate_interval, and the other is defined as max(1/30 *<rate_interval>, 10) seconds. The rate calculated over the first interval is the average event rate, whereas the second rate is the burst or instant event rate. Using both intervals allows the appliance to monitor both long-term average and short-term event rates. Note that the rate_interval value should be a multiple of 60 seconds if it is less than 3600 seconds (1 hour) or in multiples of 3600s if it's above 1 hour. The burst-rate interval is calculated automatically with the formula above.

Advanced Threat Detection can be used if granularity is desired to better track and identify the attackers. ASA supports tracking statistics for host IPs, ports, protocols, ACLs, and servers protected by TCP intercept. Advanced Threat Detection is only enabled by default for ACL statistics, but all options can be enabled with the command `threat-detection statistics` For host, port, and protocol objects, Threat Detection tracks the number of packets, bytes, and drops that were both sent and received by that object within a specific time period. For ACLs, Threat Detection tracks the top 10 ACEs (both permit and deny) that were hit the most within a specific time period. Advanced Threat Detection is more resource consuming than Basic Threat Detection because various statistics are kept in memory.

Scanning Threat Detection is used to track suspected attackers with connections to many hosts in a subnet, or many ports on a host/subnet. Scanning Threat Detection is disabled by default. It is based on the Basic Threat Detection functionality, so the

https://t.me/learningnets

rate-interval, average rate (ARI), and burst rate (BRI) settings are shared between Basic and Scanning Threat Detection. The difference between the features is that whereas Basic Threat Detection only indicates if the average or burst rate thresholds were crossed, Scanning Threat Detection maintains a database of attacker and target IP addresses. The basic threat detection recognizes network scanning attack-based drops of invalid packets (for example, FIN packets without initial SYN for nmap SYN scanning). You may configure the drop rate thresholds by using the command `threat-detection rate scanning-threat` and setting the thresholds. The basic mode does not use any advanced heuristics to recognize scanning hosts or their targets.

You may enable advanced scanning attack detection (Scanning Threat Detection) using the command `threat-detection scanning-threat [shun]` . In this mode, the firewall will collect extensive statistics on hosts and their behavior. Specifically, the appliance will be looking for hosts attempting connection (or pinging) to many destinations in sequence, probing many ports on the same machine, or trying to connect to non-existent services. Thus, the advanced network scanning detection is behavior based and does not rely on counting the invalid packet drops. At the same time, this mode might impose serious load on the firewall's resources. The shun option allows automatic shunning of the suspected attackers. The default shunning interval is one hour and can be changed by using the command `threat-detection scanning-threat shun duration <seconds>` . Additionally, you may exempt some hosts from being shunned, even if they are recognized as network scanners by the firewall. The command to excluded the hosts is `threat-detection scanning-threat shut except [<ip> <mask>|object-group <name>]` .

Threat Detection is used only to gather basic or advanced statistics and cannot restrict traffic that otherwise would have been allowed by the firewall, unless shunning is configured.

Configuration

ASA1:

```
threat-detection rate acl-drop rate-interval 7200 average-rate 10000 burst-rate 24000
!
threat-detection scanning-threat shun duration 600
threat-detection scanning-threat shun except ip-address 150.1.1.1 255.255.255.255
```

Verification

https://t.me/learningnets

Check threat detection statistics:

```
ASA1# show threat-detection rate
                         Average(eps)    Current(eps) Trigger      Total events
  10-min SYN attck:                 0               0       0                 6
  1-hour SYN attck:                 0               0       0                12
  10-min  Scanning:                 0               0       0                 6
  1-hour  Scanning:                 0               0       0                12
!
!ASA1# show threat-detection statistics

Top          Name   Id    Average(eps)    Current(eps) Trigger      Total events
  1-hour ACL hits:
01  PERMIT/1                      0               0       0                56
  8-hour ACL hits:
01  PERMIT/1                      0               0       0                56
 24-hour ACL hits:
01  PERMIT/1                      0               0       0                56
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA uRPF

Task

- You have been notified of IP spoofing attacks sourced from **TEST-PC-A**.
- Configure ASA2 to stop these attacks.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the `clear configure` command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

uRPF (Unicast RPF) has the same function and scope on both IOS (routers/switches) and ASA; it's the same concept as RPF checks in multicast traffic, but the scope is different (in multicast, the scope is to have a loop-free multicast path, whereas in unicast the scope is to stop IP spoofing attacks). When ASA has uRPF enabled on an interface, it performs a route lookup for the source of all packets received inbound on that interface and it will drop packets if the route lookup process does not find the best path through the interface packet it was received on.

https://t.me/learningnets

uRPF is disabled by default on ASA, because it is on IOS as well. Whereas IOS has two implementations for uRPF (loose and strict), ASA supports only strict mode.

Configuration

```
ASA2:

ip verify reverse-path interface VLAN19
```

Verification

Verify uRPF statistics before issuing spoofed packets:

```
ASA2# show ip verify statistics interface VLAN19
 interface VLAN19: 0 unicast rpf drops
```

Configure on R1 a Loopback interface with the IP address of R2's Loopback1; generate traffic from this new loopback toward R2's VLAN 29 IP address, and verify uRPF statistics on ASA2 (traffic will be dropped):

```
R1#ping 136.1.29.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.29.2, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R1#ping 136.1.29.2 source loopback2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.29.2, timeout is 2 seconds:
Packet sent with a source address of 150.1.22.22 .....
Success rate is 0 percent (0/5)
!
!ASA2# show ip verify statistics interface VLAN19
interface VLAN19: 5 unicast rpf drops
!
!ASA2# show asp drop

Frame drop:
  No route to host (no-route)                                    44
Reverse-path verify failed (rpf-violated)                         5
```

https://t.me/learningnets

```
Last clearing: 15:40:13 UTC Dec 4 2013 by enable_15

Flow drop:

Last clearing: 15:40:13 UTC Dec 4 2013 by enable_15
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 5: Perimeter Security and Services - ASA Firewall

ASA Asymmetric Routing

Task

- Configure ASA3 to run in multi-context mode as follows:

Context	Interface1	Interface1 Nameif	Interface1 IP	Interface2	Interface2 Nameif	I
SW1-R3	Gi0/0.34	VLAN34	136.1.34.17/24	Gi0/1	VLAN93	13
SW2-R3	Gi0/0.34	VLAN34	136.1.34.18/24	Gi0/3	VLAN94	13

- Configure ASA3 to support asymmetric routing for the shared interface, in case outbound/inbound traffic transits different contexts.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 5 Introduction** or by clicking the **Resources** button.

https://t.me/learningnets

> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the ASA. For the ASA, the best option is to use the **`clear configure`** command or just remove the commands issued in previous labs to revert back to the initial configurations for the section.

Overview

Asymmetric routing means that a packet does not follow the same logical path both ways (outbound client-to-server traffic uses one path and inbound server-to-client uses another path). Because firewalls track connection states and inspect traffic, asymmetric routing is not firewall friendly by default and traffic is dropped, TCP traffic being especially affected. Although in all real deployments you should avoid asymmetric routing (with or without a firewall in the path), there are certain cases when this is a requirement, or cases when you don't really have control over it. If a firewall is in the path, there are several options to still allow traffic through:

- If outbound traffic transits the firewall but inbound traffic does not, use TCP state bypass for the respective connection, or use static NAT with **`nailed`** option (which effectively disables TCP state tracking and sequence checking for the connection).
- If both outbound and inbound traffic transits the firewall but on different interfaces, use the same solutions as above.
- If outbound traffic transits one context of the ASA and inbound traffic transits another context of the ASA, use ASR groups; this applies only for multi-context mode and requires active-active stateful failover to be configured.

TCP-state bypass and NAT with nailed option basically allow packets to pass through the ASA because some security checks are no longer performed, but ASR groups are different. ASA acts as a full-blown firewall with ASR groups enabled (so there are no security gaps), just that if packets are incorrectly received because of routing, ASA will rewrite layer2 headers and inject the packet in the correct path. In our case, for example:

- Outbound traffic transits context **SW1-R3**; it exits the shared interface on VLAN 34.
- Inbound traffic is received on context **SW2-R3**, which should normally be dropped.
- Because the shared interface of VLAN 34 from both contexts is configured as a member of the same ASR group, ASA will rewrite the layer2 headers for inbound traffic received on context **SW2-R3** and inject it back into context **SW1-R3**.

https://t.me/learningnets

- The layer2 header rewrite only happens if the ASA finds a connection state in the **SW1-R3** context for the packet received on context **SW2-R3**.

> ASR groups are not enabled until stateful failover is configured.

Configuration

SW1:

```
interface FastEthernet1/0/16
 switchport mode access
 switchport access vlan 94
 no shutdown
!
interface FastEthernet1/0/17
 switchport mode access
 switchport access vlan 93
 no shutdown
```

SW2:

```
interface FastEthernet1/0/16
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 34
 switchport mode trunk
 no shutdown
```

ASA3 System Context:

```
mode multiple
!
mac-address auto
!
interface GigabitEthernet0/0
 no shutdown
!
interface GigabitEthernet0/0.34
 vlan 34
!
interface GigabitEthernet0/1
 no shutdown
!
interface GigabitEthernet0/2
 no shutdown
!
interface GigabitEthernet0/3
```

https://t.me/learningnets

```
 no shutdown
!
failover group 1
 primary
!
failover group 2
 secondary
!
failover lan unit primary
failover lan interface FAILOVER GigabitEthernet0/2
failover link failover GigabitEthernet0/2
failover interface ip FAILOVER 136.1.100.17 255.255.255.0 standby 136.1.100.18
failover
!
context SW1-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/1
 config-url disk0:/SW1-R3.cfg
 join-failover-group 1
!
context SW2-R3
 allocate-interface GigabitEthernet0/0.34
 allocate-interface GigabitEthernet0/3
 config-url disk0:/SW2-R3.cfg
 join-failover-group 2
```

ASA3 SW1-R3 Context:

```
changeto context SW1-R3
!
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 asr-group 1
 ip address 136.1.34.17 255.255.255.0 standby 136.1.34.178
!
interface GigabitEthernet0/1
 nameif VLAN93
 security-level 93
 ip address 136.1.93.17 255.255.255.0 standby 136.1.93.178
```

ASA3 SW2-R3 Context:

```
changeto context SW2-R3
!
```

https://t.me/learningnets

```
interface GigabitEthernet0/0.34
 nameif VLAN34
 security-level 34
 asr-group 1
 ip address 136.1.34.18 255.255.255.0 standby 136.1.34.187
!
interface GigabitEthernet0/3
 nameif VLAN94
 security-level 94
 ip address 136.1.94.18 255.255.255.0 standby 136.1.94.187
```

Verification

Initiate traffic from SW1 to R3, and verify that both outbound and inbound traffic transits **SW1-R3** context.

```
SW1#show ip route 136.1.34.0
Routing entry for 136.1.34.0/24
  Known via "static", distance 200, metric 0
  Routing Descriptor Blocks: * 136.1.93.17
      Route metric is 0, traffic share count is 1
!
!R3#show ip route 136.1.93.0
Routing entry for 136.1.93.0/24
  Known via "static", distance 200, metric 0
  Routing Descriptor Blocks: * 136.1.34.17
      Route metric is 0, traffic share count is 1
!
!SW1#telnet 136.1.34.3
Trying 136.1.34.3 ... Open



User Access Verification

Password: R3>
!
!ASA-X/SW1-R3# show conn
1 in use, 2 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:39214, idle 0:00:17, bytes 111, flags UIO
```

Configure R3 with a static route for VLAN 93 traffic, pointing to **SW2-R3** context:

https://t.me/learningnets

```
R3:
ip route 136.1.93.0 255.255.255.0 136.1.34.18
```

Initiate traffic from SW1 to R3, and verify that the session is still successful even though there is asymmetric routing:

```
SW1#telnet 136.1.34.3
Trying 136.1.34.3 ... Open



User Access Verification

Password: R3>
!
!ASA-X/SW1-R3# show conn
1 in use, 2 most used
TCP VLAN34 136.1.34.3:23 VLAN93 136.1.93.9:24657, idle 0:00:12, bytes 111, flags UIO
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

PKI Infrastructure Overview

Overview

PKI (Public Key Infrastructure) is a complex system that joins together a suite of protocols, standards, services, and policies with the goal of managing public key cryptography, also known as asymmetric cryptography. Although symmetric cryptography uses the same single key for both encryption and decryption processes, PKI uses a pair of two keys mathematically matched, one for the encryption process and one for the decryption process. One key is publicly available (the public key), and the other one is known only to its owner and should be never disclosed (the private key). One important aspect of the PKI infrastructure is that each entity should safeguard its own private key. Based on the usage of the two keys, we can achieve the following:

- **Confidentiality**. A message encrypted with a public key can only be decrypted by its corresponding private key, which is known only to its owner, thus ensuring that the message can be read/decrypted only by the private-key owner for which the message was generated.
- **Authentication/integrity**. A message encrypted with a private key can only be decrypted by its corresponding public key; being public means that anyone can get the public-key and thus read/decrypt the message, but because the private-key resides only with the owner, it guarantees identity of the message source. This process ensures authenticity and non-repudiation of the message (also known as **digital signature**); if combined with a hash function, it can also provide message integrity.

There is a design flaw in this process: If I want to send you a secret message, I should encrypt it using your public key, but how can I make sure that I received your public key and not someone else public's key? Or how can I make sure that your public key was not intercepted in transit to me and changed to another public key of an attacker (when I send the message, you will not be able to read it, but the attacker will, so the message is no longer a secret between us). We need to provide

https://t.me/learningnets

more information along with the public key to prove and guarantee the recipient's identity. The role of a Certificate Authority (CA) exists to create the necessary trust level between two peers through the use of identities, also called **digital certificates**.

The following is a simplified step-by-step guide for issuing **digital certificates** through the CA:

1. The CA generates its own pair of RSA/EC public/private keys, issues a self-signed digital certificate, and becomes a CA (the private key of the CA should be safeguarded and never be shared with others).
2. Each client requiring a digital certificate generates its own pair of public/private RSA/EC keys (the private key of the client should be safeguarded and never be shared with others).
3. The public key of the client, along with other information required to prove its identity, is submitted to the CA in the form of a Certificate Signing Request (the request is signed with the client private key so that the CA can verify the authenticity of the request).
4. The CA provides a digital signature to the submitted information using its own private key, and the result is a digital certificate.
5. The resulting digital certificate is issued back to the requesting client.

A digital certificate (also known as a public key certificate) is like a passport or an identity card containing information (similar to a personal identification number) that uniquely identifies you; to be recognized as valid, it also requires the signature of the issuer (the CA). To be universally valid, it needs a certain format or structure and a minimum amount of information, which is described in X.509 standard. At a minimum, it must make use of the following:

- **Version**. Based on the version and the certificate template being used, the certificate is valid for certain purposes (for example, it can be used for encryption but not for digital signature).
- **Serial Number**. Uniquely identifies a certificate issued by a CA.
- **Issuer**.Information about the CA, such as Common Name or FQDN.
- **Validity**. The start/end time for which the certificate is valid and can be used.
- **Subject**. Information about the client, such as Common Name or FQDN.
- **Public Key**. The client's public key data and algorithm.
- **Certificate Signature**. The CA signature data and algorithm.

https://t.me/learningnets

When two peers need to perform mutual authentication using digital certificates (such as in SSL or IPsec VPNs), the following is required:

- Both peers exchange their certificates, which need to be valid at that time.
- Certificates can be issued by different certificate authorities.
- For one peer to verify the identity from the presented certificate of the remote peer, it must read the digital signature from the certificate, so it needs the know and trust the public key of the issuing certificate authority.

The Certificate Signing Request (CSR) required to obtain an X.509 digital certificate also uses a standard format, which is described in the PKCS#10 standard. At a minimum, the CSR contains the public key and one or more of the following fields, which will appear in the Subject field of the certificate:

- Common Name or Distinguished Name (FQDN)
- Organizational Unit and Organization
- Town and Country
- E-mail Address

In general, a CSR can be generated and submitted using Active Directory Certificate Services (ADCS), Simple Certificate Enrollment Protocol/ Network Device Enrollment Service (SCEP/NDES), Web enrollment, and Manual (terminal-based) enrollment. When the CA issues/approves the certificate, it must be installed/imported by the requesting client, and in the same time it may be desirable by the client to export it for backup purposes; for these two processes to happen seamlessly regardless of the CA vendor, some standard certificate file formats must exist:

https://t.me/learningnets

- **PEM (Privacy Enhanced Mail)**. A base64 encoded file containing one or more digital certificates along with associated private keys, optionally protected by a password.
- **PFX/PKCS#12 (Personal Information Exchange)**. A binary format file containing one or more digital certificates along with associated private keys and all CA certificates in the certification chain, optionally protected by a password.
- **P7B/PKCS#7 (Cryptographic Message Syntax Standard)**. A base64 encoded file containing one or more certificates and all CA certificates in the certification chain, without support for the private keys of the certificates.
- **DER (Distinguished Encoding Rules)**. A binary format file containing a single certificate, without support for certificate chain or private keys.
- **Base64**. A binary format file containing a single certificate, without support for certificate chain or private keys.

The digital certificate provides only the public information, but for it to be usable by its owner, it needs the associated private key. For this reason, if you want to export a certificate for backup purposes and make it usable on import, you must use the PEM or PKCS#12 formats.

A digital certificate can be revoked at any time, such as when compromised or no longer used by its owner. For this reason, when a trust relationship based on certificates is established, it may be desirable to first determine whether the received certificate of a remote peer is still valid (not in terms of its validity period, because this information is present in the certificate information). The only entity that can revoke a certificate (invalidate it before it expires according to its validity period) is the issuer, and this is the CA. There are two options for verifying a certificate:

- The CA can publish a Certificate Revocation List (CRL), which contains the serial numbers of certificates in revoked or hold state. The CA refreshes this list periodically but not instantly, and any client can download and cache it for a period of time; this method of validating certificates is not 100% accurate because it does not happen in real time.
- Using Online Certificate Status Protocol (OCSP). Each time a digital certificate is received, an OCSP query is sent to ask about its status based on the serial number; this method is 100% accurate because it is not based on a periodically refreshed list, but on a real-time query/response.

https://t.me/learningnets

In large PKI deployments, you may experience performance issues caused by a high number of CSRs, or you may want to have multiple CAs geographically dispersed within the same PKI system. For these reasons, a hierarchical PKI may be desired. The components are:

- **Root CA**. This is the first CA installed in your PKI deployment.
  - Highest level in a hierarchical PKI.
  - Can act as a fully functional CA (manage client certificate's life cycle).
  - Has a self-signed digital certificate with the highest lifetime in the PKI.
  - Can delegate all the CA functions to subordinate CAs.
  - Can delegate part of its functions to RAs.
  - With Subordinate CA's, the Root CA is taken offline for security reasons and is only responsible for issuing certificates to subordinate CAs.
- **Subordinate CA**. This is the second CA installed in your PKI.
  - Second-highest level in a hierarchical PKI if subordinate to the root CA.
  - Can also be subordinate to another subordinate CA.
  - Can act as a fully functional CA (manage client certificate's life cycle).
  - Has a digital certificate received from the root/subordinate CA with a high lifetime.
  - Can delegate all the CA functions to other subordinate CAs.
  - Can delegate part of its functions to RAs.
- **Registration Authority**. This is an optional role in your PKI.
  - Acts as a proxy between clients and root/subordinate CA.
  - Its functions are performed by default by the root/subordinate CA.
  - Authenticates and authorizes client certificate requests.
  - Forwards valid and approved requests to the subordinate/root CA.
  - Cannot issue client certificates.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

PKI Server on ASA Firewall

Task

- Configure ASA4 to act as a certificate authority with the following parameters:
  - All RSA keys should use a size of 2048 bits.
  - Use the SMTP email address **ASA4@ine.com**.
  - Use the email subject **Use this URL to enroll with INE CA**.
  - All certificates should have a lifetime equal to double the default value of client-issued certificates.
  - CRL should be available to all users on VLAN 18 on port 443.
  - All VLAN 18 users should be allowed to enroll.
- Provision and allow to enroll the following users:
  - **rack1r1** with the email address **rack1r1@ine.com** and the CN **rack1r1**.
  - **rack1r2** with the email address **rack1r2@ine.com** and the CN **rack1r2**.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **PKI Infrastructure** module of this section, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.

Overview

The ASA does not currently support a CA server function with Elliptic Curve Digital Signature Algorithm (ECDS) certificates, which is required for Suite-B cryptography (Next Generation Encryption). The Suite-B protocols bring together enhanced cryptographic security and increased encryption performance:

- RSA keys are replaced by ECDSA keys, so the PKI infrastructure must support ECDSA signatures.
- DH groups are replaced by ECDH-19, ECDH-20, and ECDH-21 groups.
- AES-CBC mode is replaced by AES-GCM or AES-GMAC; GCM supports integrated

https://t.me/learningnets

encryption and integrity, whereas GMAC supports only integrity.
- SHA-1 is replaced by SHA-256, SHA-384, or SHA-512.

Because Suite-B is supported beginning with ASA 9.0 and only on the new 5500-X platforms, it will not appear in the current version of the CCIE Security Lab Exam. In accordance with RFC 6380, Suite-B protocols are supported only with IKEv2 and PKI (ECDSA signed certificates) authentication, so there is no IKEv1 or RSA/pre-shared-key authentication. The Cisco ASA implements a very basic functionality of a Certificate Authority role:

- It does not support hierarchical CA deployments; it can only act as a root CA.
- An ASA with the CA role can no longer be member of a failover pair (Active/Standby or Active/Active).
- It does not support CA High Availability configuration (one ASA acting as the active CA and one ASA acting as the standby CA).
- It only supports RSA signed certificates.
- Certificate database and configuration files can be kept locally on flash (by default) or on external CIFS/FTP storage.
- The only supported certificate enrollment method is via a dedicated internal web page (note that this is not SCEP enrollment).
- Only HTTP is supported for publishing the CRL list, and it can be locally or externally stored.
- Issued certificates are scoped to be used for SSL VPN session authentication (both clientless/browser-based and client-based/AnyConnect).

The following are the steps necessary to perform client certificate enrollment:

1. Configure minimum required and optional CA parameters.
2. Enable the CA.
3. Enable webvpn (clientless SSL VPN) on the interface where users will enroll.
4. Configure the database of users allowed to enroll with the CA; each user is automatically assigned an OTP (one-time password) that can be delivered to the user through the use of email or directly spoken.
5. Users access the web enrollment page URL, authenticate using their provided username and OTP, and download their certificate in PKCS#12 format.

To make the process as easy as possible from the user perspective, PKCS#12 certificate format is used, because:

- It does not require the user to perform any CSR before getting the certificate; the user receives both the certificate and associated private key after being successfully

https://t.me/learningnets

authenticated.
- The CA certificate can be embedded in it; the CA certificate is required by the client to validate the digital signature of the CA.

Because all required information is presented in a single file, the ASA implements this highly secure deployment method with OTP authentication. After the CA has been enabled, configuration parameters may be no longer changed: Some require disabling the CA temporarily to make changes, while others, like the CA RSA key size, cannot be changed unless the CA role is completely removed and reconfigured. For this reason, it is recommended that you fully configure the CA parameters before actually enabling it. At a minimum, you need to know the following information:

- Enter the CA server configuration mode: `crypto ca server` ; at this moment the CA service is inactive.
- Enable the CA server with its default configuration parameters: `no shutdown` , and type a password with a minimum length of 8 characters; at this moment the ASA generates its RSA public/private key pair and issues a self-signed certificate, which are bound together in an archive protected by the configured password. It creates in flash a folder named **LOCAL-CA-SERVER** for its required database files.
- LOCAL-CA-SERVER.ser is a clear-text file containing the serial number of issued certificates.
- LOCAL-CA-SERVER.cdb is clear-text file containing information about issued certificates.
- LOCAL-CA-SERVER.udb is a clear-text file containing information about enrolled users.
- LOCAL-CA-SERVER.crl is a password-encrypted file containing information about revoked certificates.
- LOCAL-CA-SERVER.p12 is a password-encrypted file containing the RSA public/private key pair and the self-signed certificate.

An RSA key pair generated on the ASA is always exportable and cannot be configured as non-exportable; the private key is stored on a hidden partition in flash. The ASA does not allow you to export the RSA key pair (no command available). If the RSA key pair is tied to a certificate issued to the ASA (self-signed by the local CA or obtained from an external CA), it can be exported along with the certificate in a PEM or PKCS#12 format. The default and nonconfigurable enrollment URL is **https://[hostname|IP Address]/+CSCOCA+/enroll.html**. The default but configurable CRL URL is **http://[hostname|IP Address]/+CSCOCA+/asa_ca.crl**. Other configurable parameters of the CA are:

https://t.me/learningnets

- Subject DN data, which will be included on all issued certificates, has no value by default and can be configured using the command `subject-name-default <dn> .`
- Issuer DN data, which will be included on all issued certificates, has a default value of `hostname.domainname` and can be changed using the command `issuer-name <dn> .`
- Certificate database location is by default on flash, and it can be moved to an external username/password protected CIFS/FTP location using the command `database path <path> .`
- CDP URL is by default hosted on the ASA and can be modified using the command `cdp-url <url> .`
- Email sender address and subject name, used when delivering the OTP via email, have default values of `admin@hostname` and `Certificate Enrollment Invitation` ; these can be changed using the commands `smtp from-address <e-mail>` and `smtp subject <subject> .`
- RSA key pair size for all keys (CA and clients) has a default value of 1024 bits and can be changed using the commands `keysize <size>` and `keysize server <size> .`
- Lifetime of CA certificate, client certificates, and CRL list have default values of: 3 years (CA), 1 year (client certificate), and 6 hours (CRL). These can be changed using the commands `lifetime ca-certificate <days> ,` `lifetime certificate <days>` , and `lifetime crl <hours> .`
- When a user has been provisioned by the ASA for certificate enrollment, its key pair and certificate file in PKCS#12 format are by default valid to be retrieved for 24 hours; this interval can be changed with the command `enrollment-retrieval <hours> .`
- When the CRL is hosted on the ASA, by default it is not accessible; this behavior can be changed using the command `publish-crl <interface_nameif> <port> .`
- The default client-issued certificate lifetime is 1 year, so you may want to notify users when their certificate is about to expire; by default, the ASA emails users 14 days before their certificates expire, and this can be changed using the command `renewal-reminder <days> .`
- By default, the OTP password, which can be used by the user for certificate enrollment, is valid for 3 days; this setting can be changed using the command `otp expiration <hours> .`

Additionally, you can revoke and un-revoke an issued certificate at any time with the command `crypto ca server {revoke|unrevoke} <SN>` . Users allowed to enroll with the ASA CA can be added using the command `crypto ca server user-db add <username> ,` with optional parameters `email` and `dn` . Users can also be removed from the database by using the command `crypto ca server user-db remove <username>` . If for various reasons the OTP did not reach the user on its email address and you want

https://t.me/learningnets

to view it, use the command **`crypto ca server user-db show-otp <username>`**. There are multiple verification commands for both certificate and user databases, all starting with **`show crypto ca server .`**

Configuration

When enabling the CA with the **`no shutdown`** command, you may receive the following error message if the CA has been previously configured on the ASA. In this case you need to first delete the **LOCAL-CA-SERVER** directory using command the **`delete /recursive /noconfirm flash:/LOCAL-CA-SERVER .`**

```
% ERROR: A certificate database file already exists.
Please remove it before creating the new CA.
```

Configure the CA server now:

ASA4:

```
crypto ca server
 keysize 2048
 keysize server 2048
 smtp from-address ASA4@ine.com
 smtp subject Use this URL to enroll with INE CA
 lifetime ca-certificate 730
 lifetime certificate 730
 publish-crl VLAN18 443
 no shutdown
!
!
webvpn
enable VLAN18
!
crypto ca server user-db add rack1r1 email rack1r1@ine.com dn CN=rack1r1
crypto ca server user-db add rack1r2 email rack1r2@ine.com dn CN=rack1r2
crypto ca server user-db allow user rack1r1
crypto ca server user-db allow user rack1r2
```

Verification

Verify that the CA server is enabled and that port 443 is opened for enrollment:

ASA4#show crypto ca server

https://t.me/learningnets

```
Certificate Server LOCAL-CA-SERVER: Status: enabled
 State: enabled
    Server's configuration is locked  (enter "shutdown" to unlock it)
    Issuer name: CN=ASA4.ine.com
    CA certificate fingerprint/thumbprint: (MD5)
        90746d63 dc114c20 4ada7c60 550a88ef
    CA certificate fingerprint/thumbprint: (SHA1)
        412c2aca 712ec29b 28409b5c 0156eec9 44d34b6e
    Last certificate issued serial number: 0x1
    CA certificate expiration timer: 08:23:20 UTC May 8 2015
    CRL NextUpdate timer: 20:22:50 UTC May 8 2013
    Current primary storage dir: flash:/LOCAL-CA-SERVER/

    Auto-Rollover configured, overlap period 30 days
    Autorollover timer: 08:23:20 UTC Apr 8 2015
!
!ASA4#show asp table socket


Protocol  Socket    Local Address               Foreign Address         State
SSL       0011de0f  136.1.18.8:443              0.0.0.0:*               LISTEN
DTLS      002606bf  136.1.18.8:443              0.0.0.0:*               LISTEN
```

Verify that users are allowed to enroll and display the OTP required for enrollment:

```
ASA4(config)#show crypto ca server user-db username rack1r1
 username: rack1r1
 email:    rack1r1@ine.com
 dn:       CN=rack1r1
 allowed:  17:14:30 UTC Sat May 11 2013
 notified: 1 times enrollment status: Allowed to Enroll
 !
 !ASA4(config)#crypto ca server user-db show-otp user rack1r1
 Username: rack1r1
 OTP: CD9BDE9C434E6060
 Enrollment Allowed Until: 17:14:30 UTC Sat May 11 2013
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

RSA Key Management on IOS Router

Task

- Configure R3 for SSH server functionality:
  - Use the key-size of 1024 bits.
  - Name the key-pair **SSH**.
  - Ensure that you can use the same key-pair on multiple routers if needed.
  - Ensure that on a router reload, SSH is disabled.
  - Use the most secure version of the protocol.
- Enable SSH public-key authentication for user **CCIE**, because its public-key is not yet known and will be later added.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **PKI Infrastructure**, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Although IOS supports Elliptic Curve (EC) key pairs starting with IOS 15.1(2)T (because these are currently scoped for use with IKEv2 and Suite-B cryptography protocols), we're not concerned with it for the scope of the current CCIE Security lab exam.

On IOS routers, RSA key-pairs can be used for multiple purposes, such as SSH server, CSR (Certificate Signing Request) generation, and PKI CA server. An RSA key pair is composed of a public-key and a private-key, and by default it is non-

https://t.me/learningnets

exportable for the security reason of protecting the private key. If the RSA key pair is tied to a certificate of the router, you might want to make the key pair exportable so you can have a safe offline backup of your certificate and keys (a certificate is useless without the associated private-key). After a key pair is generated, IOS stores the private-key in a system-protected area in NVRAM named **private-config**. The content of the area cannot be read, moved, renamed, copied, or deleted with specific commands like `more`, `rename`, `copy`, or `delete`; it is protected by the IOS subsystem.

To generate an RSA key-pair, your device must have configured a hostname and domain-name, because unless you specify a name for the key-pair, it is automatically assigned a name in the format of **hostname.domain**. You generate an RSA key-pair using the command `crypto key generate [general-keys|usage-keys] rsa` and the following options:

- `exportable`: Allows you to make an offline backup of the key-pair through export. By default, key-pair is not exportable.
- `modulus`: Allows you to specify the key strength, up to 4096 bits; unless specified, it uses a 512-bit size.
- `label`: Allows you to name the key-pair and use it for specific purposes afterward (you may want to use a key-pair for SSH and another key-pair for your CA).
- `storage`: Allows you to modify the storage of keys from the default of NVRAM; however, because of security concerns, FLASH is not a valid location.

An RSA key-pair can also be exported/imported from/to the router in PEM format using the command `crypto key [export|import] rsa <key-label> pem [terminal|url] [des|3des] <passphrase>`. Because the private-key information is highly sensitive and critical to the security of your PKI infrastructure, you may want to protect it through encryption via a configured string using the command `crypto key encrypt rsa name <key-label> passphrase <encryption-key>`. This does not affect its functionality on the router, it only encrypts it in the NVRAM. Note that if the router reloads, the key appears as locked and is not usable by the router; to unlock it, use the command** `crypto key unlock rsa name <key-label> passphrase <encryption-key>` **. You can also manually lock the key-pair (only if it has been encrypted) and make it unavailable to the router by using the command `crypto key lock rsa name <key-label> passphrase <encryption-key>`.

https://t.me/learningnets

You can verify your RSA public-key by using the command **`show crypto key mypubkey rsa <key-label>`** , and you can remove your RSA key-pair by using the command** `crypto key zeroize <key-label>` **. IOS version 15.0 added the public key authentication functionality for SSH. Basically, you can now authenticate on the router via SSH by using a username and your own RSA key-pair, without a password. For this to work you need:

- SSH server functionality configured on your router.
- SSH RSA public key authentication configured on your router.
- SSH client supporting RSA public key authentication (such as Putty or SecureCRT).
- RSA key-pair generated on each client and the associated public-key imported on the router.

You enable the SSH RSA public-key authentication and define the user by using the following commands, after which you paste its public-key in the router's config:

```
ip ssh pubkey-chain
  username <username>
   key-string
```

Configuration

Hostname and domain names have been preconfigured from the initial configuration files. Because we're asked to import the RSA key-pair on other routers if needed, we must mark it as exportable. We could ensure that SSH is disabled after a router reload through an EEM script, but this is not in the scope of the task or CCIE Security lab exam. Because SSH requires an active RSA key-pair to be functional, if we encrypt the private-key, after a router reload the RSA key-pair will be locked; therefore, SSH will be disabled.

https://t.me/learningnets

```
R3:

crypto key generate rsa general-keys exportable modulus 1024 label SSH
crypto key encrypt write rsa name SSH passphrase ciscocisco
ip ssh rsa keypair-name SSH
ip ssh version 2
!
!
ip ssh pubkey-chain
 username CCIE
  key-string
   exit
```

Verification

Verify the RSA public-key before being encrypted:

```
R3#show crypto key mypubkey rsa SSH
% Key pair was generated at: 15:01:34 UTC Apr 30 2013 Key name: SSH
Key type: RSA KEYS
 Storage Device: not specified
 Usage: General Purpose Key Key is exportable.

 Key Data:
  30819F30 0D06092A 864886F7 0D010101 05000381 8D003081 89028181 00DF116E
  BB3C9576 3C32F98C A4CA490E 53B79E01 B0FB4E35 9766DCE1 CC63BACA 59A2831A
  97C3EC03 F9FBE952 0BEF4737 967C93F3 186ABD29 26E3F792 FD5AD18F AC4C9FDF
  DB6C0A4E E058FC75 AD9C67E0 BB9E97F1 95999AF8 EFDE588F 15FB3D7F EDFEC248
  E3A29596 56E5BDFF 6E80051C A1633FF6 770CAB88 673D6684 D33DE840 07020301 0001
```

Verify the RSA public-key after being encrypted:

```
R3#show crypto key mypubkey rsa SSH
% Key pair was generated at: 15:01:34 UTC Apr 30 2013 Key name: SSH
Key type: RSA KEYS
 Storage Device: private-config
 Usage: General Purpose Key *** The key is protected and UNLOCKED. ***
Key is exportable.

 Key Data:
  30819F30 0D06092A 864886F7 0D010101 05000381 8D003081 89028181 00DF116E
  BB3C9576 3C32F98C A4CA490E 53B79E01 B0FB4E35 9766DCE1 CC63BACA 59A2831A
```

https://t.me/learningnets

```
  97C3EC03 F9FBE952 0BEF4737 967C93F3 186ABD29 26E3F792 FD5AD18F AC4C9FDF
  DB6C0A4E E058FC75 AD9C67E0 BB9E97F1 95999AF8 EFDE588F 15FB3D7F EDFEC248
  E3A29596 56E5BDFF 6E80051C A1633FF6 770CAB88 673D6684 D33DE840 07020301 0001
```

Verify that you can export the private-key from the router:

```
R3(config)#crypto key export rsa SSH pem terminal des ciscocisco
% Key name: SSH
   Usage: General Purpose Key
   Key data:-----BEGIN PUBLIC KEY-----
MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDfEW67PJV2PDL5jKTKSQ5Tt54B
sPtONZdm3OHMY7rKWaKDGpfD7AP5++lSC+9HN5Z8k/MYar0pJuP3kv1a0Y+sTJ/f
22wKTuBY/HWtnGfgu56X8ZWZmvjv3liPFfs9f+3+wkjjopWWVuW9/26ABRyhYz/2
dwyriGc9ZoTTPehABwIDAQAB-----END PUBLIC KEY-----
_-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
DEK-Info: DES-CBC,E189E28E7DE6A10C

gM0cxZ0QXRyxH2KpUcxSRbviXJH0TgwiIQAGmRvRULfjgRzyBUz9TsCLUXap0JbH
KrFDJXGtZnrsk1GoWAI8i+Q4AZy3Ss4wOcyYARXqGhu9v7UFsfBBHsJBvTOmBOuv
r786chbyn0F3gEVesiKjtLbxrNsjXNpcOW3piMjwnS2Rodbs/adDEcRhavP0tODN
L75jNIIK4RvQd504VPl7TFr9YeylDkqkQCY9ki0O766S/YUDqGpynd0pGmFfepw7
9VbahhVy0t+rvylhwWEfn9vZUSy3yHxEbhrOUFLKj2VcKDqI79+JWmYCA1gdwDgP
d/Zu95FS83aDhrj5jiMLPIA4DHV1bhUztrHO587E/iOHb9NtxRJHvcNZFnlK3pE+
Cc0CxhORhMwsGKAlh6MendOZ1GdamA9OVUvYRd5tR05bTojENSFg1yZ3nne6P21b
S0WM8r/LYZbyMAkP02WZPscXrgKyoJ1xzfrv+0ERCnE6ydOlfeNKC/px0WvhpLja
dpS97fEIfEy7XiIRQ3/j8tu0WV1UMdv3nHou4irjjdRXO7m1t9XOJ0RKF37cDseN
Ob1/jR80dhVPbtEymZ/DZHBLF+JupuvHWzq5u8b6fF/P0HFeHqKZm0QwNutgYY01
rZi8s2Tr+eh6ks93XZ/CA6lA5ay8j09MzTPPN6J66PFBONTRuiRv/DpKiOMMPDcf
9DMuj/BVL5o+3hKBsanr+5pEV9iywNNv9m7g6O3TNBJ+c0Sp8Wm8vgE3Oi9nr2GN
hKDkQNy3EvPLQ4OSmbc3LTwSCGQefhazYsUr8iDR7pJVWleOKHGcwg==-----END RSA PRIVATE KEY-----
```

Verify that SSH works before reload:

```
R1#ssh -l cisco 150.1.3.3
Password:
R3#
```

Verify that SSH is disabled after router reload:

```
R3#show ip ssh
SSH Disabled - version 2.0
```

https://t.me/learningnets

```
%Please create RSA keys to enable SSH (and of atleast 768 bits for SSH v2).
Authentication timeout: 120 secs; Authentication retries: 3
Minimum expected Diffie Hellman key size : 1024 bits
IOS Keys in SECSH format(ssh-rsa, base64 encoded): NONE
!
!R1#ssh -l cisco 150.1.3.3
% Connection refused by remote host
```

Decrypt the private key so that SSH becomes active again:

```
R3#configure terminal
R3(config)#crypto key decrypt rsa name SSH passphrase ciscocisco
WARNING: Configuration with decrypted key not saved.
  Please save it manually as soon as possible to
  save decrypted key
R3(config)#%SSH-5-ENABLED: SSH 2.0 has been enabled

R3(config)#
```

Verify SSH service again:

```
R3#show ip ssh
SSH Enabled - version 2.0
Authentication timeout: 120 secs; Authentication retries: 3
Minimum expected Diffie Hellman key size : 1024 bits
IOS Keys in SECSH format(ssh-rsa, base64 encoded):
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAAAgQC/bi3t+HGdj47oz+1I33wpJKnoCgRco6fA06Gmspzk
SKDxRAL33V9ic1n2ygeNrnZSQ3MUVmFw7qkjupiidqFER5OoqUOwHjjJznuKXetpudHFYlctSrseQ5ok
H23dK1+YlAvqNQTm6G6JJpN6dgPfG9uTbgKh/yA/ERYnXVNp4w==
!
!R1#ssh -l cisco 150.1.3.3
Password:
R3#
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

PKI Server on IOS Router

Task

- Configure R2 as a root CA with the following requirements:
  - R2 should host and publish the CRL.
  - Insert its FQDN in the Subject Name field.
  - All certificates should be signed using a 160-bit hash function.
  - All certificates should have a lifetime equal to double the default.
  - Ensure that you can revoke certificates at any time if needed.
  - Ensure an automatic backup of CA RSA keys and its self-signed certificate.
  - All certificates should be automatically renewed 36 days before expiring.
  - All client certificates should be automatically granted.
- Configure R3 as an RA with the following requirements:
  - Insert its FQDN in the Subject Name field.
  - Use an RSA key-pair with 1024-bit strength named **RA**.
  - Do not include SN or IP address information in the RA certificate.
- SW1 should request a certificate from RA and use it for its HTTPS service:
  - Insert its FQDN in the Subject Name field.
  - Do not include SN or IP address information in the certificate.
- Use R2 as the NTP server and ensure that SCEP traffic uses only Loopback0 interfaces.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **PKI Infrastructure** module of this section, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the

https://t.me/learningnets

devices. This command, however, is not supported by the ASA.

Overview

Cisco IOS does not currently support a CA server function with Elliptic Curve Digital Signature Algorithm (ECDS) certificates, which is required for Suite-B cryptography (Next Generation Encryption). The Suite-B protocols bring together enhanced cryptographic security and increased encryption performance:

- RSA keys are replaced by ECDSA keys, so PKI infrastructure must support ECDSA signatures.
- DH groups are replaced by ECDH-19, ECDH-20, and ECDH-21 groups.
- AES-CBC mode is replaced by AES-GCM or AES-GMAC; GCM supports integrated encryption and integrity, whereas GMAC supports only integrity.
- SHA-1 is replaced by SHA-256, SHA-384, or SHA-512.

Complete support for Suite-B protocols was added to IOS in release version 15.1(2)T, and depending on the platform, it may be supported or not supported in hardware; however, not all VPN features (GETVPN or FlexVPN) support it with this IOS. In accordance with RFC 6380, Suite-B protocols are supported only with IKEv2 and PKI (ECDSA signed certificates) authentication, so no IKEv1 or RSA/pre-shared-key authentication. For this reason, Suite-B will not be present in the current version of the CCIE Security lab exam, so it will not be discussed here.

The IOS CA function is more complex than the ASA CA implementation:

- Requires NTP for time synchronization to ensure that IOS does not disable its functionality.
- It supports a hierarchical CA system: an IOS CA can act as root CA, subordinate CA, or registration authority.
- It supports CA High Availability configuration (one router acting as the active CA and one router acting as the standby CA).
- It supports a single active CA instance, so it does not support multiple CA roles concurrently (such as being root in one PKI and RA in another PKI).
- Certificate database and configuration files can be kept locally on NVRAM (by default) or on external storage.
- Client certificate enrollment can be done manually through PKCS#10 CSR, or automatically through SCEP (HTTP server functionality must be enabled separately for SCEP to become active).
- The CRL list can be stored locally and retrieved by clients through SCEP or on

https://t.me/learningnets

external HTTP server (by default, the CRL list is not published by the router).

IOS CA stores information about issued certificates based on the configured level:

- **minimal**. This is the default mode. Enough information is kept only to ensure that unique certificates can be issued upon any request.
- **names**. In this mode, serial-number and Subject Name fields are stored for each issued certificate, allowing you to revoke a certificate when necessary.
- **complete**. All issued certificates are stored in the database; depending on the number of your clients, it is recommended that you configure the database on external storage.

The IOS CA server is configured using the command `crypto pki server <CA-name> [rollover]` and by default is disabled, but it does not require any optional parameters to be configured for a basic functionality. To enable it, use the command `no shutdown` . Although some optional CA parameters can be configured after CA is enabled (you must suspend the CA temporarily for it), it does not apply for all (like the RSA key-pair to be used), so it is recommended that you fully configure the CA before activating it. With the `rollover` keyword, you can create a shadow/backup certificate, which will be automatically used when the current/active CA self-signed certificate expires.

Because all certificates have a validity period, you don't want a network down incident because all of your certificates used for IPsec VPN's have expired. To prevent this, ensure that all certificates (Root CA, Subordinate CA, RA, and client certificates) are automatically refreshed (re-issued) before expiring:

- Configure all CA levels (Root, Subordinate, RA) with the `auto-rollover` command.
- Configure the Root CA with the `grant auto rollover ca-cert` command for the Subordinate CA certificate to be automatically granted on renewal requests.
- Configure the Root CA with the `grant auto rollover ra-cert` command for the RA certificate to be automatically granted on renewal requests.
- Configure all clients with the `auto-enroll` command to automatically send a certificate renewal request to the CA.
- All the above work only for SCEP-issued certificates, not for manually enrolled clients.

The IOS CA configuration can be modified through the following optional commands:

- `cdp-url` allows you to publish and make use of the CRL file; to host it on the router itself, use the dedicated URL `http://<IP Address>/cgi-bin/pkiclient.exe?operation=GetCRL` .

https://t.me/learningnets

- `eku` allows you to specify one or multiple scopes for the digital certificates. EKU stands for Extended Key Usage and basically defines the purpose(s) for which a certificate can be used. For example, it can be client-authentication or server-authentication. Depending on the technology being used, certain EKU values may be required. For example, for IKEv2 using digital certificates, IOS requires its issued certificate to have an EKU of `server authentication .`
- `hash` allows you to modify the default hash algorithm used for signing certificates, which is MD5.
- `mode [ra|sub-cs]` allows you to specify the CA server role in a hierarchical deployment.
- `issuer-name` allows you to modify the default DN value of the CA, which is equal to the RSA keypair used or the CA name if no keypair has been specified.
- `lifetime [ca-certificate|certificate]` allows you to modify the default values of 3 years for CA certificate and 1 year for client certificates.

Configuration

Based on Root CA configuration task requirements:

- SHA1 uses a 160-bit hash function.
- To revoke certificates if needed, database level must be `names .`
- The CDP URL uses the special character `?` , so make sure you use the `CTRL+V` command sequence when typing it.
- An IOS CA can automatically grant client certificates for directly received requests or for requests relayed through an RA, but not both concurrently; to automatically grant client certificates requests received through an RA, the RA certificate must be manually granted.
- If auto-rollover is configured on the RA, its certificate renewal request is automatically granted by the CA.
- To create the automatic backup (PKCS#12 or PEM) of the CA keys and its certificate after being refreshed, it must be protected by a password.

Based on RA configuration task requirements:

https://t.me/learningnets

- The RSA keys to be used by the RA when it requests a certificate from the root CA can be specified only under the `trustpoint .`
- Enable automatic authorization of client-certificate requests with `grant auto .`
- Ensure that `auto-rollover` is configured so that the RA automatically requests a new certificate for itself from the Root CA.

Configure the Root CA on R2 and RA role on R3:

```
R2:
ip http server
ntp master
!
crypto pki server ROOT
 hash sha1
 grant ra-auto
 grant auto rollover ra-cert
 lifetime certificate 730
 lifetime ca-certificate 2190
 database level names
 issuer-name CN=R2.ine.com
 cdp-url http://150.1.2.2/cgi-bin/pkiclient.exe?operation=GetCRL
 database archive pkcs12 password ciscocisco
 auto-rollover 36
 no shutdown

R3:

ip http server
ntp server 150.1.2.2
!
crypto key generate rsa general-keys modulus 1024 label RA
!
crypto pki trustpoint RA
 rsakeypair RA
 enrollment url http://150.1.2.2
 source interface Loopback0
 ip-address none
 serial-number none
 subject-name CN=R3.ine.com , OU=ioscs RA
!
crypto pki server RA
 mode ra
 grant auto
 auto-rollover 36
 database archive pkcs12 password ciscocisco
```

https://t.me/learningnets

```
 no shutdown
```

Before enabling the RA service on R3, wait for NT to synchronize between the Root CA and RA. After enabling the RA with the **`no shutdown`** command, you will be asked to configure a password for protecting the private key; accept the Root CA certificate and request a certificate from the Root CA:

```
R3(cs-server)#no shutdown
%Some server settings cannot be changed after CA certificate generation.
% Please enter a passphrase to protect the private key
% or type Return to exit Password:
 Re-enter password:
% Generating 1024 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 4 seconds)

%SSH-5-ENABLED: SSH 1.99 has been enabled
Certificate has the following attributes:
       Fingerprint MD5: 44880B96 5264C65E A383300A C316BDE7
      Fingerprint SHA1: 9F5FA127 E6683A01 27414587 1B139476 DE24B57B
 % Do you accept this certificate? [yes/no]: yes
Trustpoint CA certificate accepted.%
% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
   password to the CA Administrator in order to revoke your certificate.
   For security reasons your password will not be saved in the configuration.
   Please make a note of it.
 Password:
Re-enter password:

% The subject name in the certificate will include: CN=R3.ine.com,OU=ioscs RA
% The subject name in the certificate will include: R3.ine.com
 Request certificate from CA? [yes/no]: yes
% Certificate request sent to Certificate Authority
% The 'show crypto pki certificate verbose RA' command will show the fingerprint.
 % Enrollment in progress...
R3(cs-server)#
CRYPTO_PKI:  Certificate Request Fingerprint MD5: E5A82879 9F74BB13 956C2246 5C745FDE
CRYPTO_PKI:  Certificate Request Fingerprint SHA1: 528AAA96 84F4A8ED 6FFF8D7B 6EBA38AB 831C0076
!
!R3#show crypto pki certificates
CA Certificate
  Status: Available
  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
  Issuer:
```

https://t.me/learningnets

```
    cn=R2.ine.com
  Subject:
    cn=R2.ine.com
  Validity Date:
    start date: 18:03:08 UTC Oct 3 2013
    end   date: 18:03:08 UTC Oct 2 2019
  Associated Trustpoints: RA

Certificate
  Subject: Name: R3.ine.com
Status: Pending

   Key Usage: General Purpose
   Certificate Request Fingerprint MD5: E5A82879 9F74BB13 956C2246 5C745FDE
   Certificate Request Fingerprint SHA1: 528AAA96 84F4A8ED 6FFF8D7B 6EBA38AB 831C0076
   Associated Trustpoint: RA
```

After both root CA and RA have been configured and enabled, and before configuring SW1 as a client, make sure to manually grant the RA certificate on the root CA. You can visualize pending requests on the root CA using the command **`show crypto pki server ROOT requests`**. Because the RA request should be the first and only at this point, grant the certificate using the command **`crypto pki server ROOT grant 1`**.

```
R2#show crypto pki server ROOT requests
Enrollment Request Database:

Subordinate CA certificate requests:
ReqID  State      Fingerprint                      SubjectName
--------------------------------------------------------------

RA certificate requests:
ReqID  State      Fingerprint                      SubjectName
--------------------------------------------------------------
1      pending    E5A828799F74BB13956C22465C745FDE hostname=R3.ine.com,cn=R3.ine.com,ou=ioscs RA

Router certificates requests:
ReqID  State      Fingerprint                      SubjectName
--------------------------------------------------------------
!
!R2#crypto pki server ROOT  grant 1
```

https://t.me/learningnets

When the RA certificate has been granted, the following logs will appear on R3's console:

```
R3(cs-server)#% Exporting Certificate Server signing certificate and keys...
%PKI-6-CERTRET: Certificate received from Certificate Authority
R3(cs-server)#%PKI-6-CS_ENABLED: Certificate server now enabled.
```

Configure SW1 to request its certificate from RA:

SW1:

```
ntp server 150.1.2.2
!
crypto pki trustpoint R3-R2
 auto-enroll 95
 enrollment url http://150.1.3.3
 source interface Loopback0
 ip-address none
 serial-number none
 subject-name CN=SW1.ine.com
!
!
crypto pki authenticate R3-R2
crypto pki enroll R3-R2
!
ip http secure-trustpoint R3-R2
ip http secure-server
```

Verification

Verify root CA configuration and certificates:

```
R2#show ntp status
Clock is synchronized, stratum 8, reference is 127.127.1.1
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**21
reference time is D52E4180.C8A58056 (14:00:32.783 UTC Fri May 3 2013)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.33 msec, peer dispersion is 0.23 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s
system poll interval is 16, last update was 8 sec ago.
!
!R2#show crypto pki server
```

https://t.me/learningnets

```
Certificate Server ROOT: Status: enabled
 State: enabled
    Server's configuration is locked  (enter "shut" to unlock it) Issuer name: CN=R2.ine.com
    CA cert fingerprint: EF4A0416 022EAE18 FEB92E2A E71132C4
 Granting mode is: auto for RA-authorized requests, manual otherwise
    Last certificate issued serial number (hex): 3
    CA certificate expiration timer: 13:47:23 UTC May 2 2019
    CRL NextUpdate timer: 19:47:26 UTC May 3 2013
    Current primary storage dir: nvram:
    Database Level: Names - subject name data written as <serialnum>.cnm
    Auto-Rollover configured, overlap period 36 days
    Autorollover timer: 13:47:22 UTC Mar 27 2019
!
!R2#show crypto pki certificates
 CA Certificate
 Status: Available

  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
  Issuer:
    cn=R2.ine.com
  Subject:
    cn=R2.ine.com
  Validity Date:
    start date: 13:47:23 UTC May 3 2013
    end   date: 13:47:23 UTC May 2 2019
  Associated Trustpoints: ROOT
  Storage: nvram:Rack1R2ineco#1CA.cer
```

Note that there are three issued certificates: root CA, RA, and SW1. However, detailed data is not available because the database level is not set to `complete` :

https://t.me/learningnets

```
R2#show crypto pki server ROOT certificates

Serial Issued date              Expire date              Subject Name
1       <cert file not accessible>
2       <cert file not accessible>
3       <cert file not accessible>

!
!R2#show crypto pki server ROOT crl

Certificate Revocation List: Issuer: cn=R2.ine.com

    This Update: 13:47:26 UTC May 3 2013
    Next Update: 19:47:26 UTC May 3 2013 Number of CRL entries: 0

    CRL size: 225 bytes
```

Verify RA configuration and its certificate:

```
R3#show ntp status
Clock is synchronized, stratum 9, reference is 150.1.2.2
nominal freq is 250.0000 Hz, actual freq is 249.9996 Hz, precision is 2**24
reference time is D52E4150.7B3D469F (13:59:44.481 UTC Fri May 3 2013)
clock offset is -9.1545 msec, root delay is 1.51 msec
root dispersion is 17.71 msec, peer dispersion is 4.46 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000001398 s/s
system poll interval is 64, last update was 413 sec ago.
!
!R3#show crypto pki server
Certificate Server RA: Status: enabled
State: enabled
    Server's configuration is locked  (enter "shut" to unlock it)
Issuer name: CN=R3.ine.com , OU=ioscs RA
    CA cert fingerprint: EF4A0416 022EAE18 FEB92E2A E71132C4 Server configured in RA mode
    RA cert fingerprint: 1456A638 3C96B989 1063C7DC 532D961E Granting mode is: auto
    Current primary storage dir: nvram:
    Database Level: Minimum - no cert data written to storage
    Auto-Rollover configured, overlap period 36 days
    Autorollover timer: 13:50:57 UTC Mar 28 2015
!
!R3#show crypto pki certificates
Certificate (RA mode CS certificate)
Status: Available
```

https://t.me/learningnets

```
  Certificate Serial Number (hex): 02
  Certificate Usage: General Purpose
  Issuer:
    cn=R2.ine.com
  Subject:
    Name: R3.ine.com
    hostname=R3.ine.com
    cn=R3.ine.com
    ou=ioscs RA
  CRL Distribution Points: http://150.1.2.2/cgi-bin/pkiclient.exe?operation=GetCRL
  Validity Date:
    start date: 13:50:58 UTC May 3 2013
    end   date: 13:50:58 UTC May 3 2015
  Associated Trustpoints: RA
CA Certificate
Status: Available

  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
  Issuer:
    cn=R2.ine.com
  Subject:
    cn=R2.ine.com
  Validity Date:
    start date: 13:47:23 UTC May 3 2013
    end   date: 13:47:23 UTC May 2 2019
  Associated Trustpoints: RA
```

Note that the RA does not issue certificates, it just validates the requests and forwards them to the CA:

```
R3#show crypto pki server RA certificates
%Certificate Server is in RA Mode and does not issue certificates.
```

Verify SW1 configuration. Although the enrollment URL points to the RA, notice that its certificate is issued by the root CA; also, because of the auto-enroll configuration, SW1's certificate has a renew date that is 36 days before the expire date of its current certificate.

```
SW1#show ntp status
Clock is synchronized, stratum 9, reference is 150.1.2.2
nominal freq is 119.2092 Hz, actual freq is 119.2105 Hz, precision is 2**17
reference time is D52E4330.9ED5A369 (14:07:44.620 UTC Fri May 3 2013)
clock offset is 0.2780 msec, root delay is 2.73 msec
```

https://t.me/learningnets

```
root dispersion is 0.87 msec, peer dispersion is 0.15 msec
!
!
!SW1#show crypto pki trustpoints status
Trustpoint R3-R2: Issuing CA certificate configured:
    Subject Name:
     cn=R2.ine.com
    Fingerprint MD5: EF4A0416 022EAE18 FEB92E2A E71132C4
    Fingerprint SHA1: B40E0BBB D375F873 8050128A 305CAC81 32E40AB8
Router General Purpose certificate configured:
    Subject Name:
     hostname=SW1.ine.com,cn=SW1.ine.com
    Fingerprint MD5: 9CBF8994 946012CD 93594AB3 A4455D3D
    Fingerprint SHA1: 56A68D64 8E0784FA 02E38F90 F6861C03 2A197F54 Last enrollment status: Granted
  Next enrollment attempt:
    01:53:28 UTC Mar 28 2015
    * Configuration will not be saved after enrollment *
  State:
    Keys generated ..... Yes (General Purpose, non-exportable)
    Issuing CA authenticated ..... Yes
    Certificate request(s) ..... Yes
!
!SW1#show crypto pki certificates
Certificate
 Status: Available
  Certificate Serial Number: 03
  Certificate Usage: General Purpose
  Issuer:
    cn=Rack1R2.ine.com
  Subject:
    Name: SW1.ine.com
    hostname=SW1.ine.com
    cn=SW1.ine.com
  CRL Distribution Points: http://150.1.2.2/cgi-bin/pkiclient.exe?operation=GetCRL
  Validity Date:
    start date: 13:53:28 UTC May 3 2013 end   date: 13:53:28 UTC May 3 2015
 renew date: 01:53:28 UTC Mar 28 2015
  Associated Trustpoints: R3-R2
 CA Certificate
 Status: Available
  Certificate Serial Number: 01
  Certificate Usage: Signature Issuer:
 cn=R2.ine.com

  Subject:
    cn=R2.ine.com
```

https://t.me/learningnets

```
  Validity Date:
    start date: 13:47:23 UTC May 3 2013
    end   date: 13:47:23 UTC May 2 2019
  Associated Trustpoints: R3-R2
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

PKI Server High Availability on IOS Router

Task

- Configure R1 and R2 to function as redundant CA servers.
- In the event of a reload, R1 should always become the active router.
- Insert **ROOT-HA.ine.com** in the Subject field of the CA certificate.
- Ensure that client certificates are automatically approved.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **PKI Infrastructure** module of this section, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Cisco IOS PKI can be deployed in a High Availability mode, providing redundancy for client requests. Like other technologies supported by IOS in HA mode, such as Zone Based Firewall (ZBPF) or IPsec, PKI HA uses the Stateful Switch-Over (SSO) redundancy feature. This inter-device redundancy function relies on two protocols, HSRP and SCTP:

- HSRP determines the roles ACTIVE and STANDBY.
- SCTP ensures automatic synchronization between ACTIVE and STANDBY. For PKI, the following are automatically synchronized from the ACTIVE:
  - CA server configuration
  - CA certificate
  - CA specific files from NVRAM: CRL, issued certificates, file keeping serial

https://t.me/learningnets

numbers
- RSA keys

To ensure functionality of the IOS PKI High Availability deployment, it is recommended that you use the following configuration steps:

1. Configure and verify HSRP functionality.
2. Configure and verify inter-device SSO redundancy functionality (requires a manual reload on the STANDBY device).
3. Do not continue further unless SSO is functional.
4. Configure and activate PKI server on the ACTIVE device.
5. Disable PKI server on the ACTIVE device and enable PKI redundancy.
6. Activate PKI server on the ACTIVE device.

> The High Availability configuration from the PKI Configuration Guide of IOS 15MT is found in the section "Configuring Authorization and Revocation of Certificates in a PKI."

> Because of the high volume of data required to be synchronized, if the CA runs in complete database level, the client-issued certificate files (.crt) will not be synchronized with the standby system. The workaround is to have both CA systems point to a common external storage for these files, by using the command **`database url`** .

Configuration

R1:

```
ip http server
!
interface GigabitEthernet0/0
 standby ip 136.1.18.12
 standby priority 150
 standby preempt
 standby name PKI
!
ipc zone default
 association 1
  no shutdown
  protocol sctp
  local-port 5000
  local-ip 136.1.18.1
  remote-port 5000
```

https://t.me/learningnets

```
  remote-ip 136.1.18.2
!
redundancy inter-device
 scheme standby PKI
```

R2:

```
ip http server
!
interface GigabitEthernet0/0
 standby ip 136.1.18.12
 standby preempt
 standby name PKI
!
ipc zone default
 association 1
  no shutdown
  protocol sctp
  local-port 5000
  local-ip 136.1.18.2
  remote-port 5000
  remote-ip 136.1.18.1
!
redundancy inter-device
 scheme standby PKI
```

At this point, we need to save the configuration and reload the both devices to activate the redundancy. When R2 detects itself as standby, it will force itself another reload. The following output shows the initial required reload:

```
R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_INIT
Pending Scheme: Standby (Will not take effect until next reload)

      Pending Groupname: PKI
  Scheme: <NOT CONFIGURED>
  Peer present: UNKNOWN
  Security: Not configured
```

After SSO is functional, configure PKI only on the ACTIVE device; it will be automatically synchronized to the STANDBY:

R1:

https://t.me/learningnets

```
crypto key generate rsa general-keys redundancy label HA modulus 1024
!
crypto pki server HA
 database level names
 issuer-name CN=ROOT-HA.ine.com
 database archive pkcs12 password ciscocisco
 grant auto
 no shutdown
!
crypto pki server HA
 shutdown
 redundancy
 no shutdown
```

> If PKI functionality is not synchronized as shown in the Verification section, it may be required to perform another reload of both routers.

Verification

First, verify SSO inter-device redundancy:

```
R1#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_ACT
  Scheme: Standby Groupname: PKI Group State: Active
  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
!
!R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_STDBY
  Scheme: Standby Groupname: PKI Group State: Standby
  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
!
!R1#show redundancy states
 my state = 13 -ACTIVE
 peer state = 8  -STANDBY HOT
           Mode = Duplex
        Unit ID = 0

     Maintenance Mode = Disabled
    Manual Swact = enabled Communications = Up

   client count = 13
 client_notification_TMR = 60000 milliseconds
```

https://t.me/learningnets

```
           RF debug mask = 0x0
!
!R2#show redundancy states
my state = 8  -STANDBY HOT
peer state = 13 -ACTIVE
           Mode = Duplex
        Unit ID = 0

     Maintenance Mode = Disabled
    Manual Swact = cannot be initiated from this the standby unit Communications = Up


   client count = 13
 client_notification_TMR = 60000 milliseconds
           RF debug mask = 0x0
```

Verify that the PKI server is enabled on both routers, but R1 is the current ACTIVE PKI server:

```
R1#show crypto pki server
Certificate Server HA: Status: enabled
 State: enabled
    Server's configuration is locked  (enter "shut" to unlock it)
    Issuer name: CN=ROOT-HA.ine.com
    CA cert fingerprint: E6B92BCE A2AF733C 3EC48E05 D9873F52
    Granting mode is: auto
    Last certificate issued serial number (hex): 1
    CA certificate expiration timer: 14:51:51 UTC Oct 3 2016
    CRL NextUpdate timer: 20:51:54 UTC Oct 4 2013
    Current primary storage dir: nvram:
    Database Level: Names - subject name data written as <serialnum>.cnm
Redundancy configured.  This is active.
!
!R2#show crypto pki server
Certificate Server HA: Status: enabled
 State: enabled
    Server's configuration is locked  (enter "shut" to unlock it)
    Issuer name: cn=ROOT-HA.ine.com
    CA cert fingerprint: E6B92BCE A2AF733C 3EC48E05 D9873F52
    Granting mode is: auto
    Last certificate issued serial number (hex): 1
    CA certificate expiration timer: 14:51:51 UTC Oct 3 2016
    CRL NextUpdate timer: 20:51:54 UTC Oct 4 2013
    Current primary storage dir: nvram:
    Database Level: Names - subject name data written as <serialnum>.cnm
```

https://t.me/learningnets

```
Redundancy configured.  This is standby.
```

Verify the CA certificate on both routers:

https://t.me/learningnets

```
R1#show crypto pki certificates
CA Certificate
 Status: Available
 Certificate Serial Number (hex): 01

  Certificate Usage: Signature
  Issuer:
    cn=ROOT-HA.ine.com
  Subject:
    cn=ROOT-HA.ine.com
  Validity Date:
    start date: 14:51:51 UTC Oct 4 2013
    end   date: 14:51:51 UTC Oct 3 2016
  Associated Trustpoints: HA
  Storage: nvram:ROOT-HAineco#1CA.cer
!
!R2#show crypto pki certificates
CA Certificate
 Status: Available
 Certificate Serial Number (hex): 01

  Certificate Usage: Signature
  Issuer:
    cn=ROOT-HA.ine.com
  Subject:
    cn=ROOT-HA.ine.com
  Validity Date:
    start date: 14:51:51 UTC Oct 4 2013
    end   date: 14:51:51 UTC Oct 3 2016
  Associated Trustpoints: HA
  Storage:
!
!R1#show crypto pki server HA certificates

Serial Issued date              Expire date              Subject Name1
       <cert file not accessible>
Certificate might have been granted by other CA
!
!R2#show crypto pki server HA certificates

Serial Issued date              Expire date              Subject Name1
       <cert file not accessible>
```

https://t.me/learningnets

```
Certificate might have been granted by other CA
```

Enroll SW1 in the PKI infrastructure with R1 being the ACTIVE router (you may need to synchronize time with NTP between SW1 and R1/R2 if there is time difference between the internal clocks):

```
R1:
ntp master

R2:
ntp server 150.1.1.1

SW1:

ntp server 150.1.1.1
!
crypto pki trustpoint HA
 enrollment url http://136.1.18.12
!
!
crypto pki authenticate HA
crypto pki enroll HA
```

Verify that SW1 received the certificate and R2 is synchronized with R1 for the certificate database:

```
SW1#show crypto pki certificates
Certificate
Status: Available
Certificate Serial Number: 02
  Certificate Usage: General Purpose
  Issuer:
    cn=ROOT-HA.ine.com
  Subject:
    Name: SW1.ine.com
    hostname=SW1.ine.com
  Validity Date:
    start date: 17:55:45 UTC Oct 4 2013
    end   date: 17:55:45 UTC Oct 4 2014 Associated Trustpoints: HA

CA Certificate
  Status: Available
  Certificate Serial Number: 01
  Certificate Usage: Signature
  Issuer:
```

https://t.me/learningnets

```
    cn=ROOT-HA.ine.com
  Subject:
    cn=ROOT-HA.ine.com
  Validity Date:
    start date: 14:51:51 UTC Oct 4 2013
    end   date: 14:51:51 UTC Oct 3 2016
  Associated Trustpoints: HA
!
!R1#show crypto pki server HA certificates
Serial Issued date              Expire date               Subject Name
1       <cert file not accessible>
Certificate might have been granted by other CA 2       <cert file not accessible>
Certificate might have been granted by other CA
!
!R2#show crypto pki server HA certificates
Serial Issued date              Expire date               Subject Name
1       <cert file not accessible>
Certificate might have been granted by other CA 2       <cert file not accessible>

Certificate might have been granted by other CA
```

Move the HSRP ACTIVE role to R2, and re-enroll SW1 in the PKI (when the ACTIVE role changes, the STANDBY always issues a forced self-reload to ensure synchronization):

```
R2:
interface gigabitEthernet0/0
 standby priority 200

SW1:

crypto pki enroll HA
```

Verify that the new certificate was received by SW1, and verify the serial number difference:

```
SW1#show crypto pki certificates
Certificate
 Status: Available
 Certificate Serial Number: 03

  Certificate Usage: General Purpose
  Issuer:
    cn=ROOT-HA.ine.com
  Subject:
```

https://t.me/learningnets

```
    Name: SW1.ine.com
    hostname=SW1.ine.com
  Validity Date:
    start date: 18:00:29 UTC Oct 4 2013
    end   date: 18:00:29 UTC Oct 4 2014
  Associated Trustpoints: HA

CA Certificate
  Status: Available
  Certificate Serial Number: 01
  Certificate Usage: Signature
  Issuer:
    cn=ROOT-HA.ine.com
  Subject:
    cn=ROOT-HA.ine.com
  Validity Date:
    start date: 14:51:51 UTC Oct 4 2013
    end   date: 14:51:51 UTC Oct 3 2016
  Associated Trustpoints: HA
```

After R1 reloads, verify that R2 is now the ACTIVE router/PKI server and that the certificate database was synchronized:

```
R2#show crypto pki server HA
Certificate Server HA: Status: enabled
 State: enabled
    Server's configuration is locked  (enter "shut" to unlock it)
    Issuer name: cn=ROOT-HA.ine.com
    CA cert fingerprint: E6B92BCE A2AF733C 3EC48E05 D9873F52
    Granting mode is: auto
    Last certificate issued serial number (hex): 3
    CA certificate expiration timer: 14:51:51 UTC Oct 3 2016
    CRL NextUpdate timer: 20:51:54 UTC Oct 4 2013
    Current primary storage dir: nvram:
    Database Level: Names - subject name data written as <serialnum>.cnm
Redundancy configured.  This is active.
!
!R1#show crypto pki server HA
Certificate Server HA: Status: enabled
 State: enabled
    Server's configuration is locked  (enter "shut" to unlock it)
    Issuer name: CN=ROOT-HA.ine.com
    CA cert fingerprint: E6B92BCE A2AF733C 3EC48E05 D9873F52
    Granting mode is: auto
    Last certificate issued serial number (hex): 3
```

https://t.me/learningnets

```
    CA certificate expiration timer: 14:51:51 UTC Oct 3 2016
    CRL NextUpdate timer: 20:51:54 UTC Oct 4 2013
    Current primary storage dir: nvram:
    Database Level: Names - subject name data written as <serialnum>.cnm
Redundancy configured.  This is standby.
!
!R2#show crypto pki server HA certificates
Serial Issued date              Expire date               Subject Name
1       <cert file not accessible>
Certificate might have been granted by other CA
2       <cert file not accessible>
Certificate might have been granted by other CA 3       <cert file not accessible>
Certificate might have been granted by other CA
!
!R1#show crypto pki server HA certificates
Serial Issued date              Expire date               Subject Name
1       <cert file not accessible>
Certificate might have been granted by other CA
2       <cert file not accessible>
Certificate might have been granted by other CA 3       <cert file not accessible>
x
Certificate might have been granted by other CA
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

PC Enrollment with ASA PKI

Task

- Configure ASA4 as a CA server with a CN of **Rack1ASA4.ine.com**.
- Enroll the user **rack1.test-pc** in the PKI infrastructure from TEST-PC-A.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **PKI Infrastructure** module of this section, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The CA detailed configuration on the ASA can be found in Certificate Authority on Cisco ASA. A basic setup of the CA functionality is used in this task.

Configuration

ASA4:

```
crypto ca server
 issuer-name CN=Rack1ASA4.ine.com
 no shutdown
!
webvpn
 enable VLAN18
!
```

https://t.me/learningnets

```
crypto ca server user-db add rack1.test-pc dn CN=rack1.test-pc
crypto ca server user-db allow user rack1.test-pc
```

Verification

Display the OTP (one-time password) for the user:

```
ASA4(config)# crypto ca server user-db show-otp rack1.test-pc
Username: rack1.test-pc
OTP: 81D882AC5287922A

Enrollment Allowed Until: 16:19:38 UTC Sat May 11 2013
```

Using Internet Explorer from TEST-PC-A, access the ASA on the URL **https://136.1.18.8/+CSCOCA+/enroll.html**. The following page should appear; enter the username and OTP. If successfully authenticated, you will receive the certificate in PKCS#12 format (which contains the private key, user certificate, and CA certificate), named **rack1.test-pc.p12**, and you will be asked to save it:



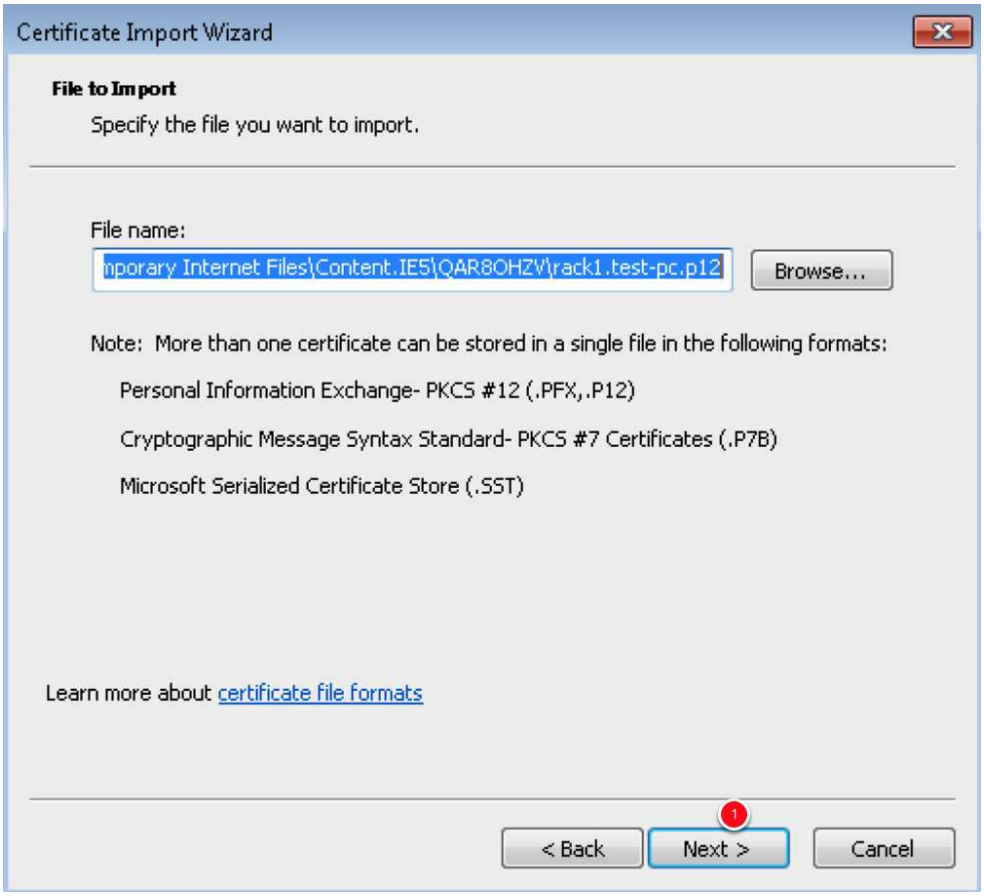




Open the file, and the Certificate Import Wizard starts:

https://t.me/learningnets

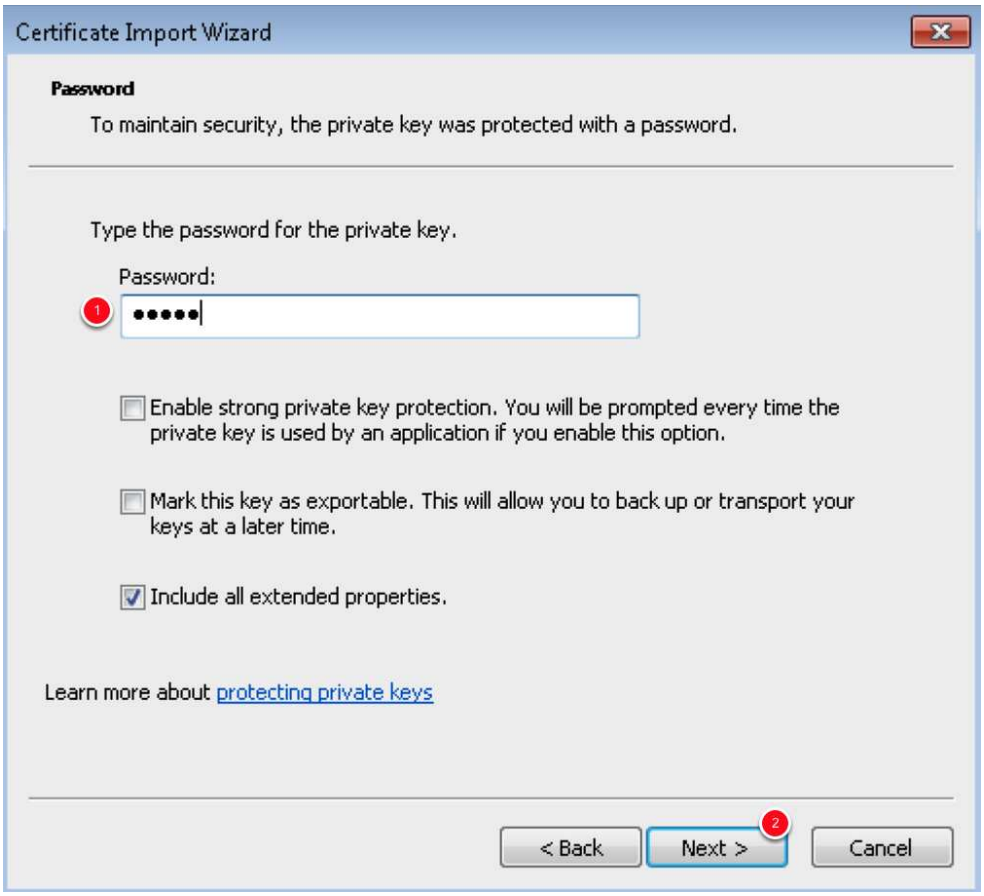




Enter the OTP again, which is also used to protect the private-key. If you want the

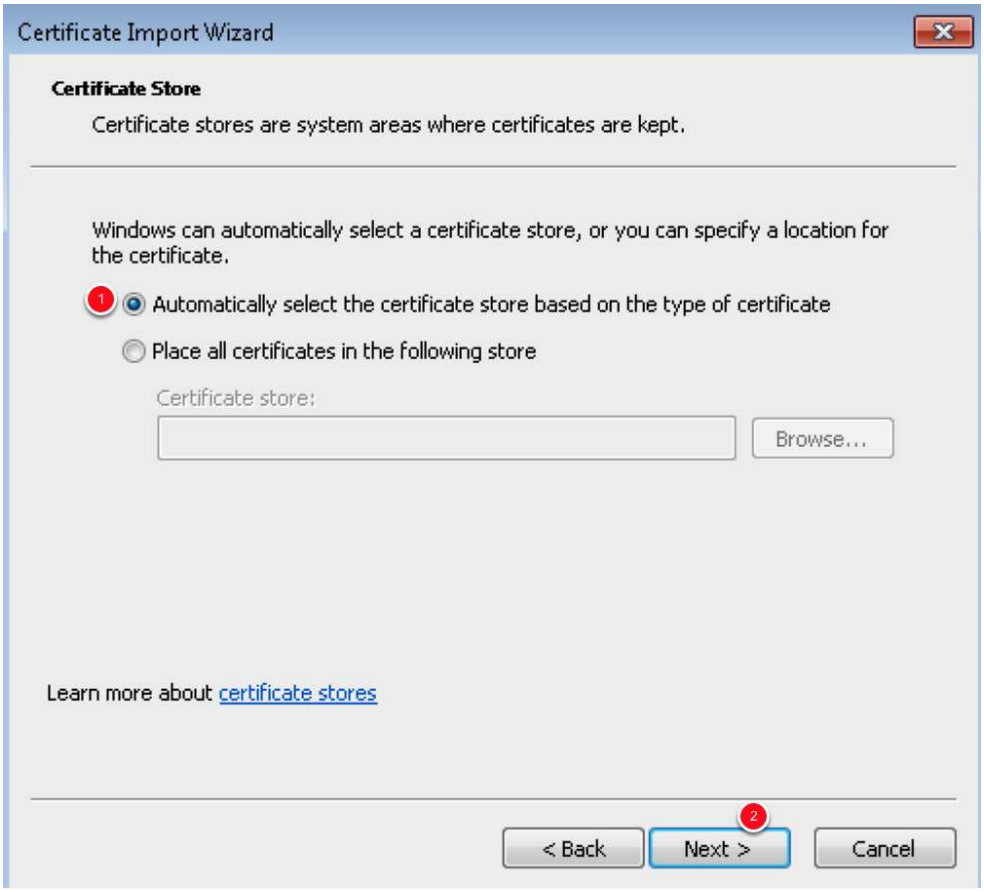
https://t.me/learningnets

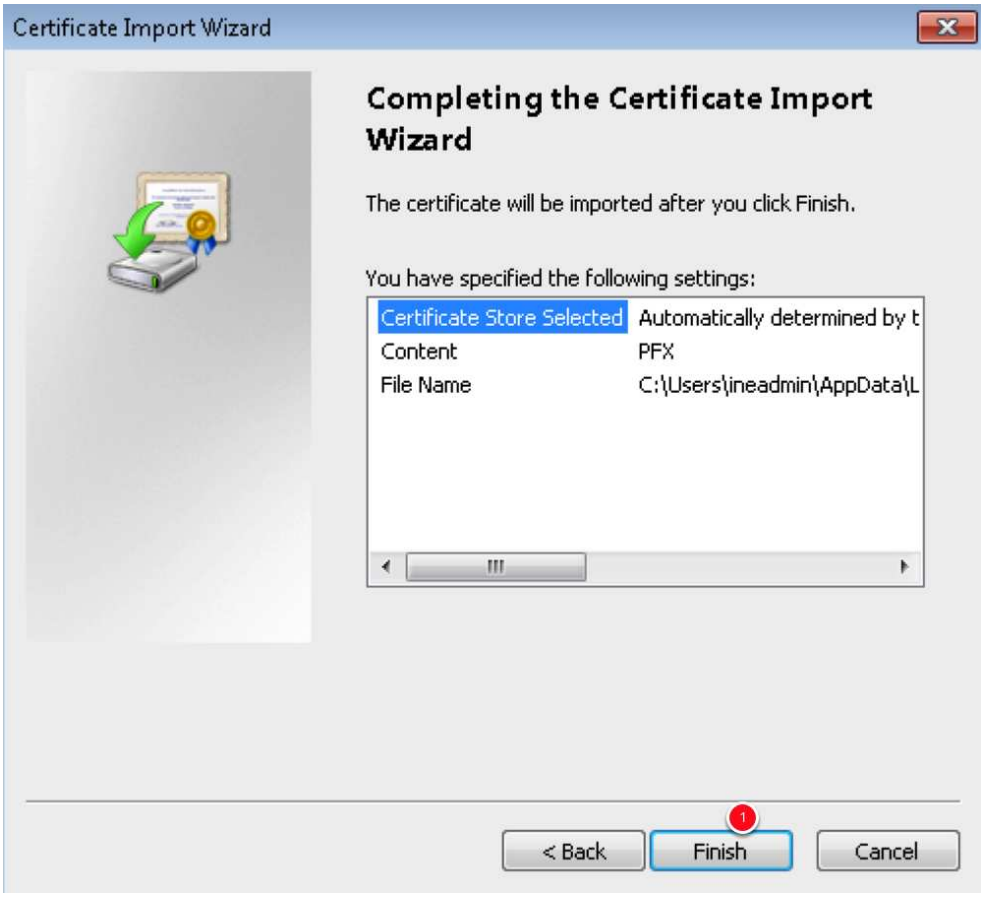
private key to be exportable, make sure to select **Mark this key as exportable**.



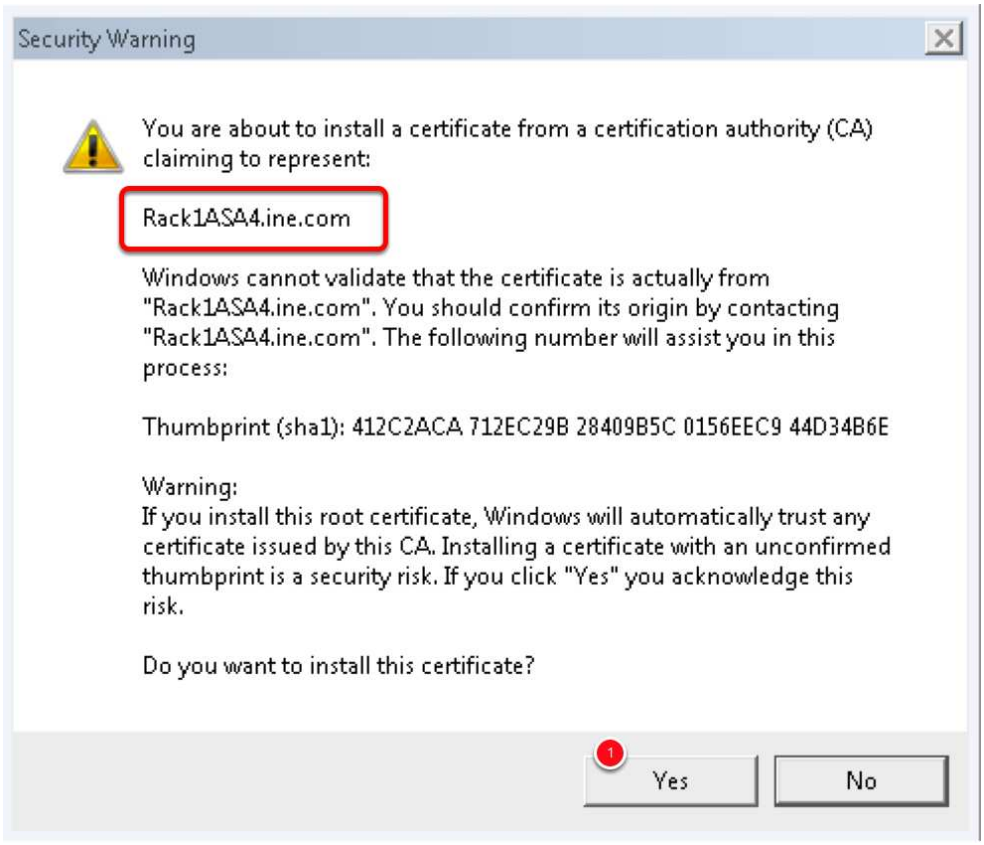
With the automatic selection option, the user certificate will be imported in the Personal User Store and the CA certificate will be imported in the Trusted Root Certificate Authority, which ensures that certificate can be used by the operating system:

https://t.me/learningnets

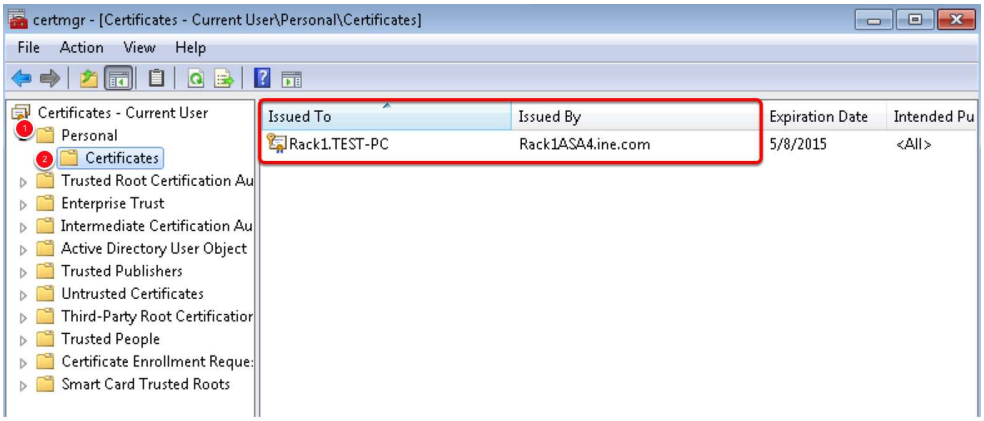




https://t.me/learningnets

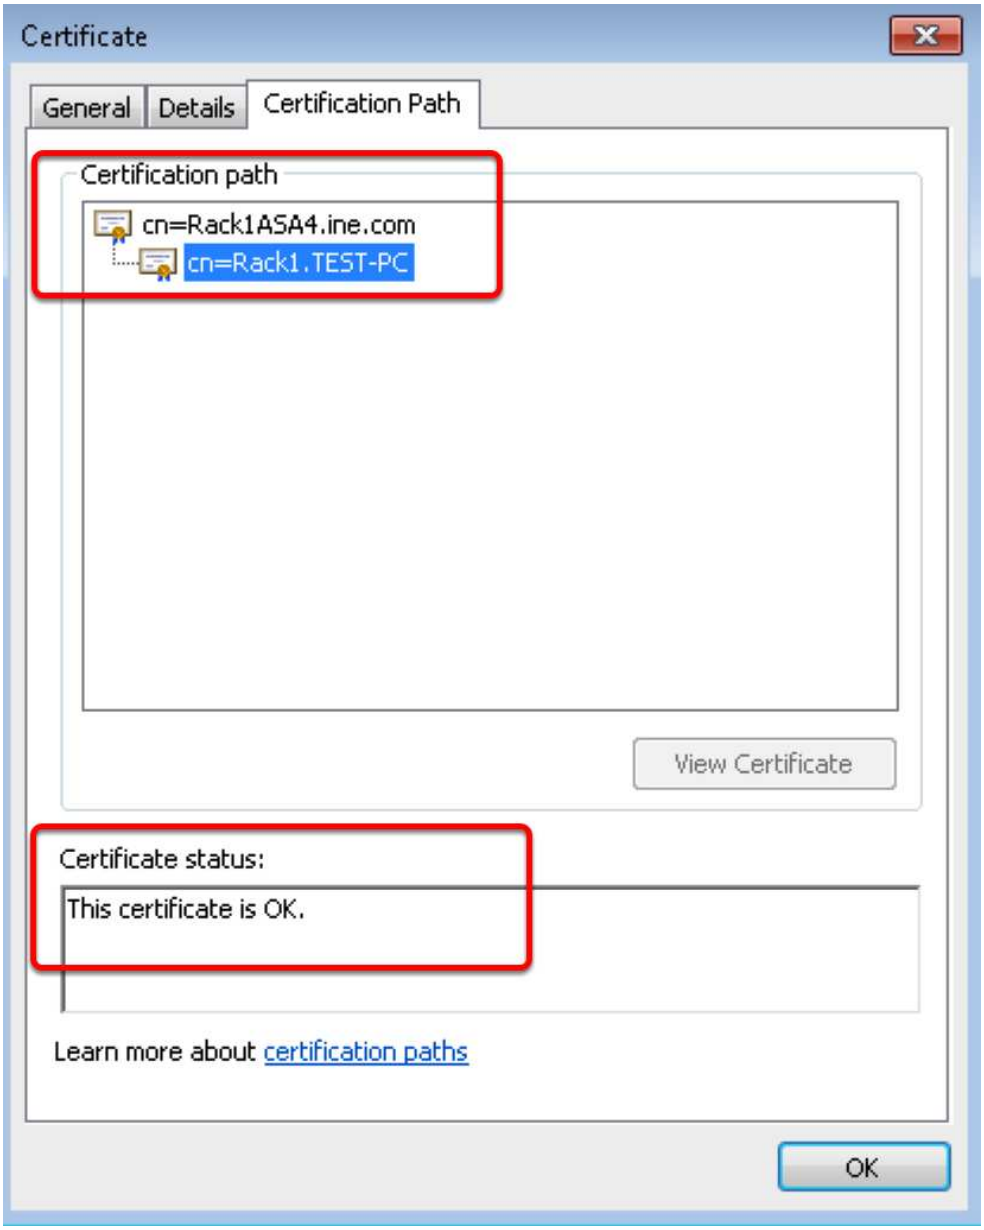


Open the User Certificate Store using the terminal command `certmgr.msc` , and expand the tree under **Personal**, **Certificates** to verify that certificate was installed:



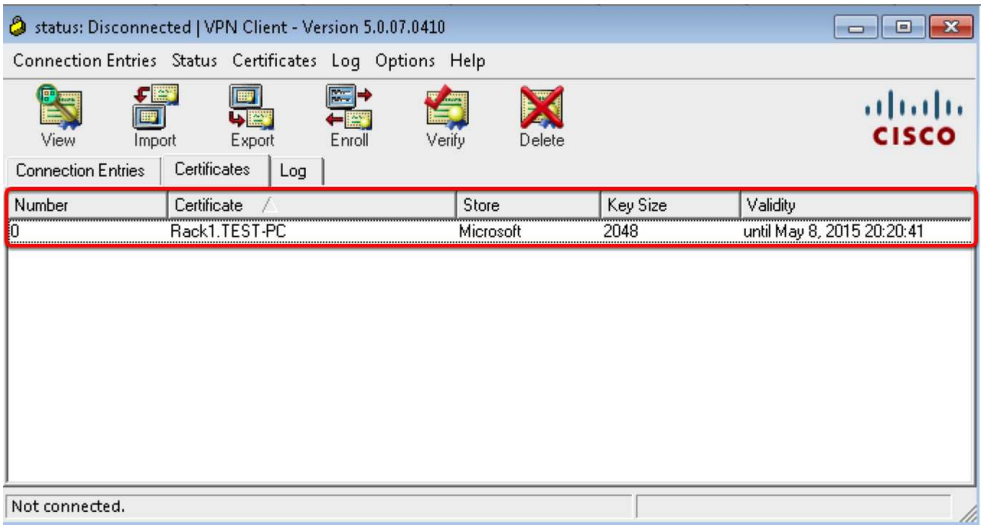
Check that the certificate is valid, so the issuer is also trusted. If the PC's clock is not synchronized with the CA's clock, your certificate may appear as invalid:

https://t.me/learningnets



After the certificate is installed, it will be accessible by both the Cisco VPN Client and the Cisco AnyConnect Secure Mobility Client; verify from the Cisco VPN Client:

https://t.me/learningnets



The PKCS#12 file can also be imported directly in the Cisco VPN Client Store by clicking the **Import** button on the **Certificates** tab of the Cisco VPN Client. With this method, the certificate will not appear in the User Certificate Store, and it will not be accessible by AnyConnect.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

PC Enrollment with IOS PKI

Task

- Configure PKI on R1 so that **TEST-PC-A** can enroll using an IP address of **136.1.18.12** in the URL.
- The certificate should be available only in the Cisco VPN Client Certificate Store.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **PKI Infrastructure** module of this section, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The Cisco VPN Client allows certificate enrollment through SCEP; the Cisco ASA CA does not support SCEP. The IOS CA does, however. The enrollment can be done directly from the VPN Client GUI by clicking the **Enroll** button on the **Certificates** tab. Make sure that you specify the same CA domain name as the one configured on the CA server and use a URL in the following format: **http://[IP Address|Hostname]/cgi-bin/pkiclient.exe**.

Configuration

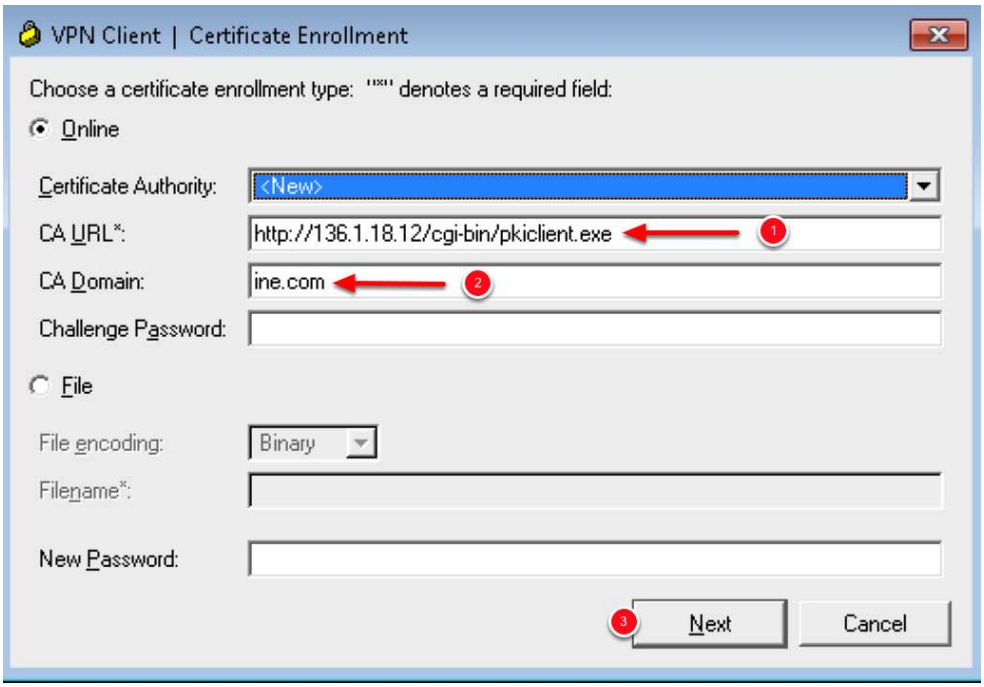
R1:

```
ip http server
!
interface GigabitEthernet0/0
```

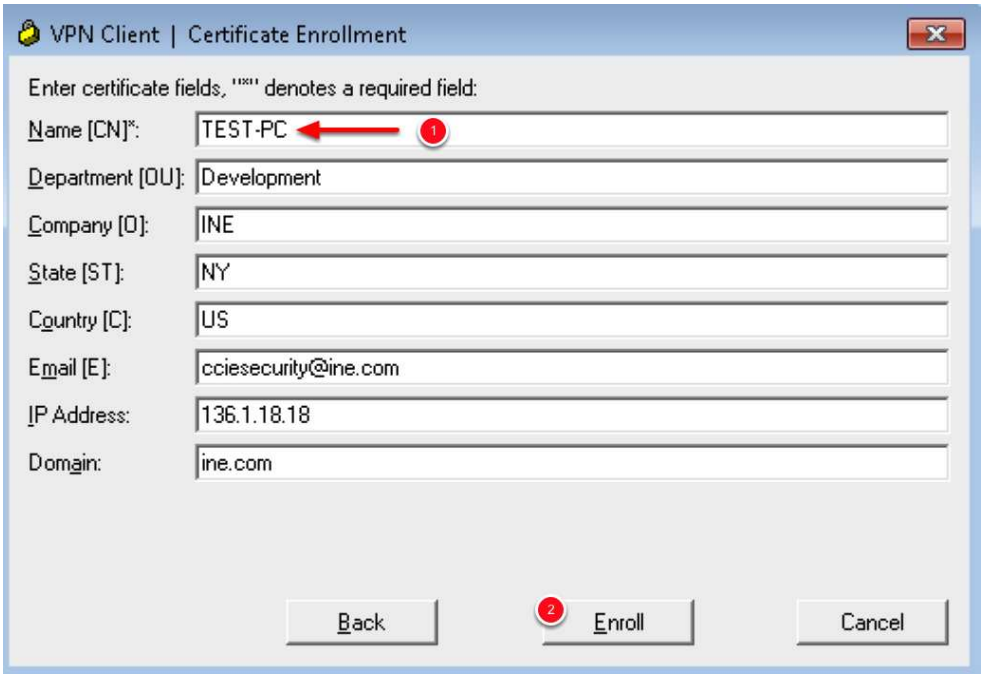
https://t.me/learningnets

```
 standby ip 136.1.18.12
!
crypto pki server HA
 issuer-name CN=ROOT-HA.ine.com
 database archive pkcs12 password ciscocisco
 grant auto
 no shutdown
```

The enrollment process consists of two phases: First specify the CA parameters (URL and domain name), and then specify the certificate attributes where only the CN is mandatory to be completed:

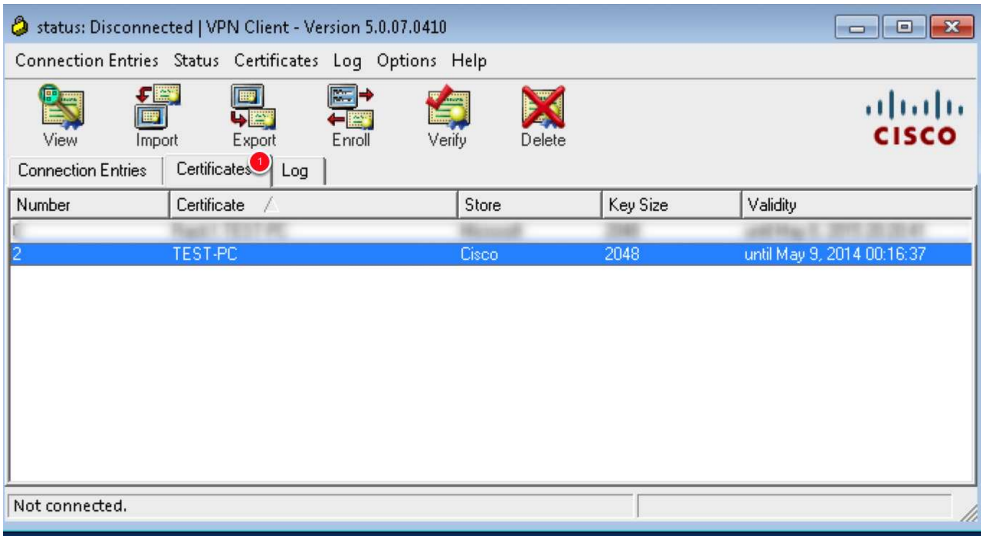


https://t.me/learningnets

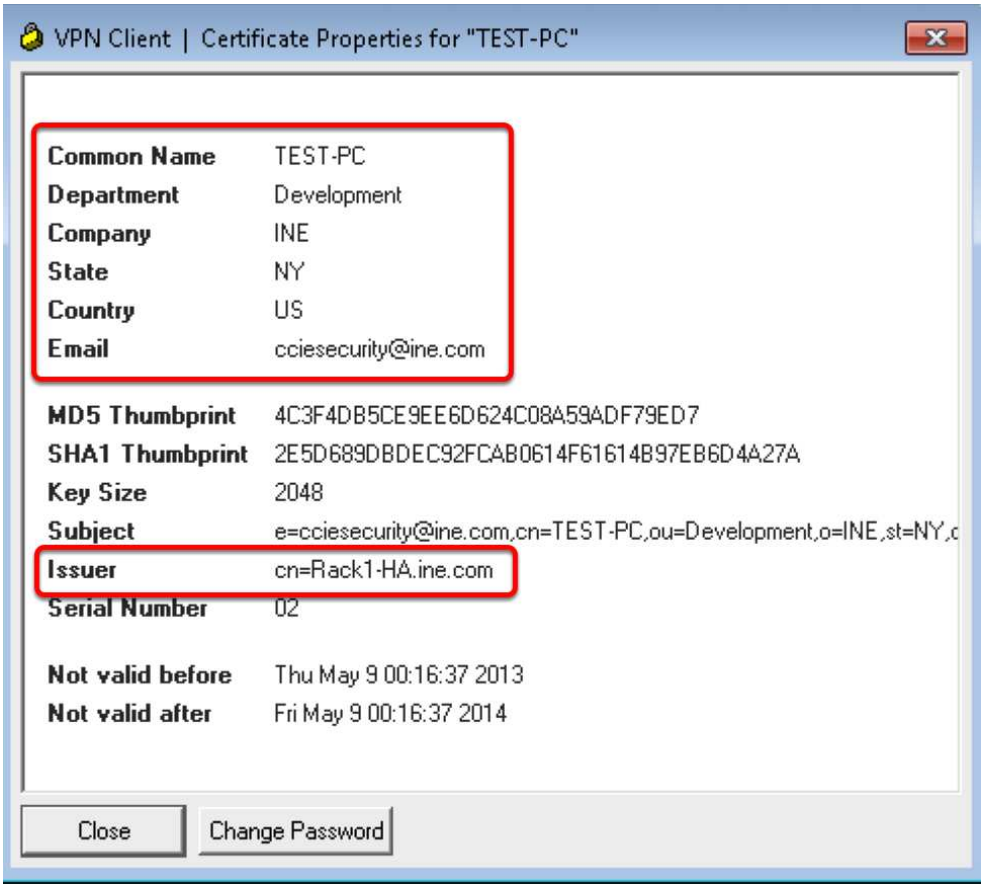


Verification

When the enrollment is completed, you will be able to see the certificate imported in the VPN Client Certificate Store:



https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS Routers with Crypto Map and PSK

Task

- Configure an IKEv1 IPsec tunnel on R1 and R3 to protect traffic between their respective Loopback0 subnets.
- Use the pre-shared key value of **CISCO** for authentication and 3DES/MD5 as cipher/hash for both IPsec phases.
- Ensure that only ICMP traffic is allowed across the tunnel, and log denied traffic.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** or by clicking the **Resources** button.

Overview

A LAN-to-LAN IKEv1 IPsec VPN involves two devices in the security negotiation process. The result of this negotiation is a set of algorithms used to protect (confidentiality, integrity, authentication) certain traffic between the two endpoints. For IKEv1, the negotiation takes place in two phases:

IKEv1 IPsec Phase 1

Devices authenticate each other using the configured method (pre-shared-key or digital certificates), but both parties must first negotiate the authentication method. During the authentication phase, devices exchange their IKEv1 identities (IP addresses, hostnames, or DN from digital certificates) and establish a secure channel called ISAKAMP SA (Security Association), which is used to protect any further management communications. The core component for establishing the secure channel is the Diffie-Hellman (DH) key exchange (KE). This procedure allows a pair of devices to derive a set of common secret shared keys

https://t.me/learningnets

(authentication and encryption) without letting any side party eavesdrop. DH involves discrete calculations on certain cyclic groups. IPsec settings allow you to select the group number, with the larger group being slower in computations but more secure.

You may see the terms ISAKMP (Internet Security Association Key Management Protocol) and IKE (Internet Key Exchange) used interchangeably. However, this is not strictly correct: ISAKMP is an abstract framework, and IKE is its actual implementation. The IKEv1 IPSec Phase 1 may run in two modes:

- **Main Mode** uses a six-message exchange procedure; prior to exchanging device identities and authenticating them, Main Mode ensures that the DH produces a shared encryption key to protect the authentication phase.
- **Aggressive Mode** uses just three messages to establish the ISAKMP SA, and it exchanges device identities in parallel with the shared encryption key generation. This is less secure, but it has some unique advantages when using pre-shared keys for authentication. As we will see later, IKE Main Mode with pre-shared keys has some limitations, because identities are exchanged only after the channel has been secured.

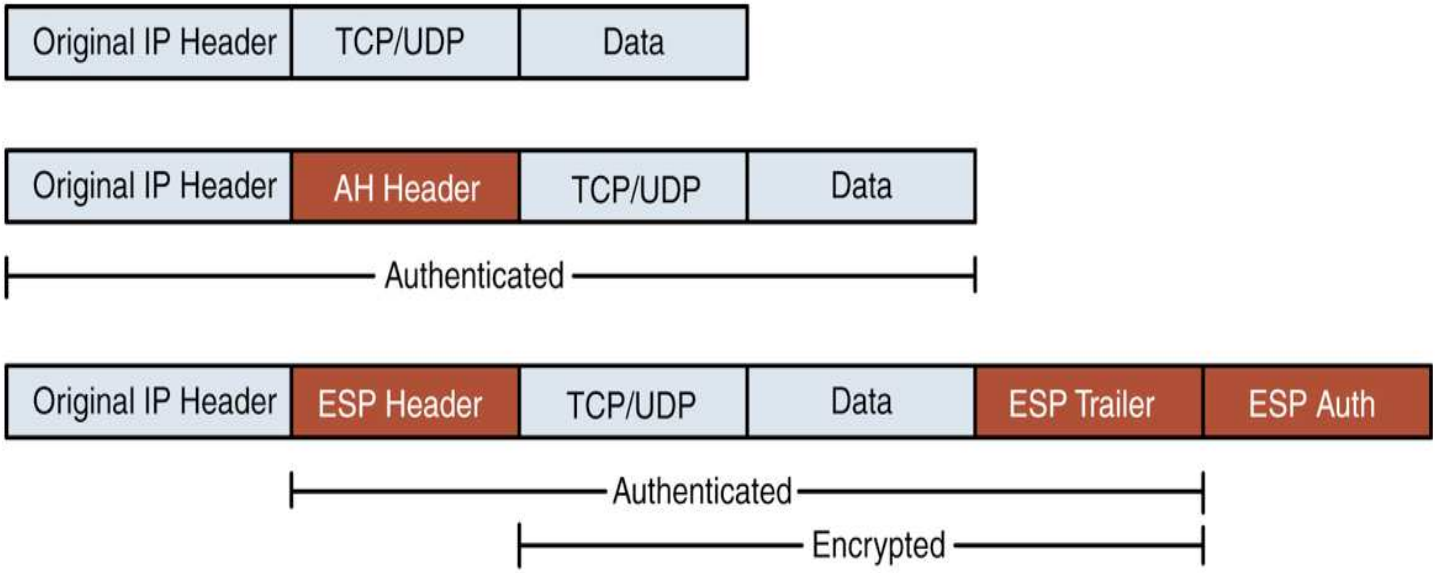
To establish the secure management tunnel, Phase 1 also negotiates encryption and hashing algorithms.

IKEv1 IPsec Phase 2

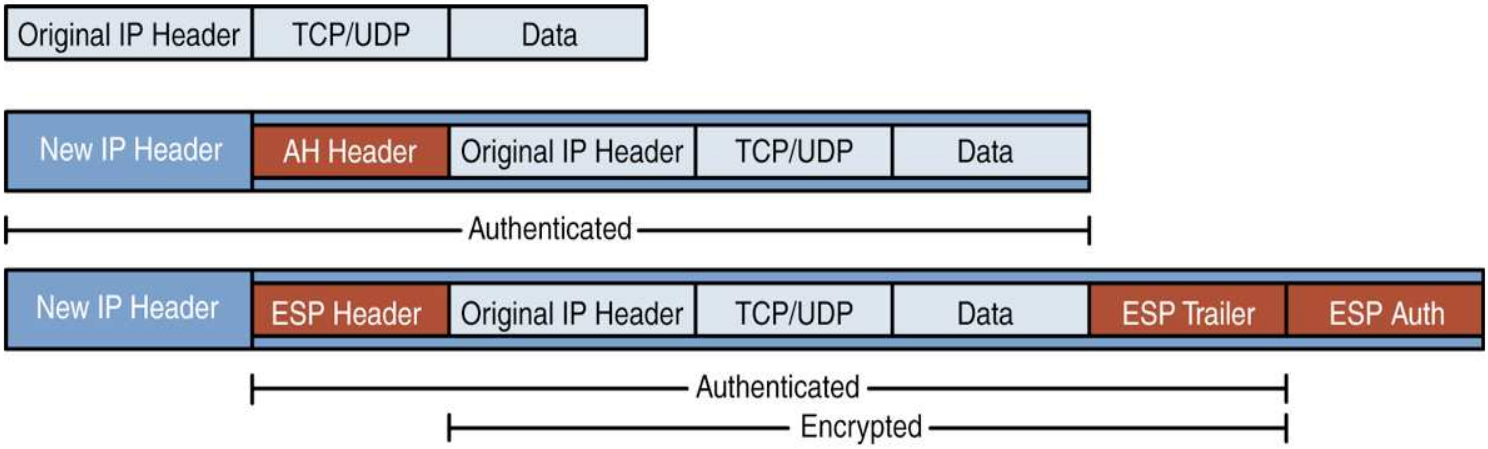
Two endpoints agree on the traffic they are going to encrypt and the cipher/hash functions to use. Both parties exchange the so-called Proxy Identities, which are essentially access-lists defining the traffic that each side wants to encrypt. Both sides check that their Proxy IDs are non-conflicting (they don’t define mismatching traffic sets). The encapsulating protocol (AH or ESP) is also negotiated (AH is not supported by Cisco ASA and does not provide encryption). Except the encapsulating protocol and algorithms used to provide integrity/confidentiality, the two peers also negotiate the encapsulating method, which can be Tunnel or Transport mode; the difference between the two is that Tunnel mode adds an additional IP header as overhead representing the IP addresses of VPN gateways:

Transport mode:

https://t.me/learningnets



Tunnel mode:



IPSec Phase 2 has only one mode of operation, called Quick Mode. This mode uses three messages to establish the IPSec SA. All Quick Mode communications and negotiations are protected by the initial management ISAKMP SA established in Phase 1.

This is the basic scenario for establishing a LAN-to-LAN IPSec tunnel between two IOS routers, using classical IPsec crypto-map configuration. Note that although the control and data plane of IPsec passes through an ASA firewall, from the initial configuration files traffic is allowed to flow bidirectionally unrestricted. From the configuration perspective, the following steps should be used as a guideline:

1. Similar to ASA configuration, define an ISAKMP policy (unless you decide to use the default set of eight ISAKMP policies automatically enabled on the router). Make sure

https://t.me/learningnets

to set the DH group to 2 when connecting to an ASA firewall (or configure the ASA firewall to use DH group 1) because the default DH group for IOS routers is group 1, whereas on the ASA the default is group 2. Other settings must match the settings configured in the remote endpoint, except for the lifetime parameter which is negotiated to the lowest value between the two peers.

2. If you are using pre-shared keys for authentication, you should define one using the global mode command `crypto isakmp key <KEY> [address|hostname] [no-xauth]` or by using a keyring, which would be mapped to an ISAKMP profile. This differs from the tunnel-group settings used in the ASA firewall. In some advanced cases, you may want to set additional Phase 1 settings using ISAKMP profiles. We will cover those in separate tasks.
3. Create a transform set, as you did on the ASA (unless you decide to use the default set of one/two policies, depending on IOS version, automatically enabled on the IOS). Make sure the cipher and the hash match the values used in the remote endpoint.
4. Create an extended access-list that defines the traffic to be encrypted. The locally defined access-list should mirror the access-list entries used at the remote VPN gateway endpoint.
5. Create a crypto map that matches the access-list created above, set the remote peer IP address, and configure the transform set to be applied to the traffic. This completes the configuration of IPSec Phase 2 settings.

ASA firewall can force decrypted IPsecVPN traffic to be checked against the interface or global ACL using the command `no sysopt connection permit-vpn` ; by default it does not. Since version 12.3(8)T, IOS does not match the decrypted traffic against the interface access-list, and this is not configurable. If you want to filter IPsec traffic flowing through the VPN tunnel on IOS, you should apply additional access-lists using the crypto map-level command `set ip access-group <acl-name|acl-number> [in|out]` . This feature could be matched against the ASA `vpn-filter value <acl-name>` command configured under the group-policy for the respective tunnel, but notice that the ASA does not allow you to specify the direction. Based on task requirements, we applied the VPN filter outbound so that only allowed traffic is sent over the tunnel; if it was applied inbound, traffic would be allowed through the tunnel based on the proxy-ACL configuration but it would be filtered inbound at the remote end.

Configuration

The recommended order of operation for IOS L2L IKEv1 IPsec VPN configuration is:

https://t.me/learningnets

1. Configure IPsec Phase1 parameters (IKE policy and PSK).
2. Configure IPsec Phase2 parameters (Proxy-ACL, Filter-ACL and Transform-Set).
3. Bind all configuration into a static crypto-map.
4. Apply the crypto-map on the physical interface where encrypted traffic is sent and received, closest to the remote peer.

Thus, the configuration is:

R1:

```
crypto isakmp policy 100
 authentication pre-share
 hash md5
 encryption 3des
!
!
crypto isakmp key CISCO address 136.1.38.3
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
ip access-list extended LO1_TO_LO3
 permit ip 150.1.1.1 0.0.0.0 150.1.3.3 0.0.0.0
!
ip access-list extended VPN_R1_TO_R3_FILTER
 permit icmp any any
 deny ip any any log
!
crypto map VPN 10 ipsec-isakmp
 match address LO1_TO_LO3
 set transform 3DES_MD5
 set peer 136.1.38.3
 set ip access-group VPN_R1_TO_R3_FILTER out
!
interface GigabitEthernet0/0
 crypto map VPN
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 hash md5
 encryption 3des
!
!
crypto isakmp key CISCO address 136.1.18.1
```

https://t.me/learningnets

```
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
ip access-list extended LO3_TO_LO1
  permit  ip 150.1.3.3 0.0.0.0 150.1.1.1 0.0.0.0
!
ip access-list extended VPN_R3_TO_R1_FILTER
 permit icmp any any
 deny ip any any log
!
crypto map VPN 10 ipsec-isakmp
 match address LO3_TO_LO1
 set transform 3DES_MD5
 set peer 136.1.18.1
 set ip access-group VPN_R3_TO_R1_FILTER out
!
interface FastEthernet0/0
 crypto map VPN
```

Verification

Generate interesting traffic to trigger the tunnel initiation process. The first few packets will not be successful because of the tunnel negotiation process. Verify that non-ICMP traffic is not allowed through the tunnel, and is therefore not functional:

```
R1#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
!
!Rack1R1#telnet 150.1.3.3 /source-interface loopback0
Trying 150.1.3.3 ... % Connection timed out; remote host not responding
```

Verify that the IKEv1 Phase1 and Phase2 is established using the required parameters:

```
R1#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
```

https://t.me/learningnets

```
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1001  136.1.18.1      136.1.38.3 ACTIVE 3des md5
    psk  1  23:55:02
       Engine-id:Conn-id =  SW:1

IPv6 Crypto ISAKMP SA
!
!R1#show crypto ipsec sa peer 136.1.38.3

interface: GigabitEthernet0/0
    Crypto map tag: VPN, local addr 136.1.18.1

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
   current_peer 136.1.38.3 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.1, remote crypto endpt.: 136.1.38.3
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0
     current outbound spi: 0xCA838380(3397616512)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0xBF9BC979(3214657913)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2001, flow_id: Onboard VPN:1, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4356139/3276)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
 outbound esp sas:
      spi: 0xCA838380(3397616512)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2002, flow_id: Onboard VPN:2, sibling_flags 80000040, crypto map: VPN
```

https://t.me/learningnets

```
        sa timing: remaining key lifetime (k/sec): (4356139/3276)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:


     outbound pcp sas:
```

IKEv1 Phase1 and Phase2 configuration can be verified using the following commands:

```
R1#show crypto isakmp policy
 Global IKE policy
Protection suite of priority 100 encryption algorithm:   Three key triple DES
 hash algorithm:          Message Digest 5
 authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #1 (768 bit)
        lifetime:               86400 seconds, no volume limit
!
!R1#show crypto map tag VPN
Crypto Map IPv4 "VPN" 10 ipsec-isakmp Peer = 136.1.38.3
        Extended IP access list LO1_TO_LO3
 access-list LO1_TO_LO3 permit ip host 150.1.1.1 host 150.1.3.3
        Extended IP access check OUT list VPN_R1_TO_R3_FILTER
 access-list VPN_R1_TO_R3_FILTER permit icmp any any
 access-list VPN_R1_TO_R3_FILTER deny ip any any
        Current peer: 136.1.38.3
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={  3DES_MD5:  { esp-3des esp-md5-hmac  } ,


        }
        Interfaces using crypto map VPN:
                GigabitEthernet0/0
```

The following is a very useful command for visualizing the tunnel status, as well as negotiated Proxy-ACL and traffic statistics:

```
R1#show crypto session remote 136.1.38.3 detail
Crypto session current status


Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
```

https://t.me/learningnets

```
X - IKE Extended Authentication, F - IKE Fragmentation

Interface: GigabitEthernet0/0
Uptime: 00:09:32 Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500 fvrf: (none) ivrf: (none)
      Phase1_id: 136.1.38.3
      Desc: (none)
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
         Capabilities:(none) connid:1001 lifetime:23:50:27
 IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3
  Active SAs: 2, origin: crypto map

        Inbound:  #pkts dec'ed 4 drop 0 life (KB/Sec) 4356139/3027
        Outbound: #pkts enc'ed 4 drop 0 life (KB/Sec) 4356139/3027
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS Routers with Crypto Map and PKI

Task

- Configure an IKEv1 IPsec tunnel on R1 and R3 to protect traffic between their respective Loopback0 subnets, use digital certificate authentication.
- Consider R1 as the CA server, it is pre-configured from the loaded initial configurations.
- Ensure that R1 allows the tunnel to be established only if R3 presents a certificate with a CN value of **R3**.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The certificate-based authentication can also be called digital-signature based or RSA-signature based, because RSA is the most common underlying public-key encryption algorithm (remember from the PKI tasks that EC signatures are not of interest for the current CCIE Security lab). To configure the authentication based on digital signatures, perform the following steps:

1. Generate RSA private/public keys on every router or firewall that should participate in the VPN. This is necessary to create a key-pair that uniquely identifies the router. Notice that IOS routers and ASA firewall need a domain-name configured before

https://t.me/learningnets

generating the RSA keys.

2. Configure the CA trustpoint, which is required to retrieve the CA certificate and obtain a personal digital certificate. You define a trustpoint by using the command `crypto ca trustpoint <name>` on the ASA firewall and `crypto pki trustpoint <name>` on the IOS routers. Cisco devices use HTTP-based SCEP protocol to interact with CA, so you must define the URL to access the CA (IOS routers can also use non-SCEP HTTP enrollment through the use of enrollment profiles defined with the command `crypto pki enrollment profile <name>`). The command to define a URL is `enrollment url` under the CA trustpoint configuration mode. CAs usually define and store CRL's (Certificate Revocation List) URL to allow the endpoints to poll the CA for the list of invalidated (revoked) certificates.
3. After you define the required trustpoint settings, you must retrieve the CA certificate to validate VPN peers certificates based on the CA's digital signature. Use the command `crypto ca authenticate <name>` on ASA firewalls or `crypto pki authenticate <name>` on IOS routers for CA certificate retrieval and authentication. Before you authenticate the CA trustpoint, make sure the time is coordinated between the CA server and CA client. Most often, you will want to use NTP protocol to accomplish this.
4. After you have obtained the CA certificate, you must request a personal digital certificate from the CA using the command `crypto ca enroll <name>` on ASA firewalls or `crypto pki enroll <name>` on IOS routers. When you issue one of these commands, the router or the firewall will send a CSR request to the CA using SCEP protocol. The request contains the public key of the router and identity information, such as the router's hostname and domain name, signed with its private key. Depending on the CA configuration, the certificate request will remain in pending state until the administrator approves it, or the certificate could be issued automatically. In our topology, the CA is configured for automatic certificate granting. You may want to use the command `debug crypto pki transactions` on the client side or `debug crypto pki server` on server side if you are experiencing any problems enrolling with the CA.
5. After you have obtained the certificate, you may configure ISAKMP policy for authentication based on RSA signatures.
6. All other IPSec settings remain the same. Only the ISAKMP (Phase 1) is affected by certificate configuration.

From the IPsec configuration perspective, we need only modify the authentication method from PSK to PKI if compared with previous task. Note that although R1 has a self-signed certificate from its CA role, this certificate is used solely for signing client certificates and cannot be used for other purposes (the scope of the certificate is available in the **Certificate Usage** certificate field). For this reason, we must

https://t.me/learningnets

enroll both R1 and R3 with the CA and receive a default **General Purpose** certificate that can be used for IPsec. The trustpoint used by the CA function on R1 is locked (cannot be accessed or edited) as long as the CA role is active.

By default, IOS routers use the source IP address of the IKEv1 packets as an IKEv1 identity (the IP address of IKEv1 identity can be changed using the command `crypto-map <map-name> local-address <interface>`); it can be changed to hostname or DN using the command `crypto isakmp identity [hostname|dn]` , and DN should be used only when using certificate authentication. For R1 to validate the DN value or R3's certificate, we must:

- Configure R3 to use a DN IKEv1 identity.
- Configure R1 to match on the DN value in the crypto-map using crypto identities.

Crypto identities can be configured to match on either DN or FQDN value from the digital certificate used in the IKEv1 authentication process.

Configuration

Although not required by the task, R1 was also configured with a CN value in its certificate. Extra commands for disallowing serial-number and IP address to be present in the certificate were configured to ease the enrollment process; otherwise, IOS will ask for it.

R1:

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1-CLIENT
 subject-name CN=R1
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1-CLIENT
crypto pki enroll R1-CLIENT
!
crypto isakmp policy 100
 authentication rsa-sig
 hash md5
 encryption 3des
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
```

https://t.me/learningnets

```
ip access-list extended LO1_TO_LO3
 permit ip 150.1.1.1 0.0.0.0 150.1.3.3 0.0.0.0
!
crypto identity R1_R3_DN
 dn CN=R3
!
crypto map VPN 10 ipsec-isakmp
 match address LO1_TO_LO3
 set transform 3DES_MD5
 set peer 136.1.38.3
 set identity R1_R3_DN
!
interface GigabitEthernet0/0
 crypto map VPN
```

R3:

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1
 subject-name CN=R3
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1
crypto pki enroll R1
!
crypto isakmp identity dn
crypto isakmp policy 100
 authentication rsa-sig
 hash md5
 encryption 3des
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
ip access-list extended LO3_TO_LO1
  permit  ip 150.1.3.3 0.0.0.0 150.1.1.1 0.0.0.0
!
crypto map VPN 10 ipsec-isakmp
 match address LO3_TO_LO1
 set transform 3DES_MD5
 set peer 136.1.18.1
!
```

https://t.me/learningnets

```
interface FastEthernet0/0
 crypto map VPN
```

Verification

If continuing from the previous task, make sure to bring down the tunnel that used PSK authentication. Clear both IKEv1 phases on both routers using the commands **`clear crypto isakmp`** and **`clear crypto sa .`** Generate traffic to trigger the IPsec process:

```
R1#ping 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 .!!!!


Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
```

Verify that the tunnel was established using certificate authentication:

https://t.me/learningnets

```
R1#show crypto isakmp sa detail

Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
 1003  136.1.18.1      136.1.38.3 ACTIVE 3des md5    rsig
 1  23:51:14
       Engine-id:Conn-id =  SW:3

IPv6 Crypto ISAKMP SA
!
!R1#show crypto session remote 136.1.38.3

Crypto session current status

Interface: GigabitEthernet0/0 Session status: UP-ACTIVE

Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3  Active SAs: 2, origin: crypto map
```

Verify that the CA self-signed certificate has a limited usage scope:

https://t.me/learningnets

```
R1#show crypto pki certificates R1
CA Certificate
 Status: Available

  Certificate Serial Number (hex): 01 Certificate Usage: Signature

  Issuer:
    cn=R1.ine.com
  Subject:
    cn=R1.ine.com
  Validity Date:
    start date: 05:56:01 UTC May 2 2013
    end   date: 05:56:01 UTC May 1 2016
  Associated Trustpoints: R1-CLIENT R1
  Storage: nvram:Rack1R1ineco#1CA.cer
```

Verify that the required IPsec certificates have been issued to both R1 and R3:

```
R1#show crypto pki certificates R1-CLIENT
Certificate
 Status: Available
  Certificate Serial Number (hex): 03 Certificate Usage: General Purpose
  Issuer:
    cn=R1.ine.com
  Subject:
    Name: R1.ine.com
    hostname=R1.ine.com cn=R1
  Validity Date:
    start date: 09:08:57 UTC May 2 2013
    end   date: 09:08:57 UTC May 2 2014
  Associated Trustpoints: R1-CLIENT

CA Certificate
  Status: Available
  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
  Issuer:
    cn=R1.ine.com
  Subject:
    cn=R1.ine.com
  Validity Date:
    start date: 05:56:01 UTC May 2 2013
    end   date: 05:56:01 UTC May 1 2016
  Associated Trustpoints: R1-CLIENT R1
```

https://t.me/learningnets

```
  Storage: nvram:Rack1R1ineco#1CA.cer
!
!R3#show crypto pki certificates R1
Certificate
 Status: Available
  Certificate Serial Number (hex): 02 Certificate Usage: General Purpose
  Issuer:
    cn=R1.ine.com
  Subject:
    Name: R3.ine.com
    hostname=R3.ine.com cn=R3

  Validity Date:
    start date: 08:59:45 UTC May 2 2013
    end   date: 08:59:45 UTC May 2 2014
  Associated Trustpoints: R1

CA Certificate
  Status: Available
  Certificate Serial Number (hex): 01
  Certificate Usage: Signature
  Issuer:
    cn=R1.ine.com
  Subject:
    cn=R1.ine.com
  Validity Date:
    start date: 05:56:01 UTC May 2 2013
    end   date: 05:56:01 UTC May 1 2016
  Associated Trustpoints: R1
```

Verify that packets were both encrypted and decrypted:

```
R1#show crypto session remote 136.1.38.3 detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation

Interface: GigabitEthernet0/0
Uptime: 00:01:40 Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500 fvrf: (none) ivrf: (none)
      Phase1_id: hostname=R3.ine.com,cn=R3
      Desc: (none)
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
```

https://t.me/learningnets

```
          Capabilities:(none) connid:1002 lifetime:23:58:18
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3
        Active SAs: 2, origin: crypto map Inbound:  #pkts dec'ed 4 drop 0 life (KB/Sec) 4275160/3499
Outbound: #pkts enc'ed 4 drop 0 life (KB/Sec) 4275160/3499
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS Routers with ISAKMP Defaults

Task

- Configure an IKEv1 IPsec tunnel on R1 and R2 to protect traffic between their respective Loopback0 subnets.
- Use the pre-shared key value **CISCO** for authentication and AES-128/SHA1 as cipher/hash for both IPsec phases.
- Adhere to the following restrictions:
  - Use the minimum necessary commands on both routers.
  - Do not use a `set peer` command on R1.
- Configure an IKEv1 ISAKMP policy on R1 using 3DES/MD5/DH Group 2 and pre-shared-key authentication.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is a basic crypto-map-based L2L IPsec VPN, but we need to understand the restrictions. On R1, we are not allowed to use a `set peer` command, so we cannot use static crypto-maps because to be active it requires both `set peer` and `match address` commands. The solution is to use dynamic crypto-maps. Note the following:

- Dynamic crypto-maps are used for cases when the IP address of the remote end is

https://t.me/learningnets

unknown or dynamically allocated (both remote-access and LAN-to-LAN).

- Dynamic crypto-maps can also be used for cases when the proxy-ACL (traffic allowed through the tunnel) is not defined/known locally (but it is on the remote-end) and is therefore negotiated.
- If separate dynamic crypto-map entries are used for remote-access and LAN-to-LAN VPN types, it is required to use a `match address` command for the LAN-to-LAN entry and a lower dynamic crypto-map sequence number for the LAN-to-LAN entry; otherwise, you will not be able to differentiate between the two entries and negotiation could also fail depending on the configuration, or all incoming IKEv1 sessions will match on the first entry in the dynamic crypto-map.
- Entries in static crypto-maps are processed top-down based on their sequence numbers, so the entry referencing a dynamic crypto-map should have the highest sequence number; otherwise, all incoming IKEv1 sessions could match on the dynamic crypto-map (static entries will not be processed).

The task requires that we use the minimum necessary commands. Starting with IOS 12.4(20)T, IOS uses a default of eight ISAKMP policies and one or two (depending on the exact version) transform-sets, which can be visualized with the commands `show crypto isakmp default policy` and `show crypto ipsec transform-set .` These can be disabled with the commands `no crypto isakmp default policy` and `no crypto ipsec default transform-set .` We need to determine the situations in which the router uses them (considering that they are not manually disabled):

- Default ISAKMP policies are used only if no manually defined ISAKMP policies exist.
- Default transform-sets are used by all crypto-map entries and IPsec profiles for which no transform-sets have been applied using the command `set transform-set <name> .`

Because R1 has non-default ISAKMP policy defined, the default ISKAMP policies can no longer be used, so an ISAKMP policy with AES-128/SHA1 to match the default from R2 is configured. On R2, we use a static crypto-map, so the proxy-ACL and peer must be defined. On both routers, we used the default transform-set.

Configuration

R1:

```
crypto isakmp policy 10
 authentication pre-share
 encryption aes
 hash sha
 group 5
!
```

https://t.me/learningnets

```
crypto isakmp policy 20
 encr 3des
 hash md5
 authentication pre-share
 group 2
!
crypto isakmp key CISCO address 136.1.18.2
!
crypto dynamic-map R1_TO_R2 100
crypto map VPN 100 ipsec-isakmp dynamic R1_TO_R2
!
interface GigabitEthernet0/0
 crypto map VPN
```

R2:

```
crypto isakmp key CISCO address 136.1.18.1
!
ip access-list extended LO2_TO_LO1
  permit ip 150.1.2.2 0.0.0.0 150.1.1.1 0.0.0.0
!
crypto map VPN 10 ipsec-isakmp
 match address LO2_TO_LO1
 set peer 136.1.18.1
!
interface GigabitEthernet0/0
 crypto map VPN
```

Verification

When IKEv1 is initiated using default policies, the router will log a message similar to **`%CRYPTO-6-IKMP_POLICY_DEFAULT: Using ISAKMP Default policies`** . Generate interesting traffic to initiate the negotiations and verify that the tunnel is UP using required algorithms. Because R1 does not have a proxy-ACL configured and uses a dynamic crypto-map, interesting traffic must be initiated from R2:

```
R2#ping 150.1.1.1 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 150.1.2.2 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
!
!R2#show crypto session  remote 136.1.18.1
```

https://t.me/learningnets

```
Crypto session current status

Interface: GigabitEthernet0/0 Session status: UP-ACTIVE
Peer: 136.1.18.1 port 500
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.18.1/500 Active
  IPSEC FLOW: permit ip host 150.1.2.2 host 150.1.1.1  Active SAs: 2, origin: crypto map
!
!R2#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1001  136.1.18.2      136.1.18.1 ACTIVE aes  sha    psk  5
  23:56:07
       Engine-id:Conn-id =  SW:1

IPv6 Crypto ISAKMP SA
```

Verify that the default transform-set is used by the ACTIVE SAs. Note that although it is configured in **`Transport mode`** , this is negotiated and actually Tunnel mode is used; transport mode should be used when the router acts as both the VPN gateway and VPN traffic source/destination, such as for GRE over IPSec.

```
R2#show crypto ipsec transform-set
Transform set default: { esp-aes esp-sha-hmac  }
   will negotiate = { Transport,  },
!
!R2#show crypto ipsec sa peer 136.1.18.1

interface: GigabitEthernet0/0
    Crypto map tag: VPN, local addr 136.1.18.2

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (150.1.2.2/255.255.255.255/0/0)
   remote ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   current_peer 136.1.18.1 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
 #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
```

https://t.me/learningnets

```
    #send errors 0, #recv errors 0

  local crypto endpt.: 136.1.18.2, remote crypto endpt.: 136.1.18.1
  path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0
  current outbound spi: 0x7FA44BE0(2141473760)
  PFS (Y/N): N, DH group: none
 inbound esp sas:
       spi: 0xD83E133B(3627946811) transform: esp-aes esp-sha-hmac ,
 in use settings ={Tunnel, }
         conn id: 2001, flow_id: Onboard VPN:1, sibling_flags 80000040, crypto map: VPN
         sa timing: remaining key lifetime (k/sec): (4373645/3308)
         IV size: 16 bytes
         replay detection support: Y Status: ACTIVE(ACTIVE)

      inbound ah sas:

      inbound pcp sas:
 outbound esp sas:
       spi: 0x7FA44BE0(2141473760) transform: esp-aes esp-sha-hmac ,
 in use settings ={Tunnel, }
         conn id: 2002, flow_id: Onboard VPN:2, sibling_flags 80000040, crypto map: VPN
         sa timing: remaining key lifetime (k/sec): (4373645/3308)
         IV size: 16 bytes
         replay detection support: Y Status: ACTIVE(ACTIVE)


      outbound ah sas:

      outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS Routers with ISAKMP Profile

Task

- Configure an IKEv1 IPsec tunnel on R1 and R2 to protect traffic between their respective Loopback0 subnets:
  - Use R1-issued certificates for authentication; the CA role on R1 is pre-configured from initial configurations.
  - R1 should establish the VPN tunnel only if the remote peer's domain equals **ine.com**.
  - R2 should establish the VPN tunnel only if the peer's certificate issuer equals **R1.ine.com**.
  - Do not use any `crypto identity` -related commands.
- Make use of the default crypto policies and do not use a `set peer` command on R1.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Before an IKEv1 IPsec VPN data-plane is functional, the control-plane passes through two or three phases: Phase1, Phase1.5, and Phase2. The problem is that there is no logical connection or link in the router configuration between Phase1/Phase1.5 and Phase2; basically, the router parses the configuration and makes independent decisions for each phase. Although this does not seem to be a

https://t.me/learningnets

problem in most VPN deployments, on a router configured with multiple VPN types, where each requires a different set of rules, it may become challenging. This is the reason why ISAKMP policies exist. You configure some conditions for Phase1/Phase1.5, and if these are matched, the IKEv1 connection is associated to a particular Phase2. For the tunnel to be established, Phase2 conditions must also be matched. Some common scenarios in which ISAKMP profiles are required or recommended include:

- VRF-aware IPsec
- Per-IPsec tunnel QoS settings
- Per-IPsec tunnel Phase1 settings
- Certificate-to-IPsec tunnel mappings

Some common parameters that can be configured under the ISAKMP profile include the following:

- `client authentication list <AAA_LIST>` enables XAUTH using the `AAA_LIST` login method.
- `client configuration group <EZVPN_GROUP>` specifies the EzVPN configuration group for a remote-access VPN; the group is used to retrieve Phase 1.5 attributes.
- `client configuration address [respond|initiate]` determines the functional mode for assigning an IP address to a remote peer for a remote-access VPN.
- `isakmp authorization list <AAA_LIST>` specifies a network authorization list to be used for retrieving authorization attributes for the session (for example, DNS servers and domain name).
- `virtual-template <number>` specifies the virtual-template interface of type **tunnel** to be used for remote access VPN connections; this is detailed in the DVTI task.
- `initiate mode [aggressive]` allows the router to start IKEv1 Aggressive Mode instead of Main Mode but only if it's the VPN session initiator. All other settings configured in ISAKMP profiles apply for when the router is the VPN session responder. To activate the initiator mode, you must do either of the following:
  - Associate the ISAKMP profile with a crypto-map using the command `crypto map <NAME> isakmp-profile <PROFILE_NAME>` or `set isakmp-profile <PROFILE_NAME>` at crypto-map level. Every connection attempted using this crypto-map will use the respective profile's settings.
  - Associate the ISAKMP profile with an IPSec profile bound to a VTI using the IPSec profile command `set isakmp-profile <PROFILE_NAME> .`
- `self-identity [address|fqdn|user-fqdn]` allows you to modify the IKEv1 identity per IPsec tunnel, regardless of the global configuration.
- `local-address` is identical in functionality to `crypto map <name> local-address` and

https://t.me/learningnets

specifies which IP address IKEv1 should use for the session; by default it is the IP address of the interface where crypto map is applied or the tunnel source if VTI/GRE is used.

- `keyring [default|<name>]` is useful in the case of pre-shared-key authentication and allows you to make use of one or more keys defined in a keyring; with the use of keyring and ISAKMP profiles, the command `crypto isakmp key <string> address <IP> <mask> [no-xauth]` becomes obsolete.

When first enabling an ISAKMP profile, the router will log a message similar to `% A profile is deemed incomplete until it has match identity statements` , which means that for the ISAKMP profile to be active and usable by the router, the remote peer IKE identity must be configured. Imagine that you have two LAN-to-LAN and one remote-access IPSec VPNs configured; if ISKAMP profiles are used, you must ensure that each of the three incoming sessions maps to a unique IKSAMP profile. This can be achieved by matching on the IKEv1 ID of the remote peer or matching certain fields of the peer's certificate by using the commands:

- `match identity [address|group|host|user-fqdn]` : The `group` option is used for remote-access VPNs and refers to group name if pre-shared-keys are used for authentication, or the OU field if certificates are used for authentication.
- `match certificate <certificate-map>` : A certificate map is configured separately using the command `crypto pki certificate map <name> <seq-nr>` , and one or more certificate attributes can be matched, such as subject-name, alt-subject-name, issuer-name, or serial-number.

To bind the conditions configured in an ISAKMP profile to the Phase2, the profile must be applied at the crypto-map or IPsec profile level. The most important thing to remember is that ISAKMP profiles are not dependent on the type of VPN connection configured (DMVPN, GETVPN, GRE-over-IPsec, Crypto-Map, VTI). A certificate map, also known as certificate access-list, can also be applied at the trustpoint level, creating restrictions on certificates issued by a certain CA. For example, if your router is enrolled with a CA and you want to accept peer certificates issued by this CA only if the valid start date is greater than or equal to 12-May-2013, the configuration would be:

https://t.me/learningnets

```
crypto pki certificate map MYMAP 10
 valid-start ge 12 MAY 2013
!
crypto pki trustpoint MYCA
 match certificate MYMAP
```

Configuration

R1 requires an ISAKMP profile to match on the domain-name from the IKE identity presented by R2. R2's identity must be modified from the default of IP address to hostname, so that R1 can match on its domain-name; this can be globally modified with the command **`crypto isakmp identity hostname`** or only for this specific VPN tunnel using the ISAKMP profile-level command **`self-identity fqdn`** . R2 requires an ISKAMP profile to match on the issuer of R1's certificate.

R1:

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1-CLIENT
 subject-name CN=R1
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1-CLIENT
crypto pki enroll R1-CLIENT
!
crypto isakmp profile R1_R2
 match identity host domain ine.com
!
crypto dynamic-map R1_TO_R2 100
 set isakmp-profile R1_R2
!
crypto map VPN 100 ipsec-isakmp dynamic R1_TO_R2
!
interface GigabitEthernet0/0
 crypto map VPN
```

R2:

https://t.me/learningnets

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1
 subject-name CN=R2
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1
crypto pki enroll R1
!
crypto isakmp identity hostname
!
crypto pki certificate map R2_R1 10
 issuer-name eq R1.ine.com
!
crypto isakmp profile R2_R1
 match certificate R2_R1
!
ip access-list extended LO2_TO_LO1
 permit ip 150.1.2.2 0.0.0.0 150.1.1.1 0.0.0.0
!
crypto map VPN 10 ipsec-isakmp
 match address LO2_TO_LO1
 set peer 136.1.18.1
 set isakmp-profile R2_R1
!
interface GigabitEthernet0/0
 crypto map VPN
```

Verification

If you have completed the previous task, bring down the IPsec tunnel between R1 and R2 to ensure renegotiation by using the `clear crypto sa` and `clear crypto isakmp` commands. Generate interesting traffic from R2:

```
R2#ping 150.1.1.1 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 150.1.2.2
%CRYPTO-6-IKMP_POLICY_DEFAULT: Using ISAKMP Default policies
.!!!!
```

https://t.me/learningnets

```
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Verify that certificate authentication is used and that the tunnel is established by using dynamic crypto map:

```
R1#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1016  136.1.18.1      136.1.18.2 ACTIVE aes  sha    rsig 5
  23:05:15
      Engine-id:Conn-id =  SW:16

IPv6 Crypto ISAKMP SA
!
!R1#show crypto session  remote 136.1.18.2
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R1_R2 Session status: UP-ACTIVE
Peer: 136.1.18.2 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.18.2/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.2.2  Active SAs: 2, origin: dynamic crypto map
```

Verify that ISAKMP profiles are matched on both routers by viewing the **Ref Count** number:

```
R2#show crypto isakmp profile

IKEv1 PROFILE R2_R1 Ref Count = 4
   Identities matched are:
   Certificate maps matched are:
      R1
   keyring(s): <none>
   trustpoint(s): <all>
!
!R1#show crypto isakmp profile

IKEv1 PROFILE R1_R2 Ref Count = 4
```

https://t.me/learningnets

```
   Identities matched are:
    domain ine.com
   Certificate maps matched are:
   keyring(s): <none>
   trustpoint(s): <all>
```

Verify that Phase2 actually matched the ISAKMP profile and check encrypted/decrypted packets:

```
R1#show crypto session isakmp profile R1_R2 detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation

Interface: GigabitEthernet0/0 Profile: R1_R2
Uptime: 00:02:28 Session status: UP-ACTIVE
Peer: 136.1.18.2 port 500 fvrf: (none) ivrf: (none)
      Phase1_id: R2.ine.com
      Desc: (none)
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.18.2/500 Active
          Capabilities:(none) connid:1008 lifetime:23:57:30
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.2.2
        Active SAs: 2, origin: dynamic crypto map
 Inbound:  #pkts dec'ed 4 drop 0 life (KB/Sec) 4251304/3451
 Outbound: #pkts enc'ed 4 drop 0 life (KB/Sec) 4251304/3451
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS Routers with SVTI and PSK

Task

- Configure an IPsec tunnel between R2 and R3 without using `crypto map` commands.
- Use a pre-shared-key of **CISCO** and AES-128/SHA as algorithms for both IPsec phases.
- Use default algorithms for Phase1/Phase2 only on R2.
- Do not use protocol 47 in your configuration.
- Ensure that R2 establishes the IPsec tunnel only if R3 provides an IKE ID of **R3.ine.com**.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

When you configure IPSec VPNs using crypto maps, you effectively apply the filter and encrypt/decrypt logic through the use of proxy-ACL. The crypto-map inspects all outgoing traffic for the interface where is applied and triggers ISAKMP negotiations when there is a match. Based on this logic, a crypto-map must always apply to the physical interface where the traffic leaves the router. At the same time, when processing an incoming ISAKMP session, the router will inspect the crypto-map on the ingress interface to find the matching Phase 2 policy. Basically, everything is triggered by the access-lists used in crypto-map entries, and the router parses crypto-map entries to determine whether traffic should or should not be protected.

The IPsec VPN is merely a point-to-point tunnel. To reflect this idea, a concept of virtual tunnel interface (VTI) was introduced into IOS. A VTI is very similar to a regular GRE tunnel that has source and destination tunnel IP addresses and additional tunnel header. However, VTI does not use any additional encapsulation; all packets are tunneled using the regular IPsec ESP headers. The ISAKMP negotiation is self-triggered, and negotiated Phase 2 SA has proxy IDs that cover all possible traffic, equivalent to a `permit ip any any` access-list statement. However, this SA is bound to the VTI itself and only affects traffic going through the tunnel. When the router processes ingress encrypted traffic, it binds it to a certain VTI based on the SPI values. Thus, if you want certain flows to be encrypted, you just need to make sure they are routed across the VTI. There is no more need to construct complex access-lists to filter encrypted traffic; all you need is a route through the tunnel interface.

The first thing that comes to mind is the use of IGP for dynamic route propagation across the VTI. As you remember, using dynamic IGP routing was not allowed with classic IPsec and crypto maps. However, VTIs are point-to-point by design, and therefore it is possible to route multicast traffic (OSPF, RIP, EIGRP) across the tunnel by encapsulating all packets blindly. This allows for establishment of routing protocol adjacencies across the VTI tunnels and permits flexible control over the encrypted traffic. In most cases, you will simply use a different routing protocol for the VPN traffic and advertise the subnets that should be encrypted into the new IGP. Remember that it is generally a bad idea to use the same IGP for non-tunneled and tunneled packets, because this might result in routing loops or recursive lookup errors, where the prefix matching the tunnel destination ends up being learned over the tunnel.

Except for the change in defining protected traffic (ACL based in crypto-maps vs. routing based in VTI), SVTI offers additional benefits when compared with crypto-maps:

- You have an interface that can be monitored, and its status defines the state of the IPsec tunnel; if the line-protocol of the interface is DOWN, encryption is not functional.

https://t.me/learningnets

- It provides a routable interface.
- Crypto-maps require that interesting traffic be detected for the tunnel negotiation to start, and the tunnel is UP based on the configured lifetime parameters and traffic presence; SVTI, because of the `permit ip any any` proxy-ACL, automatically starts tunnel negotiation and keeps it UP as long as there is connectivity between tunnel endpoints and applied configuration is correct.
- It is more scalable because it requires fewer SAs to be negotiated. With ACL-based encryption, routers negotiate a pair of SAs for each entry in the ACL of an IPsec tunnel; with SVTI, because the proxy-ACL is a non-configurable `permit ip any any`, a single pair of SAs is negotiated for any single IPsec tunnel.
- It supports multicast encryption, both control-plane and data-plane.
- Because you have a logical interface, traffic policies such as NAT, ACL, PBR, uRPF, CBAC, and ZBF can be easily applied.

There are two types of VTIs: static (SVTI) and dynamic (DVTI). As the names imply, SVTI is used for static LAN-to-LAN tunnels, and DVTI is used for dynamic LAN-to-LAN or remote-access tunnels; you could say that SVTI replaces static crypto-maps and DVTI replaces dynamic crypto-maps. SVTI offers per-tunnel configuration granularity, because each tunnel is assigned a dedicated VTI. To achieve the same granularity with DVTI, a static virtual-template interface is used for all remote-access sessions, and an on-demand dynamic virtual-access interface is created for each remote-access session where per-tunnel or per-user policies are applied. The dynamic virtual-access interface exists only as long as its associated SAs are active. DVTI uses an on-demand tunnel establishment process similar to crypto-maps.

Although there are some similarities in features and functionality between GRE and VTI, each technology has its own design benefits and flaws and should be used accordingly. For example, GRE is clearly required for DMVPNs, whereas DVTI is clearly required for remote-access VPNs. Considering SVTI and GRE over IPsec, the GRE option will add an additional 4 bytes of overhead, but it supports encapsulation of non-IP protocols (like NHRP; this is the reason why DMVPN uses GRE). GRE is protocol number 47, which is not allowed in this task.

To configure a VTI, you must first create a regular tunnel interface and specify the source/destination IP addresses as usual. Next, you should specify the tunnel mode as `ipsec ipv4`, configure IP addresses to enable IP processing, and apply a tunnel protection profile. The tunnel protection profile is previously configured using the command `crypto ipsec profile`. IPsec profiles specify common Phase2 settings to be reused for multiple interfaces or VPNs.

https://t.me/learningnets

Configuration

On R2, a configuration using ISKAMP profiles has been used to match on the IKE ID provided by R3; default ISAKMP policies and transform-set are used. On R3, there was no need to use ISAKMP profiles, but ISAKMP policies and transform-set to match the defaults of R2 were required.

R2:

```
crypto keyring R2_R3
 pre-shared-key address 136.1.38.3 255.255.255.255 key CISCO
!
crypto isakmp profile R2_R3
 keyring R2_R3
 match identity host R3.ine.com
!
crypto ipsec profile R2_R3
 set isakmp-profile R2_R3
!
interface Tunnel0
 ip address 136.1.23.2 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.38.3
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile R2_R3
```

R3:

```
crypto isakmp policy 100
 encryption aes
 hash sha
 authentication pre-share
 group 5
!
crypto isakmp key CISCO address 136.1.18.2
!
crypto ipsec transform-set AES128_SHA esp-aes esp-sha-hmac
!
crypto isakmp profile R2_R3
 keyring default
 self-identity fqdn
 match identity address 136.1.18.2 255.255.255.255
!
```

https://t.me/learningnets

```
crypto ipsec profile R2_R3
 set transform-set AES128_SHA
 set isakmp-profile R2_R3
!
interface Tunnel0
 ip address 136.1.23.3 255.255.255.0
 tunnel source FastEthernet0/0
 tunnel destination 136.1.18.2
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile R2_R3
```

Verification

If configuration is correct, the tunnel should be automatically established because SVTI tunnel does not require interesting traffic to be established; it is self-triggered. Test connectivity:

```
R2#ping 136.1.23.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.23.3, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Note that even though the configuration used SVTIs, the outputs from `show` commands indicate that a crypto-map has been used with a proxy-ACL of `permit ip any any` :

```
R2#show crypto session remote 136.1.38.3
Crypto session current status

Interface: Tunnel0
Profile: R2_R3Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 Active SAs: 2, origin: crypto map
!
!R2#show crypto ipsec sa | section interface
interface: Tunnel0    Crypto map tag:Tunnel0-head-0
, local addr 136.1.18.2
!
!R2#show crypto map tag Tunnel0-head-0
Crypto Map IPv4 "Tunnel0-head-0" 65536 ipsec-isakmp
        ISAKMP Profile: R2_R3
        ISAKMP profile R2_R3
```

https://t.me/learningnets

```
        Profile name: R2_R3
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={
                default:  { esp-aes esp-sha-hmac  } ,
        }

Crypto Map IPv4 "Tunnel0-head-0" 65537 ipsec-isakmp
        Map is a PROFILE INSTANCE. Peer = 136.1.38.3
ISAKMP Profile: R2_R3
Extended IP access list
access-list  permit ip any any
        ISAKMP profile R2_R3 Current peer: 136.1.38.3
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={  default:  { esp-aes esp-sha-hmac  } ,

        }
        Always create SAs
        Interfaces using crypto map Tunnel0-head-0:
                Tunnel0
```

Verify that required algorithms have been used for Phase1:

https://t.me/learningnets

```
R2#show crypto isakmp sa detail

Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
 1009  136.1.18.2      136.1.38.3 ACTIVE aes  sha    psk  5
  23:46:24
       Engine-id:Conn-id =  SW:9
 1004  136.1.18.2      136.1.18.1 ACTIVE aes  sha    rsig 5
  23:46:06
       Engine-id:Conn-id =  SW:4

IPv6 Crypto ISAKMP SA
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS Routers with SVTI and PKI

Task

- Configure an IPsec tunnel between R2 and R3 without using `crypto map` commands.
- Use digital certificate authentication with R1 as the CA, which is pre-configured from the loaded initial configurations.
- Use default algorithms for Phase1/Phase2 only on R2.
- Do not use protocol 47 in your configuration.
- Ensure that R2 establishes the IPsec tunnel only if R3 provides an IKE ID of **R3.ine.com**.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The only difference from the previous task is enrolling with the CA from R1 and configuring R3's ISAKMP policy to use digital certificate authentication. Also, the key-ring and pre-shared keys are no longer required. Additionally, R3 was globally configured to use its hostname as the IKE ID, so no ISKAMP profile configuration is required.

Configuration

https://t.me/learningnets

R2:

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1
 subject-name CN=R2
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1
crypto pki enroll R1
!
crypto isakmp profile R2_R3
 match identity host R3.ine.com
!
crypto ipsec profile R2_R3
 set isakmp-profile R2_R3
!
interface Tunnel0
 ip address 136.1.23.2 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.38.3
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile R2_R3
```

R3:

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1
 subject-name CN=R3
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1
crypto pki enroll R1
!
crypto isakmp policy 100
 encryption aes
 hash sha
 authentication rsa-sig
```

https://t.me/learningnets

```
 group 5
!
crypto isakmp identity hostname
!
crypto ipsec transform-set AES128_SHA esp-aes esp-sha-hmac
!
crypto ipsec profile R2_R3
 set transform-set AES128_SHA
!
interface Tunnel0
 ip address 136.1.23.3 255.255.255.0
 tunnel source FastEthernet0/0
 tunnel destination 136.1.18.2
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile R2_R3
```

Verification

If you have completed the previous task, to renegotiate the tunnel algorithms, first clear it using the commands **`clear crypto isakmp`** and **`clear crypto sa`** . The tunnel should be automatically established:

```
R2#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1009  136.1.18.2      136.1.38.3 ACTIVE aes  sha    rsig 5
  23:58:46
      Engine-id:Conn-id =  SW:9
1010  136.1.18.2      136.1.38.3 ACTIVE aes  sha    rsig 5
  23:58:47
      Engine-id:Conn-id =  SW:10

IPv6 Crypto ISAKMP SA
!
!R2#show crypto session isakmp profile R2_R3
Crypto session current status

Interface: Tunnel0 Profile: R2_R3
Session status: UP-ACTIVE
```

https://t.me/learningnets

```
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.38.3/500 Active
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 Active SAs: 2, origin: crypto map
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS and ASA with PSK in Main Mode

Task

- Configure a LAN-to-LAN IPsec tunnel between ASA4 and R3 using the following information:
  - Protect traffic between Loopback0 interfaces of R3 and SW1.
  - Use 3DES encryption.
  - Use MD5 hash.
  - Use DH group5.
  - Use the pre-shared-key **CISCO** for authentication.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

LAN-to-LAN VPN configuration on ASA firewall consists of the following steps:

1. Define global ISAKMP (Phase 1) policy using the command `crypto isakmp policy <priority>` . You must set the authentication method, the cipher to protect the ISAKMP SA, and the hash function for integrity checks. You may change the DH group number if you want; the default group is 2. After you have defined the ISAKMP policy, you should enable it on the interface

https://t.me/learningnets

where the VPN tunnel is to be terminated using the command `crypto isakmp enable <ifname>` . Notice that the default ISAKMP policy uses RSA-signatures for authentication, and that the policy-list is scanned from lowest numbers to highest when matching the incoming proposals from the remote peer.

2. Create a tunnel-group (connection profile) for the LAN-to-LAN tunnel using the command `tunnel-group <IP> type ipsec-l2l` and configure parameters using the command `tunnel-group <IP> ipsec-attributes` . The tunnel-group is essentially an object that defines the administrative policy to be applied to the LAN-to-LAN tunnel. It does NOT define the traffic to be protected; rather, most attributes are related to IPsec Phase 1. Up to some point, it could be compared with the ISAKMP profile concept in IOS routers. The tunnel group name must be the IP address of the remote endpoint if you are using pre-shared keys for authentication. A symbolic name could be used in some cases, as we will see in other tasks. When the firewall establishes an IKEv1 VPN tunnel, it looks through the list of local tunnel-groups based on the remote endpoint IKE ID (which by default is an IP address). At the very least, the tunnel group must specify the peer authentication information, such as pre-shared-key, if the global ISAKMP policy uses pre-shared-keys for authentication.

3. Define a crypto transform-set for IPsec Phase 2 using the command `crypto ipsec transform-set` , just like in IOS. This command defines the security parameters for the IPsec tunnel, specifically the cipher, hash function, and, optionally, the mode of the IPsec protection, tunnel or transport.

4. Define a subset of traffic for IPsec protection using an extended access-list. The syntax is `permit <local-ip> <local-mask> <remote-ip> <remote-mask>` and should mirror the entries configured on the remote endpoint; there can be multiple entries in the access-list.

5. Create a crypto-map using the command `crypto-map <NAME> <SEQ-NR> [set|match]` that matches the above-created access-list and sets the remote peer and the transform-set. This completes the settings for IPSec Phase 2. Setting the remote peer is important, because this is how the firewall binds the proxy IDs in the access-list to the tunnel group. There are some optional crypto map parameters that are only supported by ASA

https://t.me/learningnets

firewall. You may use the command `crypto-map <NAME> <SEQ_NR> set connection-type [answer-only|originate-only|bidirectional]` to specify the allowed session mode:

- Answer-only entry will not attempt to negotiate the VPN tunnel for outgoing traffic.
- Originate-only tunnel will not attempt to negotiate the tunnel for incoming traffic.
- The default of bidirectional allows the tunnel to be initiated by any of the two peers.

When you're done with the crypto-map configuration, apply the crypto-map to the interface where you expect the VPN tunnel to be terminated using the command `crypto-map <MAP_NAME> interface <nameif> .`

Make sure that the command `sysopt connection permit-vpn` is configured, which by default it is. When this option is enabled, the decrypted VPN traffic is not subject to the inbound access-list checks on the interface where the tunnel has been terminated, or to the global access-list. For example, if you have this command disabled and the tunnel terminates on the outside interface, the decrypted traffic will be checked against the outside interface inbound access-list or the global access-list if it exists.

Starting with ASA version 8.4, when IKEv1 is enabled on any interface, the ASA activates a default IKEv1 ISAKMP policy with a priority of 65535 (highest) and the following settings:

- authentication pre-share
- encryption 3des
- hash sha
- group 2
- lifetime 86400

Configuration

```
R3:
crypto isakmp policy 10
 encryption 3des
 authentication pre-share
 hash md5
 group 5
!
crypto isakmp key CISCO address 136.1.38.8 255.255.255.255
```

https://t.me/learningnets

```
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
ip access-list extended LOR3_LOSW1
 permit ip host 150.1.3.3 host 150.1.7.7
!
crypto map VPN 20 ipsec-isakmp
 match address LOR3_LOSW1
 set transform 3DES_MD5
 set peer 136.1.38.8
!
interface FastEthernet0/0
 crypto map VPN
```

ASA4:

```
crypto isakmp policy 10
 authentication pre-share
 encryption 3des
 hash md5
 group 5
!
crypto isakmp enable VLAN38
!
tunnel-group 136.1.38.3 type ipsec-l2l
tunnel-group 136.1.38.3 ipsec-attributes
 pre-shared-key CISCO
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
access-list LOSW1_LOR3 permit ip host 150.1.7.7 host 150.1.3.3
!
crypto map VPN 10 match address LOSW1_LOR3
crypto map VPN 10 set peer 136.1.38.3
crypto map VPN 10 set transform-set 3DES_MD5
!
crypto map VPN interface VLAN38
```

Note that the above configuration is for ASA running code 8.2.x. Some of these commands are no longer available (they are hidden) on ASA running code 8.4.x and above, so if still used, ASA will automatically convert it to new commands. (Using the new commands is recommended; the ASA converts it automatically to facilitate the upgrade and migration process.)

https://t.me/learningnets

ASA4:

```
crypto ikev1 policy 10
 authentication pre-share
 encryption 3des
 hash md5
 group 5
!
crypto ikev1 enable VLAN38
!
tunnel-group 136.1.38.3 type ipsec-l2l
tunnel-group 136.1.38.3 ipsec-attributes
 ikev1 pre-shared-key CISCO
!
crypto ipsec ikev1 transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
access-list LOSW1_LOR3 permit ip host 150.1.7.7  host 150.1.3.3
!
crypto map VPN 10 match address LOSW1_LOR3
crypto map VPN 10 set peer 136.1.38.3
crypto map VPN 10 set ikev1 transform-set 3DES_MD5
!
crypto map VPN interface VLAN38
```

Verification

Generate interesting traffic to trigger the IPsec process:

```
SW1#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.7.7 .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/4/9 ms
```

Verify that the tunnel is established with the required parameters; note that state is **MM_ACTIVE**, so IKEv1 Phase1 Main Mode is used:

```
ASA4#show crypto ikev1 sa detail

IKEv1 SAs:

   Active SA: 1
```

https://t.me/learningnets

```
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

1   IKE Peer: 136.1.38.3
    Type    : L2L             Role    : initiator     Rekey   : no State   : MM_ACTIVE
 Encrypt : 3des              Hash    : MD5
    Auth    : preshared       Lifetime: 86400
    Lifetime Remaining: 86151
!
!R3#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF    Status Encr Hash Auth DH Lifetime Cap.
1009  136.1.38.3      136.1.38.8 ACTIVE 3des md5  psk  5
  23:58:54
       Engine-id:Conn-id =  SW:9

IPv6 Crypto ISAKMP SA
!
!ASA4# show vpn-sessiondb l2l filter ipaddress 136.1.38.3

Session Type: LAN-to-LAN

Connection   : 136.1.38.3
Index        : 1                       IP Addr      : 136.1.38.3
Protocol     : IKEv1 IPsec Encryption   : 3DES                    Hashing      : MD5

Bytes Tx     : 400                     Bytes Rx     : 400
Login Time   : 17:50:47 UTC Sat May 4 2013
Duration     : 0h:12m:34s
```

Verify that ICMP packets were encrypted/decrypted:

```
ASA4#show crypto ipsec sa identity
interface: VLAN38
    Crypto map tag: VPN, seq num: 10, local addr: 136.1.38.8
 access-list LOSW1_LOR3 extended permit ip host 150.1.7.7 host 150.1.3.3
      local ident (addr/mask/prot/port): (150.1.7.7/255.255.255.255/0/0)
      remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0) current_peer: 136.1.38.3
```

https://t.me/learningnets

```
 #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4

      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 4, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 136.1.38.8/0, remote crypto endpt.: 136.1.38.3/0
      path mtu 1500, ipsec overhead 58, media mtu 1500
      current outbound spi: 064983E8
      current inbound spi : 07414A1E
```

Verify IKEv1 Phase2 algorithms:

```
ASA4# show crypto ipsec sa map VPN
cryptomap: VPN
    Crypto map tag: VPN, seq num: 10, local addr: 136.1.38.8

      access-list LOSW1_LOR3 extended permit ip host 150.1.7.7 host 150.1.3.3
      local ident (addr/mask/prot/port): (150.1.7.7/255.255.255.255/0/0)
      remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
      current_peer: 136.1.38.3

      #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
      #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 4, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 136.1.38.8/0, remote crypto endpt.: 136.1.38.3/0
      path mtu 1500, ipsec overhead 58, media mtu 1500
      current outbound spi: 58A5D602
      current inbound spi : DD5A6F3F
 inbound esp sas:
      spi: 0xDD5A6F3F (3713691455) transform: esp-3des esp-md5-hmac no compression
         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 4096, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4373999/3328)
         IV size: 8 bytes
         replay detection support: Y
         Anti replay bitmap:
```

https://t.me/learningnets

```
      0x00000000 0x0000001F outbound esp sas:
    spi: 0x58A5D602 (1487263234) transform: esp-3des esp-md5-hmac no compression

      in use settings ={L2L, Tunnel, }
      slot: 0, conn_id: 4096, crypto-map: VPN
      sa timing: remaining key lifetime (kB/sec): (4373999/3328)
      IV size: 8 bytes
      replay detection support: Y
      Anti replay bitmap:
       0x00000000 0x00000001
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS and ASA with PSK in Aggressive Mode

Task

- Configure a LAN-to-LAN IPsec tunnel between ASA4 and R3 using the following information:
  - Protect traffic between Loopback0 interfaces of R3 and SW1.
  - Use 3DES encryption.
  - Use MD5 hash.
  - Use DH group5.
  - Use the pre-shared-key **CISCO** for authentication.
  - Define the connection profile on ASA as `tunnel-group R3.ine.com .`
- ASA4 should initiate the session in aggressive mode.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

When configuring a connection-profile on ASA by using the `tunnel-group` command, we can specify either an IP address or a name, and this information can be used by the ASA to bind inbound IKEv1 VPN session negotiations to certain connection-profiles. ASA has multiple options to bind incoming VPN sessions to certain profiles:

- Based on one or more digital certificate attributes (not enabled by default).

https://t.me/learningnets

- Based on the OU (Organizational Unit) field of the peer's presented digital certificate.
- Based on the IKE ID presented by the remote peer, the peer that can be an IP address (default), the hostname, or the DN (Distinguished Name), or in some cases an EzVPN group name (for example, when you are using the Cisco EzVPN client or the EzVPN Remote feature).
- Based on the remote endpoint source IP address of the VPN session.

The order in which the ASA processes incoming VPN connections can be seen in the following output. It stops on first match (note that the first option is disabled by default).

```
ASA4# show running-config all tunnel-group-map

no tunnel-group-map enable rules
tunnel-group-map enable ou
tunnel-group-map enable ike-id
tunnel-group-map enable peer-ip
```

Depending on Phase1 mode (Main or Aggressive) and type of authentication (pre-shared-key or digital certificate), not all matching options are valid. The first two options are valid only when using digital certificate authentication, regardless of the Phase1 mode. The last option (IP address) acts as a gateway of last resort in case the connection was not matched based on any of the three previous options and is valid for both Phase1 modes and authentication methods. The third option (IKE ID) does not work for Main mode with pre-shared-key authentication; let's find out why.

https://t.me/learningnets

In addition to the different number of exchanged messages between Main mode and Aggressive mode, another important aspect is that Main mode protects the IKE IDs of the peers, whereas with Aggressive Mode, these are sent in clear-text before DH exchange is performed and the secure communication channel is created. When pre-shared-keys are used for authentication, they are also used to generate the shared encryption key for ISAKMP SA (along with the DH generated key). When an ISAKMP responder receives a Main mode proposal from initiator and negotiates authentication based on pre-shared-keys, it should generate the shared encryption key. This requires knowing the pre-shared-key of the remote peer in advance; however, the responder does not know the IKE ID of the initiator yet, but only its IP address. Therefore, the only way to select the proper pre-shared-key in Main mode is to look up the key in the database based on the initiator's IP address. Even if you use hostnames for IKE IDs with pre-shared-key authentication, the keys and tunnel-group names are still matched based on the IP addresses. This is the only supported connection-profile matching mode for Main mode with pre-shared-key authentication.

The task requires that you configure a connection-profile based on the name **R3.ine.com** instead of the IP address of the remote peer. Of the four options for matching a VPN session to a connection profile, we cannot use the first two because these are valid only for digital certificate authentication, and the last option (IP address) can be used only if the tunnel-group is configured based on IP address. The third option is the only one left, but matching on IKE ID with pre-shared-key is possible only with aggressive mode. We will modify the ASA to run in aggressive mode for its VPN tunnel with R3. However, note that by default ASA forces it to use DH group2 (regardless of the configured/negotiated ISAKMP policy), so we need to modify it to match the configured DH group5 from the ISAKMP policy.

We also must ensure that R3 uses an IKE ID of **R3.ine.com**, and we can set it globally or through an ISAKMP profile. Because we did not use ISAKMP profiles in our configuration, we have two options: `crypto isakmp identity dn` and `crypto isakmp identity hostname` . Note that even though DN IKE ID should be used when authentication is based on digital certificates, it also works with pre-shared-keys, and the router uses its FQDN (hostname.domain) instead of the DN value from the certificate.

Configuration

ASA4:

```
crypto isakmp policy 10
 authentication pre-share
 encryption 3des
 hash md5
```

https://t.me/learningnets

```
 group 5
!
crypto isakmp enable VLAN38
!
tunnel-group R3.ine.com type ipsec-l2l
tunnel-group R3.ine.com ipsec-attributes
 pre-shared-key CISCO
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
access-list LOSW1_LOR3 permit ip host 150.1.7.7 host 150.1.3.3
!
crypto map VPN 10 match address LOSW1_LOR3
crypto map VPN 10 set peer 136.1.38.3
crypto map VPN 10 set transform-set 3DES_MD5
crypto map VPN 10 set phase1-mode aggressive group5
!
crypto map VPN interface VLAN38
```

Note that the above configuration is for ASA running code 8.2.x. Some of these commands are no longer available (they are hidden) on ASA running code 8.4.x and above; if the commands are still used, ASA will automatically convert them to new commands:

ASA4:

```
crypto ikev1 policy 10
 authentication pre-share
 encryption 3des
 hash md5
 group 5
!
crypto ikev1 enable VLAN38
!
tunnel-group R3.ine.com type ipsec-l2l
tunnel-group R3.ine.com ipsec-attributes
 ikev1 pre-shared-key CISCO
!
crypto ipsec ikev1 transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
access-list LOSW1_LOR3 permit ip host 150.1.7.7  host 150.1.3.3
!
crypto map VPN 10 match address LOSW1_LOR3
crypto map VPN 10 set peer 136.1.38.3
crypto map VPN 10 set ikev1 transform-set 3DES_MD5
crypto map VPN 10 set ikev1 phase1-mode aggressive group5
```

https://t.me/learningnets

```
!
crypto map VPN interface VLAN38



R3:

crypto isakmp policy 10
 encryption 3des
 authentication pre-share
 hash md5
 group 5
!
crypto isakmp identity hostname
crypto isakmp key CISCO address 136.1.38.8 255.255.255.255
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
ip access-list extended LOR3_LOSW1
 permit ip host 150.1.3.3 host 150.1.7.7
!
crypto map VPN 20 ipsec-isakmp
 match address LOR3_LOSW1
 set transform 3DES_MD5
 set peer 136.1.38.8
!
interface FastEthernet0/0
 crypto map VPN
```

Verification

If you've done the previous task, to renegotiate the tunnel settings, first clear it using **`clear crypto isakmp sa`** or **`clear crypto ikev1 sa`**, depending on the ASA running code. Then generate interesting traffic and verify that the tunnel is established; note that the state is **AM_MODE**, which means Phase 1 was completed using aggressive mode:

```
SW1#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.7.7 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
!
```

https://t.me/learningnets

```
!ASA4#show crypto ikev1 sa detail


IKEv1 SAs:

   Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

1   IKE Peer: 136.1.38.3
    Type    : L2L             Role    : initiator     Rekey   : no State   : AM_ACTIVE
Encrypt : 3des            Hash    : MD5
    Auth    : preshared       Lifetime: 86400
    Lifetime Remaining: 86384
!
!R3#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF    Status Encr Hash Auth DH Lifetime Cap.
1010  136.1.38.3      136.1.38.8ACTIVE 3des md5  psk  5
  23:59:28
       Engine-id:Conn-id =  SW:10

IPv6 Crypto ISAKMP SA
```

Check the connection profile being matched:

```
ASA4#show vpn-sessiondb l2l filter ipaddress 136.1.38.3

Session Type: LAN-to-LAN
Connection   : R3.ine.com
Index        : 9                       IP Addr      : 136.1.38.3
Protocol     : IKEv1 IPsecEncryption   : 3DES                    Hashing      : MD5

Bytes Tx     : 400                     Bytes Rx     : 400
Login Time   : 06:39:35 UTC Sun May 5 2013
Duration     : 0h:00m:58s
```

https://t.me/learningnets

Verify that Phase2 is established and check negotiated parameters:

```
ASA4# show crypto ipsec sa  user R3.ine.com
username: R3.ine.com
    Crypto map tag: VPN, seq num: 10, local addr: 136.1.38.8

      access-list LOSW1_LOR3 extended permit ip host 150.1.7.7 host 150.1.3.3
      local ident (addr/mask/prot/port): (150.1.7.7/255.255.255.255/0/0)
      remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
      current_peer: 136.1.38.3
 #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
 #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 4, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 136.1.38.8/0, remote crypto endpt.: 136.1.38.3/0
      path mtu 1500, ipsec overhead 58, media mtu 1500
      current outbound spi: 60C29DC3
      current inbound spi : 72FD53AD
 inbound esp sas:
      spi: 0x72FD53AD (1929204653) transform: esp-3des esp-md5-hmac no compression
         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 8192, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4373999/3392)
         IV size: 8 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x0000001F outbound esp sas:
      spi: 0x60C29DC3 (1623367107) transform: esp-3des esp-md5-hmac no compression

         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 8192, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4373999/3392)
         IV size: 8 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x00000001
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS and ASA with PKI

Task

- Configure a LAN-to-LAN IPsec tunnel between ASA3 and R3 to match following requirements:
  - Protect traffic between Loopback0 interfaces of R3 and SW2.
  - Use AES-256 encryption.
  - Use SHA1 hash.
  - Use DH group5.
  - Use R1-issued certificates for authentication.
- Ensure that ASA3 can act as both IKEv1 initiator and responder.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

By default, the ASA performs some additional security checks when using certificate authentication for IKEv1 IPsec VPNs. These are related to the IKEv1 IKE ID presented by the remote peer. By default, without configuration modifications, the following IKEv1 IDs are used by the ASA/IOS:

- Pre-shared-key authentication:
  - IOS uses an IKE ID of IP address.

https://t.me/learningnets

    - ASA uses an IKE ID of IP address.
- Digital certificate authentication:
    - IOS uses an IKE ID of FQDN (hostname.domain-name).
    - ASA uses an IKE ID of DN (Distinguished Name field of the certificate).

By default, when creating a CSR on the ASA, it automatically copies its FQDN in the SAN (Subject Alternate Name) field, and this is not configurable. IOS routers behave differently. By default, no SAN information is specified in the CSR; however, IOS allows you to manually specify it using the trustpoint-level command `subject-alt-name`, but only for self-signed certificates. SAN scope is to allow an entity for which a digital certificate has been issued to be known by multiple DNS/FQDN names, in addition to the information from the CN field.

By default, when the ASA needs to establish an IKEv1 IPsec VPN using digital certificate authentication, it performs the following validation checks:

- If the remote peer–presented IKE ID is DN, it verifies that the presented IKE ID value matches the DN information from the received certificate.
- If the remote peer–presented IKE ID is FQDN, it verifies that the presented IKE ID value matches the SAN information from the received certificate.

If the compared values don't match, the ASA rejects the IKEv1 IPsec session. This is controlled by the tunnel-group-level command `peer-id-validate`, which has the following options:

- `req`. This is the default and basically enforces the two compared values to match, which means that the certificate must have information in the CN or SAN fields.
- `cert`. This means that the two values are compared only if the certificate has information in the CN or SAN fields.
- `nocheck`. This means that the ASA no longer performs the validation.

With the default settings, a digital certificate–based IKEv1 IPsec VPN authentication between an IOS router and ASA firewall will not be successful. IOS will use an IKE ID of FQDN, and ASA will compare the FQDN with the SAN field from the certificate, which does not exist. To fix it, there are three options:

- Configure the IOS router to use an IKE ID of DN.
- Configure the ASA with the `cert` option. Because the SAN field does not exist, the verification is no longer performed.
- Configure the ASA with the `nocheck` option.

When the IOS router uses an FQDN IKE ID, if you issue a `debug crypto ikev1 128` on

https://t.me/learningnets

the ASA, you'll see the following messages being logged:

```
[IKEv1 DEBUG]Group = 136.1.38.3, IP = 136.1.38.3, peer ID type 2 received (FQDN)
[IKEv1]Group = 136.1.38.3, IP = 136.1.38.3, Unable to compare IKE ID against peer cert Subject Alt Name
```

If you configure the IOS router to use a DN IKE ID, you'll see the following message being logged, indicating the change in IKE ID:

```
[IKEv1 DEBUG]Group = 136.1.38.3, IP = 136.1.38.3, peer ID type 9 received (DER_ASN1_DN)
```

Additionally, the ASA requires the following configurations for successful tunnel establishment, depending on its role (initiator or responder):

- For the ASA to initiate an IKEv1 IPsec tunnel using digital certificate authentication, you must specify at crypto-map level the certificate to be used based on the name of the trustpoint using the command **`crypto map <NAME> <SEQ_NR> set trustpoint <NAME> .`**
- For the ASA to respond to an IKEv1 IPsec tunnel using digital certificate authentication, you must specify at tunnel-group level the certificate to be used based on the trustpoint using the command **`trust-point <NAME> .`**

Configuration

ASA3:

```
crypto key generate rsa general-keys label IPSEC modulus 1024
!
crypto ca trustpoint R1
 keypair IPSEC
 enrollment url http://150.1.1.1
 revocation-check crl none
 subject-name CN=ASA3.ine.com
 no ip-address
 no serial-number
!
crypto ca authenticate R1
crypto ca enroll R1
!
crypto isakmp policy 10
 authentication rsa-sig
 encryption aes-256
 hash sha
 group 5
!
```

https://t.me/learningnets

```
crypto isakmp enable VLAN78
!
tunnel-group 136.1.38.3 type ipsec-l2l
tunnel-group 136.1.38.3 ipsec-attributes
  trust-point R1
!
crypto ipsec transform-set AES256_SHA esp-aes-256 esp-sha-hmac
!
access-list LOSW2_LOR3 permit ip host 150.1.8.8 host 150.1.3.3
!
crypto map VPN 10 match address LOSW2_LOR3
crypto map VPN 10 set peer 136.1.38.3
crypto map VPN 10 set transform-set AES256_SHA
crypto map VPN 10 set trustpoint R1
!
crypto map VPN interface VLAN78
```

Note that the above configuration is for ASA running code 8.2.x. Some of these commands are no longer available (they are hidden) on ASA running code 8.4.x and above; if the commands are still used, ASA will automatically convert them to new commands:

ASA3:

```
crypto key generate rsa general-keys label IPSEC modulus 1024
!
crypto ca trustpoint R1
 keypair IPSEC
 enrollment url http://150.1.1.1
 revocation-check crl none
 subject-name CN=ASA3.ine.com
 no ip-address
 no serial-number
!
crypto ca authenticate R1
crypto ca enroll R1
!
crypto ikev1 policy 10
 authentication rsa-sig
 encryption aes-256
 hash sha
 group 5
!
crypto ikev1 enable VLAN78
!
tunnel-group 136.1.38.3 type ipsec-l2l
```

https://t.me/learningnets

```
tunnel-group 136.1.38.3 ipsec-attributes
  ikev1 trust-point R1
!
crypto ipsec ikev1 transform-set AES256_SHA esp-aes-256 esp-sha-hmac
!
access-list LOSW2_LOR3 permit ip host 150.1.8.8 host 150.1.3.3
!
crypto map VPN 10 match address LOSW2_LOR3
crypto map VPN 10 set peer 136.1.38.3
crypto map VPN 10 set ikev1 transform-set AES256_SHA
crypto map VPN 10 set trustpoint R1
!
crypto map VPN interface VLAN78
```

R3:

```
crypto key generate rsa general-keys modulus 1024 label IPSEC
!
crypto pki trustpoint R1
 subject-name CN=R3
 enrollment url http://150.1.1.1
 rsakeypair IPSEC
 ip-address none
 serial-number none
!
crypto pki authenticate R1
crypto pki enroll R1
!
crypto isakmp policy 10
 encryption aes 256
 authentication rsa-sig
 hash sha
 group 5
!
crypto ipsec transform-set AES256_SHA esp-aes 256 esp-sha-hmac
crypto isakmp identity dn
!
ip access-list extended LOR3_LOSW2
 permit ip host 150.1.3.3 host 150.1.8.8
!
crypto map VPN 30 ipsec-isakmp
 match address LOR3_LOSW2
 set transform AES256_SHA
 set peer 136.1.78.7
```

https://t.me/learningnets

```
!
interface FastEthernet0/0
 crypto map VPN
```

Verification

Generate interesting traffic to trigger the IPsec process.

```
SW2#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
```

Verify tunnel establishment with required parameters:

```
ASA3#show crypto ikev1 sa detail

IKEv1 SAs:

   Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1

1   IKE Peer: 136.1.38.3
    Type    : L2L             Role    : initiator  Rekey   : no              State   : MM_ACTIVE
Encrypt : aes-256         Hash    : SHA
    Auth    : rsa             Lifetime: 86400
    Lifetime Remaining: 86363
!
!ASA3#show vpn-sessiondb l2l filter ipaddress 136.1.38.3

Session Type: LAN-to-LAN

Connection   : 136.1.38.3
Index        : 21                    IP Addr      : 136.1.38.3
Protocol     : IKEv1 IPsecEncryption   : AES256                Hashing      : SHA1
Bytes Tx     : 400                   Bytes Rx     : 400
Login Time   : 07:44:07 UTC Mon May 6 2013
Duration     : 0h:01m:26s
!
!R3#show crypto isakmp sa detail
```

https://t.me/learningnets

```
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF    Status Encr Hash Auth DH Lifetime Cap.
1011  136.1.38.3      136.1.78.7 ACTIVE aes  sha  rsig 5
  23:58:48
       Engine-id:Conn-id =  SW:11

IPv6 Crypto ISAKMP SA
```

Verify Phase2 negotiated parameters and that traffic was encrypted/decrypted:

```
ASA3# show crypto ipsec sa peer 136.1.38.3
peer address: 136.1.38.3
    Crypto map tag: VPN, seq num: 10, local addr: 136.1.78.7

      access-list LOSW2_LOR3 extended permit ip host 150.1.8.8 host 150.1.3.3
      local ident (addr/mask/prot/port): (150.1.8.8/255.255.255.255/0/0)
      remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
      current_peer: 136.1.38.3
 #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 4, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 136.1.78.7/0, remote crypto endpt.: 136.1.38.3/0
      path mtu 1500, ipsec overhead 74, media mtu 1500
      current outbound spi: FDD9174A
      current inbound spi : D6A4AA5F
 inbound esp sas:
      spi: 0xD6A4AA5F (3601115743) transform: esp-aes-256 esp-sha-hmac no compression
         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 4096, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4373999/3408)
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x0000001F
```

https://t.me/learningnets

```
outbound esp sas:

      spi: 0xFDD9174A (4258862922) transform: esp-aes-256 esp-sha-hmac no compression

         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 4096, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4373999/3408)
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x00000001
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 L2L Between IOS and ASA with Certificate Map

Task

- Configure ASA3 to bind any incoming IKEv1 IPsec sessions to its tunnel with R3 if the following conditions are true:
  - Session uses digital certificate authentication.
  - Remote peer certificate has a CN value for the issuer of **r1.ine.com**.
  - Remote peer certificate has a CN value for the subject-name of **r3.ine.com**.
- If the remote peer certificate has the string **ccie-security** in its subject name, the IPsec session should also be mapped to its tunnel with R3.

> **A Note on Task Initial Configuration Files**: This task requires the previous task to have been configured, so the configurations that resulted from the previous task are loaded on R3 and ASA3.

Overview

ASA allows binding of any IKEv1 IPsec sessions to a certain tunnel-group (connection profile) based on remote peer certificate attributes, but this is disabled by default. The required configuration steps are:

1. Enable the functionality using the global command `tunnel-group-map enable rules .`
2. Define the certificate attributes rules in a certificate map using the global command `crypto ca certificate map <NAME> <SEQ_NR> .`
3. Bind the defined rules to a tunnel-group using the global command `tunnel-group-map <MAP_NAME> <MAP_SEQ_NR> <TUNNEL_NAME> .`

Certificate map functionality is very similar to a route-map:

- It is defined by a name.

https://t.me/learningnets

- Multiple entries can exist within the same certificate map based on sequence numbers.
- Multiple rules can exist within the same entry; between these rules there is an **AND** logic applied unconditionally, so all rules must be matched.

Note that a certificate map is not active unless applied to a certain tunnel-group. Multiple entries of the same map can be applied to the same tunnel-group. The strings matched in the rules are not case-sensitive.

Configuration

ASA3:

```
crypto ca certificate map R3 10
 issuer-name attr cn eq r1.ine.com
 subject-name attr cn eq r3.ine.com
!
crypto ca certificate map R3 20
 subject-name co ccie-security
!
tunnel-group-map enable rules
tunnel-group-map R3 10 136.1.38.3
tunnel-group-map R3 20 136.1.38.3
```

Verification

Clear the SAs with the command `clear crypto ikev1 sa` . If you issue a `debug crypto ikev1 128` on the ASA and generate interesting traffic to trigger the IPsec process, you will see the following log messages that indicate ASA matching the VPN session based on certificate map rules:

https://t.me/learningnets

```
[IKEv1]IP = 136.1.38.3, Trying to find group via cert rules...
[IKEv1]IP = 136.1.38.3, Connection landed on tunnel_group 136.1.38.3



SW2#ping 150.1.3.3 source loopback0


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/3/9 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 GRE over IPsec with Crypto Map and PSK

Task

- Configure an IKEv1 IPsec tunnel between R1 and R2 without using `crypto ipsec profile` commands.
- Use a pre-shared-key of **CISCO** and AES-128/SHA as algorithms for both IPsec phases.
- Use default algorithms for Phase1/Phase2 on R2.
- Use protocol number 47 and ensure that both routers negotiate the minimum possible overhead.
- Ensure that R1 and R2 establish an OSPF adjacency over the tunnel.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Before the introduction of VTI, using GRE over IPSec was the only way to configure dynamic routing over encrypted tunnels. GRE allows encapsulation of multicast traffic, including routing updates. However, the use of GRE imposes an additional 4 bytes of overhead for VPN traffic, and is therefore less efficient if compared with SVTI.

There are two ways to encrypt GRE traffic:

- Configure and apply an IPSec encryption profile. This is the most convenient and

https://t.me/learningnets

recommended way, and it's the same configuration you'll be using when configuring SVTI.

- Create an access-list that matches GRE encapsulated traffic, (traffic sourced off the tunnel endpoints) and apply a crypto-map that encrypts this traffic. This is the legacy configuration for IPSec with GRE.

Before IOS 12.2(13)T, the crypto-map had to be applied to both the physical interface and the logical interfaces, such as the P2P GRE tunnel interfaces. After 12.2(13)T , the crypto-map must be applied only to the physical interface, like in any regular crypto-map-based IKEv1 IPsec VPN. Without configuring GRE keepalives, which by default are disabled, a GRE tunnel is in the UP state as long as the router has a route toward the tunnel destination; this behavior is completely different if compared with SVTI where the tunnel is in the UP state only when the IKEv1 IPsec tunnel is established.

Because using GRE encapsulation makes the router both the VPN gateway and the protected traffic source/destination, it allows configuring and using IPsec in transport. The end result is that less overhead is used for the IPsec tunnel. To use transport mode, configure the IPsec transform set appropriately.

Configuration

```
R1:
crypto isakmp policy 50
 authentication pre-share
 encryption aes
 hash sha
 group 5
!
ip access-list extended R1_R2_GRE
 permit gre host 136.1.18.1 host 136.1.18.2
!
crypto isakmp key CISCO address 136.1.18.2 255.255.255.255
crypto ipsec transform-set AES128_SHA esp-aes esp-sha-hmac
 mode transport
!
interface Tunnel1
 ip address 136.1.12.1 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.18.2
 tunnel mode gre ip
!
router ospf 1
 network 136.1.12.0 0.0.0.255 area 0
```

https://t.me/learningnets

```
!
crypto map VPN 20 ipsec-isakmp
 match address R1_R2_GRE
 set peer 136.1.18.2
 set transform-set AES128_SHA
!
interface GigabitEthernet0/0
 crypto map VPN


R2:

ip access-list extended R1_R2_GRE
 permit gre host 136.1.18.2 host 136.1.18.1
!
crypto isakmp key CISCO address 136.1.18.1 255.255.255.255
!
interface Tunnel1
 ip address 136.1.12.2 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.18.1
 tunnel mode gre ip
!
router ospf 1
 network 136.1.12.0 0.0.0.255 area 0
!
crypto map VPN 20 ipsec-isakmp
 match address R1_R2_GRE
 set peer 136.1.18.1
!
interface GigabitEthernet0/0
 crypto map VPN
```

Verification

Because OPSF was configured over the tunnel, IPsec SAs will be automatically established because there is interesting traffic initiated from both sides. Verify that OSPF adjacency is UP:

https://t.me/learningnets

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.2.2         0   FULL/  -        00:00:37    136.1.12.2      Tunnel1
```

Verify that the IPsec tunnel is established using the requested parameters and that transport mode is used to minimize overhead:

https://t.me/learningnets

```
R1#show crypto isakmp sa detail

Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
 1015  136.1.18.1      136.1.18.2 ACTIVE aes sha     psk  5
  23:48:53
       Engine-id:Conn-id =  SW:15

IPv6 Crypto ISAKMP SA
!
!R1#show crypto ipsec sa peer 136.1.18.2 | sec esp sas
 inbound esp sas:

      spi: 0x89FA71DA(2314891738) transform: esp-aes esp-sha-hmac ,
 in use settings ={Transport, }

        conn id: 2033, flow_id: Onboard VPN:33, sibling_flags 80000000, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4366090/3430)
        IV size: 8 bytes
        replay detection support: Y
        Status: ACTIVE(ACTIVE)
 outbound esp sas:

      spi: 0xE2E6ADBC(3806768572) transform: esp-aes esp-sha-hmac ,
 in use settings ={Transport, }

        conn id: 2034, flow_id: Onboard VPN:34, sibling_flags 80000000, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4366090/3430)
        IV size: 8 bytes
        replay detection support: Y
        Status: ACTIVE(ACTIVE)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 GRE over IPsec with IPsec Profiles and PSK

Task

- Configure an IPsec tunnel between R1 and R2 using IPsec profiles.
- Use a pre-shared-key of **CISCO** and AES-128/SHA as algorithms for both IPsec phases.
- Use default algorithms for Phase1/Phase2 on R2.
- Ensure that both routers initiate the IPsec session in Aggressive mode.
- Ensure that R1 and R2 establish an OSPF adjacency over the tunnel.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IPsec profiles have also been used in the VTI section. We must use ISAKMP profiles to set the Phase1 mode to Aggressive. The `keyring default` command has been used to use the globally defined pre-shared-key, without configuring any keyrings. Note that a `match identity` is required to make the ISAKMP profile valid; otherwise, the router will log a message similar to `ISAKMP:(0):ISAKMP profile R1_R2_GRE is not complete` , and will not be able to use it.

Configuration

https://t.me/learningnets

R1:

```
crypto isakmp policy 50
 authentication pre-share
 encryption aes
 hash sha
 group 5
!
crypto isakmp profile R1_R2_GRE
 match identity address 136.1.18.2 255.255.255.255
 initiate mode aggressive
 keyring default
!
crypto ipsec profile R1_R2_GRE
 set transform-set AES128_SHA
 set isakmp-profile R1_R2_GRE
!
crypto isakmp key CISCO address 136.1.18.2 255.255.255.255
crypto ipsec transform-set AES128_SHA esp-aes esp-sha-hmac
!
interface Tunnel1
 ip address 136.1.12.1 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.18.2
 tunnel mode gre ip
 tunnel protection ipsec profile R1_R2_GRE
!
router ospf 1
 network 136.1.12.0 0.0.0.255 area 0
```

R2:

```
crypto isakmp profile R1_R2_GRE
 match identity address 136.1.18.1 255.255.255.255
 initiate mode aggressive
 keyring default
!
crypto ipsec profile R1_R2_GRE
 set isakmp-profile R1_R2_GRE
!
crypto isakmp key CISCO address 136.1.18.1 255.255.255.255
!
interface Tunnel1
 ip address 136.1.12.2 255.255.255.0
 tunnel source GigabitEthernet0/0
```

https://t.me/learningnets

```
 tunnel destination 136.1.18.1
 tunnel mode gre ip
 tunnel protection ipsec profile R1_R2_GRE
!
router ospf 1
 network 136.1.12.0 0.0.0.255 area 0
```

Verification

To verify that Aggressive mode is used, use a **`debug crypto isakmp`** on either router. You will see that negotiation starts in AM mode; when Phase1 is completed, the router should log a message similar to the following:

```
ISAKMP:(1006):Input = IKE_MESG_FROM_PEER, IKE_AM_EXCH
ISAKMP:(1006):Old State = IKE_I_AM1  New State = IKE_P1_COMPLETE
```

Verify OSPF adjacency and IP connectivity over the tunnel:

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.2.2         0   FULL/  -        00:00:35    136.1.12.2      Tunnel1
!
!R1#ping 136.1.12.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.12.2, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Verify that ISAKMP profile is being matched; note the implicit ACL used by the IPsec profile, with a match on GRE between the two tunnel endpoints:

```
R1#show crypto session isakmp profile R1_R2_GRE detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation

Interface: Tunnel1Profile: R1_R2_GRE
Uptime: 00:05:13
Session status: UP-ACTIVE
Peer: 136.1.18.2 port 500 fvrf: (none) ivrf: (none)
```

https://t.me/learningnets

```
      Phase1_id: 136.1.18.2
      Desc: (none)
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.18.2/500 Active
          Capabilities:(none) connid:1010 lifetime:23:54:46
  IPSEC FLOW: permit 47 host 136.1.18.1 host 136.1.18.2
        Active SAs: 2, origin: crypto map Inbound:  #pkts dec'ed 46 drop 0 life (KB/Sec) 4311064/3286
Outbound: #pkts enc'ed 49 drop 0 life (KB/Sec) 4311063/3286
!
!R1#show crypto isakmp profile

IKEv1 PROFILE R1_R2_GRE Ref Count = 5

   Identities matched are:
    ip-address 136.1.18.2 255.255.255.255
   Certificate maps matched are:
   keyring(s): default
   trustpoint(s): <all>
```

R1's configured transform-set will negotiate tunnel mode, and R2's default transform-set will negotiate transport mode. Because of the use of GRE, both routers will negotiate transport mode to minimize overhead:

```
R1#show crypto ipsec transform-set  AES128_SHA
{ esp-aes esp-sha-hmac  } will negotiate = { Tunnel,  },
!
!R2#show crypto ipsec transform-set
Transform set default: { esp-aes esp-sha-hmac  } will negotiate = { Transport,  },
!
!R1#show crypto ipsec sa | section esp sas
inbound esp sas:
      spi: 0x231FB029(589279273)
        transform: esp-aes esp-sha-hmac , in use settings ={Transport, }
        conn id: 2027, flow_id: Onboard VPN:27, sibling_flags 80000000, crypto map: Tunnel1-head-0
        sa timing: remaining key lifetime (k/sec): (4178555/3522)
        IV size: 16 bytes
        replay detection support: Y
        Status: ACTIVE(ACTIVE)
 outbound esp sas:
      spi: 0x5EB3693A(1588816186)
        transform: esp-aes esp-sha-hmac , in use settings ={Transport, }

        conn id: 2028, flow_id: Onboard VPN:28, sibling_flags 80000000, crypto map: Tunnel1-head-0
        sa timing: remaining key lifetime (k/sec): (4178555/3522)
        IV size: 16 bytes
        replay detection support: Y
```

https://t.me/learningnets

```
        Status: ACTIVE(ACTIVE)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VRF Aware VPN with Crypto Map and Global FVRF

Task

- Configure an IKEv1 IPsec tunnel between R1 and R3 to provide connectivity between Loopback0 interfaces.
- Use the following parameters:
  - Pre-shared-key of **CISCO**
  - 3DES/MD5 for both phases
  - DH with 1024-bit strength
- On both routers, configure a VRF named **IVRF-CLIENT** and assign Loopback0 interfaces to it.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

A Virtual Routing and Forwarding (VRF) configuration without the involvement of MPLS is also known as VRF Lite. It's basically a method of virtualizing the IOS router, somewhat similar but not identical to security contexts from firewalls. Just like with security contexts, there are some features that don't work with VRFs. The following steps are required for VRF configuration:

1. Enable CEF.

https://t.me/learningnets

2. Configure VRFs (virtual router instances) using the command `ip vrf <NAME>` or `vrf definition <NAME> .`
3. Assign interfaces to VRFs using the command `ip vrf forwarding <NAME>` or `vrf forwarding <NAME> .`

Notice that when an interface is assigned to a VRF, its IP address configuration is lost and must be reconfigured. After interfaces are assigned to their VRFs, the immediate result is having multiple routing tables, one per each VRF; you can visualize it using the command `show ip route vrf <NAME> .` With VRF Lite, without some tweaks, you cannot route packets from one VRF to another VRF on the same router, because that would break the scope of configuring VRFs. The VRF configuration from one router is not correlated to the VRF configuration from a another router; it is locally significant, just like with Private VLANs before VTP version 3.

Some functionalities (such as ACLs) are not VRF dependent (you configure the ACL normally and apply it at the interface level), whereas others (such as routing protocols and IPsec) are VRF dependent (you need to specify the VRF for which you want the feature configured). One important aspect to remember is that all non-VRF-assigned interfaces remain in the default VRF, called `global VRF .`

From the IPsec perspective, VRFs help us achieve one very simple task: binding two VRFs together. Why would that be of interest? Imagine that your Internet/WAN connection where IPsec sessions are terminated is in one VRF, and the data required to be accessed after the tunnel is established is in another VRF. VRF-aware IPsec allows you to have encrypted packets flow in one VRF while decrypted packets flow in another VRF, and the bridge between the two VRFs is active only after the IPsec tunnel is established.

Before configuration, familiarize yourself with two new terms:

- **FVRF** (Front VRF, or Front-door VRF): Where the IPsec negotiation takes place and encrypted packets flow.
- **IVRF** (Inside VRF, or Back-door VRF): Where the data being accessed securely over the IPSec tunnel exists, where decrypted packets flow.

VRF-aware IPsec is supported with both classic crypto-map configuration and newer SVTI/GRE and IPsec profiles configuration. For each of these, the following scenarios are possible (VRF1 and VRF2 are random names for VRFs):

- FVRF=IVRF=global VRF (regular non-VRF-aware IPsec)
- FVRF=IVRF=VRF1 (when the encrypted/decrypted packets flow within the same VRF)

https://t.me/learningnets

- FVRF=VRF1 and IVRF=VRF2 (when encrypted packets are in one VRF and decrypted packets are in another VRF)
- FVRF=global VRF and IVRF=VRF1
- FVRF=VRF1 and IVRF=global VRF

With crypto-map configuration, the last option is not supported, whereas with SVTI/GRE, all scenarios are supported. The required configuration steps for crypto-map scenario are as follows:

1. Create the FVRF and/or IVRF.
2. Assign interface to FVRF/IVRF and configure necessary routing.
3. Configure IKEv1 Phase1 parameters (ISAKMP policies).
4. Configure the authentication method (pre-shared or digital certificate) and map it to the FVRF.
5. Configure the proxy-ACL (define interesting traffic).
6. Configure a keyring and map it to the FVRF.
7. Configure an ISAKMP profile and map it to the IVRF.
8. Configure the crypto-map and apply it to the FVRF interface.

With crypto-maps, the entity doing the bridging between the two VRFs is the ISAKMP profile. To populate the IVRF with routing information about the remote network accessible over the IPsec tunnel (required to trigger the IPsec process and for the data-plane to be functional afterward), there are three options:

- Configure Reverse Route Injection (RRI) at crypto-map level using the `reverse-route static` command; the `static` keyword forces the networks from the proxy-ACL to be added to the routing table unconditionally and before the IPsec tunnel is established. This works only when the FVRF is the default global VRF.
- Configure RRI at crypto-map level using the `reverse-route remote-peer <IP_ADDRESS> static` command; the `static` keyword forces the networks from the proxy-ACL to be added to the routing table unconditionally and before the IPsec tunnel is established. This works only when the FVRF is not the global VRF.
- Configure VRF route leaking (inter-VRF routing) using the command `ip route vrf <IVRF_NAME> <SUBNET> <MASK> <FVRF_INTF> <FVRF_NEXT_HOP> [global]` ; the `global` keyword should be used only if the FVRF used is the default global VRF.

Configuration

```
R1:
```

https://t.me/learningnets

```
ip cef
ip vrf IVRF-CLIENT
!
crypto keyring FVRF
 pre-shared-key address 136.1.38.3 key CISCO
!
crypto isakmp policy 10
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp profile IVRF
 vrf IVRF-CLIENT
 keyring FVRF
 match identity address 136.1.38.3 255.255.255.255
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
ip access-list extended LO1_TO_LO3
 permit ip host 150.1.1.1 host 150.1.3.3
!
crypto map VPN 5 ipsec-isakmp
 set peer 136.1.38.3
 set transform-set 3DES_MD5
 set isakmp-profile IVRF
 match address LO1_TO_LO3
 reverse-route static
!
interface GigabitEthernet0/0
 crypto map VPN
!
interface Loopback0
 ip vrf forwarding IVRF-CLIENT
 ip address 150.1.1.1 255.255.255.255
```

R3:

```
ip cef
ip vrf IVRF-CLIENT
!
crypto keyring  FVRF
 pre-shared-key address 136.1.18.1 key CISCO
!
crypto isakmp policy 10
 authentication pre-share
```

https://t.me/learningnets

```
 encryption 3des
 hash md5
 group 2
!
crypto isakmp profile IVRF
 vrf IVRF-CLIENT
 keyring FVRF
 match identity address 136.1.18.1 255.255.255.255
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
ip access-list extended LO3_TO_LO1
 permit ip host 150.1.3.3 host 150.1.1.1
!
crypto map VPN 5 ipsec-isakmp
 set peer 136.1.18.1
 set transform-set 3DES_MD5
 set isakmp-profile IVRF
 match address LO3_TO_LO1
!
interface FastEthernet0/0
 crypto map VPN
!
interface Loopback0
 ip vrf forwarding IVRF-CLIENT
 ip address 150.1.3.3 255.255.255.255
!
ip route vrf IVRF-CLIENT 150.1.1.1 255.255.255.255 FastEthernet0/0 136.1.38.8 global
```

Verification

Note that although both routes in the VRF appear as static, only on R1 is it known from RRI:

```
R1#show ip route vrf IVRF-CLIENT static | b Gateway
Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 2 subnetsS        150.1.3.3 [1/0] via 136.1.38.3
!
!R3#show ip route vrf IVRF-CLIENT static | b Gateway
Gateway of last resort is not set

      150.1.0.0/32 is subnetted, 2 subnetsS        150.1.1.1 [1/0] via 136.1.38.8
!
!R1#show crypto route
```

https://t.me/learningnets

```
VPN Routing Table: Shows RRI and VTI created routes
Codes: RRI - Reverse-Route, VTI- Virtual Tunnel Interface
        S - Static Map ACLs

Routes created in table IVRF-CLIENT150.1.3.3/255.255.255.255 [1/0] via 136.1.38.3 tag 0 count 2 rtid 2
 on GigabitEthernet0/0 RRI  S
!
!R3#show crypto route
 No VPN routes to display

R3#
```

Generate interesting traffic, and verify the IPsec session using specific VRF commands:

```
R1#ping vrf IVRF-CLIENT 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 .!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#show crypto session ivrf IVRF-CLIENT
Crypto session current status

Interface: GigabitEthernet0/0 Profile: IVRF
Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3  Active SAs: 2, origin: crypto map
```

Notice from the following output that the IVRF appears as **protected**, indicating that in this VRF the decrypted data flows:

```
R1#show crypto ipsec sa vrf IVRF-CLIENT

interface: GigabitEthernet0/0
    Crypto map tag: VPN, local addr 136.1.18.1
 protected vrf: IVRF-CLIENT
   local  ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
   current_peer 136.1.38.3 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
```

https://t.me/learningnets

```
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4

    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.1, remote crypto endpt.: 136.1.38.3
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0
     current outbound spi: 0x98079763(2550634339)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0x10AA3AE1(279591649) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2007, flow_id: Onboard VPN:7, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4349056/320)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0x98079763(2550634339) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2008, flow_id: Onboard VPN:8, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4349055/320)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VRF Aware VPN with Crypto Map and Custom FVRF

Task

- Configure an IKEv1 IPsec tunnel between R1 and R3 to provide connectivity between Loopback0 interfaces.
- Use the following parameters:
  - Pre-shared-key of **CISCO**
  - 3DES/MD5 for both phases
  - DH with 1024-bit strength
- On both routers, configure a VRF named **IVRF-CLIENT** and assign Loopback0 interfaces to it.
- Configure a VRF named **FVRF-PROVIDER** and assign VLAN 18 and VLAN 38 interfaces of R1 and R3 to it.

> **A Note On Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is a crypto-map-based, VRF-aware IPsec where the FVRF is non-global, so it is configured at crypto keyring level and `match identity` setting from the ISAKMP profile.

https://t.me/learningnets

Configuration

```
R1:
ip cef
ip vrf IVRF-CLIENT
ip vrf FVRF-PROVIDER
!
crypto keyring FVRF vrf FVRF-PROVIDER
 pre-shared-key address 136.1.38.3 key CISCO
!
crypto isakmp policy 10
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp profile IVRF
 vrf IVRF-CLIENT
 keyring FVRF
 match identity address 136.1.38.3 255.255.255.255 FVRF-PROVIDER
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
ip access-list extended LO1_TO_LO3
 permit ip host 150.1.1.1 host 150.1.3.3
!
crypto map VPN 5 ipsec-isakmp
 set peer 136.1.38.3
 set transform-set 3DES_MD5
 set isakmp-profile IVRF
 match address LO1_TO_LO3
 reverse-route remote-peer 136.1.38.3 static
!
interface GigabitEthernet0/0
 ip vrf forwarding FVRF-PROVIDER
 ip address 136.1.18.1 255.255.255.0
 crypto map VPN
!
interface Loopback0
 ip vrf forwarding IVRF-CLIENT
 ip address 150.1.1.1 255.255.255.255
!
ip route vrf FVRF-PROVIDER  0.0.0.0 0.0.0.0 136.1.18.8
```

https://t.me/learningnets

R3:

```
ip cef
ip vrf IVRF-CLIENT
ip vrf FVRF-PROVIDER
!
crypto keyring FVRF vrf FVRF-PROVIDER
 pre-shared-key address 136.1.18.1 key CISCO
!
crypto isakmp policy 10
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp profile IVRF
 vrf IVRF-CLIENT
 keyring FVRF
 match identity address 136.1.18.1 255.255.255.255 FVRF-PROVIDER
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
ip access-list extended LO3_TO_LO1
 permit ip host 150.1.3.3 host 150.1.1.1
!
crypto map VPN 5 ipsec-isakmp
 set peer 136.1.18.1
 set transform-set 3DES_MD5
 set isakmp-profile IVRF
 match address LO3_TO_LO1
!
interface FastEthernet0/0
 ip vrf forwarding FVRF-PROVIDER
 ip address 136.1.38.3 255.255.255.0
 crypto map VPN
!
interface Loopback0
 ip vrf forwarding IVRF-CLIENT
 ip address 150.1.3.3 255.255.255.255
!
ip route vrf FVRF-PROVIDER 0.0.0.0 0.0.0.0 136.1.38.8
ip route vrf IVRF-CLIENT 0.0.0.0 0.0.0.0 FastEthernet0/0 136.1.38.8
```

https://t.me/learningnets

Verification

Generate interesting traffic to trigger the IPsec process:

```
R3#ping vrf IVRF-CLIENT 150.1.1.1 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 150.1.3.3 .!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Verify the IPsec session at both FVRF/IVRF levels:

```
R1#show crypto session fvrf FVRF-PROVIDER
Crypto session current status

Interface: GigabitEthernet0/0
Profile: IVRF Session status: UP-ACTIVE

Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3 Active SAs: 2, origin: crypto map
!
!
R1#show crypto session ivrf IVRF-CLIENT
Crypto session current status

Interface: GigabitEthernet0/0
Profile: IVRF Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3 Active SAs: 2, origin: crypto map
```

Verify that packets are encrypted/decrypted, and check that both VRFs are bound to this tunnel through the ISAKMP profile:

```
R1#show crypto session detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
```

https://t.me/learningnets

```
X - IKE Extended Authentication, F - IKE Fragmentation

Interface: GigabitEthernet0/0 Profile: IVRF
Uptime: 00:01:29
Session status: UP-ACTIVE     Peer: 136.1.38.3 port 500 fvrf: FVRF-PROVIDER ivrf: IVRF-CLIENT
      Phase1_id: 136.1.38.3
      Desc: (none)
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
         Capabilities:(none) connid:1014 lifetime:23:58:30
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3
       Active SAs: 2, origin: crypto map Inbound:  #pkts dec'ed 4 drop 0 life (KB/Sec) 4150768/3510
Outbound: #pkts enc'ed 4 drop 0 life (KB/Sec) 4150768/3510
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VRF Aware VPN with IPsec Profiles and Global FVRF

Task

- Configure an IKEv1 IPsec tunnel between R2 and R3 to provide connectivity between Loopback0 interfaces.
- Use the following parameters:
  - Pre-shared-key of **CISCO**
  - 3DES/SHA for both phases
  - DH with 1024-bit strength
- On both routers, configure a VRF named **IVRF-TUNNEL** and assign Loopback0 interfaces to it.
- Do not use any `crypto map` -related commands.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The required configuration steps for SVTI/GRE VRF-aware scenario are as follows:

1. Create the FVRF and/or IVRF.
2. Assign interface to FVRF/IVRF and configure necessary routing.
3. Configure IKEv1 Phase1 parameters (ISAKMP policies).

https://t.me/learningnets

4. Configure the authentication method (pre-shared or digital certificate) and map it to the FVRF.
5. Configure a keyring and map it to the FVRF.
6. Configure an ISAKMP profile.
7. Configure the tunnel interface, assign it to the IVRF, and source it from the FVRF.

Important differences if compared with the crypto-map scenario are:

- There is no need for proxy-ACL, because SVTI uses a non-configurable `permit ip any any` proxy-ACL.
- The ISAKMP profile is no longer mapped to the IVRF; this is because the tunnel interface does the bridging between the two VRFs now.
- The tunnel interface is assigned to the IVRF but sourced from the FVRF.
- There is no RRI possible for connectivity over the tunnel; static and dynamic routing are the only options.

Both GRE and SVTI tunnels can be used; for the below configuration, SVTI was selected.

Configuration

Ensure that R3's VLAN 38 and R2's VLAN 18 interfaces are not assigned to any VRF. Both should be in the default global VRF.

```
R2:
ip cef
ip vrf IVRF-TUNNEL
!
crypto keyring FVRF-TUNNEL
 pre-shared-key address 136.1.38.3 key CISCO
!
crypto isakmp policy 20
 authentication pre-share
 encryption 3des
 hash sha
 group 2
!
crypto isakmp profile IVRF-TUNNEL
 keyring FVRF-TUNNEL
 match identity address 136.1.38.3 255.255.255.255
!
crypto ipsec transform-set 3DES_SHA esp-3des esp-sha-hmac
!
```

https://t.me/learningnets

```
crypto ipsec profile IVRF-TUNNEL
 set transform-set 3DES_SHA
 set isakmp-profile IVRF-TUNNEL
!
interface Tunnel23
 ip vrf forwarding IVRF-TUNNEL
 ip address 136.1.23.2 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.38.3
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile IVRF-TUNNEL
!
interface Loopback0
 ip vrf forwarding IVRF-TUNNEL
 ip address 150.1.2.2 255.255.255.255
!
ip route vrf IVRF-TUNNEL 150.1.3.3 255.255.255.255 Tunne23
```

R3:

```
ip cef
ip vrf IVRF-TUNNEL
!
crypto keyring FVRF-TUNNEL
  pre-shared-key address 136.1.18.2 key CISCO
!
crypto isakmp policy 20
 authentication pre-share
 encryption 3des
 hash sha
 group 2
!
crypto isakmp profile IVRF-TUNNEL
  keyring  FVRF-TUNNEL
  match identity address 136.1.18.2 255.255.255.255
!
crypto ipsec transform-set 3DES_SHA  esp-3des esp-sha-hmac
!
crypto ipsec profile IVRF-TUNNEL
 set transform-set 3DES_SHA
 set isakmp-profile  IVRF-TUNNEL
!
interface Tunnel23
 ip vrf forwarding IVRF-TUNNEL
 ip address 136.1.23.3 255.255.255.0
```

https://t.me/learningnets

```
 tunnel source FastEthernet0/0
 tunnel destination 136.1.18.2
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile IVRF-TUNNEL
!
interface Loopback0
 ip vrf forwarding IVRF-TUNNEL
 ip address 150.1.3.3 255.255.255.255
!
ip route vrf IVRF-TUNNEL 150.1.2.2 255.255.255.255 Tunne23
```

Verification

Remember that when using SVTI, the tunnel will be automatically established if the configuration is correct. There is no need to generate interesting traffic. Verify connectivity over the tunnel between Loopback0 interfaces:

```
R2#ping vrf IVRF-TUNNEL 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.2.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R2#show ip route vrf IVRF-TUNNEL | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        136.1.23.0/24 is directly connected, Tunnel23
L        136.1.23.2/32 is directly connected, Tunnel23
      150.1.0.0/32 is subnetted, 2 subnets
C        150.1.2.2 is directly connected, Loopback0S        150.1.3.3 is directly connected, Tunnel23
```

Verify the tunnel state and parameters:

```
R2#show crypto session remote 136.1.38.3 detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation
 Interface: Tunnel23
Profile: IVRF-TUNNEL
```

https://t.me/learningnets

```
Uptime: 10:24:39 Session status: UP-ACTIVE
 Peer: 136.1.38.3 port 500 fvrf: (none) ivrf: IVRF-TUNNEL

      Phase1_id: 136.1.38.3
      Desc: (none)
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.38.3/500 Active
         Capabilities:(none) connid:1001 lifetime:13:35:18
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0
       Active SAs: 2, origin: crypto map Inbound:  #pkts dec'ed 5 drop 0 life (KB/Sec) 4347062/540
 Outbound: #pkts enc'ed 5 drop 0 life (KB/Sec) 4347062/540
!
!R2#show crypto map tag Tunnel23-head-0
Crypto Map IPv4 "Tunnel23-head-0" 65536 ipsec-isakmp
        ISAKMP Profile: IVRF-TUNNEL
        ISAKMP profile IVRF-TUNNEL
        Profile name: IVRF-TUNNEL
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={
                3DES_SHA:  { esp-3des esp-sha-hmac  } ,
        }

Crypto Map IPv4 "Tunnel23-head-0" 65537 ipsec-isakmp
        Map is a PROFILE INSTANCE.
        Peer = 136.1.38.3 ISAKMP Profile: IVRF-TUNNEL
 Extended IP access list
 access-list  permit ip any any
        ISAKMP profile IVRF-TUNNEL
        Current peer: 136.1.38.3
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N Transform sets={
 3DES_SHA:  { esp-3des esp-sha-hmac  } ,

        }
        Always create SAs
        Interfaces using crypto map Tunnel23-head-0:
                Tunnel23
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VRF Aware VPN with IPsec Profiles and Custom FVRF

Task

- Configure an IKEv1 IPsec tunnel between R2 and R3 to provide connectivity between Loopback0 interfaces.
- Use the following parameters:
  - Pre-shared-key of **CISCO**
  - 3DES/SHA for both phases
  - DH with 1024-bit strength
- On both routers, configure a VRF named **IVRF-TUNNEL** and assign Loopback0 interfaces to it.
- Configure a VRF named **FVRF-TUNNEL** and assign VLAN 18 and VLAN 38 interfaces of R2 and R3 to it.
- Do not use any `crypto map` -related commands.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **L2L IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is an SVTI-based, VRF-aware IPsec where the FVRF is non-global, so it is configured at crypto keyring level and `match identity` setting from the ISAKMP profile. Note also that by default the router looks for the specified tunnel source in

https://t.me/learningnets

the default global VRF; because the tunnel source is now in the **FVRF-TUNNEL** VRF, it must be specified using the tunnel-level command `tunnel vrf <VRF_NAME> .`

Configuration

R2:

```
ip cef
ip vrf IVRF-TUNNEL
ip vrf FVRF-TUNNEL
!
crypto keyring FVRF-TUNNEL vrf FVRF-TUNNEL
 pre-shared-key address 136.1.38.3 key CISCO
!
crypto isakmp policy 20
 authentication pre-share
 encryption 3des
 hash sha
 group 2
!
crypto isakmp profile IVRF-TUNNEL
 keyring FVRF-TUNNEL
 match identity address 136.1.38.3 255.255.255.255 FVRF-TUNNEL
!
crypto ipsec transform-set 3DES_SHA esp-3des esp-sha-hmac
!
crypto ipsec profile IVRF-TUNNEL
 set transform-set 3DES_SHA
 set isakmp-profile IVRF-TUNNEL
!
interface Tunnel23
 ip vrf forwarding IVRF-TUNNEL
 ip address 136.1.23.2 255.255.255.0
 tunnel source GigabitEthernet0/0
 tunnel destination 136.1.38.3
 tunnel mode ipsec ipv4
 tunnel vrf FVRF-TUNNEL
 tunnel protection ipsec profile IVRF-TUNNEL
!
interface Loopback0
 ip vrf forwarding IVRF-TUNNEL
 ip address 150.1.2.2 255.255.255.255
!
interface GigabitEthernet0/0
 ip vrf forwarding FVRF-TUNNEL
```

https://t.me/learningnets

```
 ip address 136.1.18.2 255.255.255.0
!
ip route vrf FVRF-TUNNEL 0.0.0.0 0.0.0.0 136.1.18.8
ip route vrf IVRF-TUNNEL 150.1.3.3 255.255.255.255 Tunnel23
```

R3:

```
ip cef
ip vrf IVRF-TUNNEL
ip vrf FVRF-TUNNEL
!
crypto keyring FVRF-TUNNEL vrf FVRF-TUNNEL
  pre-shared-key address 136.1.18.2 key CISCO
!
crypto isakmp policy 20
 authentication pre-share
 encryption 3des
 hash sha
 group 2
!
crypto isakmp profile IVRF-TUNNEL
 keyring  FVRF-TUNNEL
 match identity address 136.1.18.2 255.255.255.255 FVRF-TUNNEL
!
crypto ipsec transform-set 3DES_SHA esp-3des esp-sha-hmac
!
crypto ipsec profile IVRF-TUNNEL
 set transform-set 3DES_SHA
 set isakmp-profile  IVRF-TUNNEL
!
interface Tunnel23
 ip vrf forwarding IVRF-TUNNEL
 ip address 136.1.23.3 255.255.255.0
 tunnel source FastEthernet0/0
 tunnel destination 136.1.18.2
 tunnel mode ipsec ipv4
 tunnel vrf FVRF-TUNNEL
 tunnel protection ipsec profile IVRF-TUNNEL
 !
interface Loopback0
 ip vrf forwarding IVRF-TUNNEL
 ip address 150.1.3.3 255.255.255.255
!
interface FastEthernet0/0
 ip vrf forwarding FVRF-TUNNEL
```

https://t.me/learningnets

```
 ip address 136.1.38.3 255.255.255.0
!
ip route vrf FVRF-TUNNEL 0.0.0.0 0.0.0.0 136.1.38.8
ip route vrf IVRF-TUNNEL 150.1.2.2 255.255.255.255 Tunnel23
```

Verification

Verify connectivity over the tunnel and VRF routing tables:

```
R2#ping vrf IVRF-TUNNEL 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.2.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R2#show ip route vrf IVRF-TUNNEL | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        136.1.23.0/24 is directly connected, Tunnel23
L        136.1.23.2/32 is directly connected, Tunnel23
      150.1.0.0/32 is subnetted, 2 subnets
C        150.1.2.2 is directly connected, Loopback0S        150.1.3.3 is directly connected, Tunnel23
!
!R2#show ip route vrf FVRF-TUNNEL | b Gateway
Gateway of last resort is 136.1.18.8 to network 0.0.0.0
S*    0.0.0.0/0 [1/0] via 136.1.18.8

      136.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        136.1.18.0/24 is directly connected, GigabitEthernet0/0
L        136.1.18.2/32 is directly connected, GigabitEthernet0/0
```

Verify IPsec session state:

```
R2#show crypto session fvrf FVRF-TUNNEL
Crypto session current status

Interface: Tunnel23
Profile: IVRF-TUNNELSession status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 Active SAs: 2, origin: crypto map
!
!R2#show crypto session ivrf IVRF-TUNNEL
```

https://t.me/learningnets

```
Crypto session current status

Interface: Tunnel23
Profile: IVRF-TUNNEL Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.2/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0  Active SAs: 2, origin: crypto map
```

Note that both Phase1 and Phase2 are actually mapped to the IVRF, not to the FVRF, because IVRF describes the traffic protected by the IPsec tunnel:

```
R2#show crypto isakmp sa detail vrf FVRF-TUNNEL
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.

IPv6 Crypto ISAKMP SA
 R2#
!
!R2#show crypto isakmp sa detail vrf IVRF-TUNNEL
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
 1014  136.1.18.2      136.1.38.3 IVRF-T ACTIVE 3des sha    psk  2
  23:57:28
       Engine-id:Conn-id =  SW:14

IPv6 Crypto ISAKMP SA
!
!R2#show crypto ipsec sa vrf FVRF-TUNNEL
 R2#
!
!R2#show crypto ipsec sa vrf IVRF-TUNNEL
```

https://t.me/learningnets

```
interface: Tunnel23
    Crypto map tag: Tunnel23-head-0, local addr 136.1.18.2
 protected vrf: IVRF-TUNNEL
   local  ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
   remote ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
   current_peer 136.1.38.3 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
#pkts decaps: 5, #pkts decrypt: 5, #pkts verify: 5
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.2, remote crypto endpt.: 136.1.38.3
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0
     current outbound spi: 0xF6C2858B(4139943307)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0x17852235(394601013) transform: esp-3des esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2047, flow_id: Onboard VPN:47, sibling_flags 80000040, crypto map: Tunnel23-head-0
        sa timing: remaining key lifetime (k/sec): (4260436/3368)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
 outbound esp sas:
      spi: 0xF6C2858B(4139943307) transform: esp-3des esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2048, flow_id: Onboard VPN:48, sibling_flags 80000040, crypto map: Tunnel23-head-0
        sa timing: remaining key lifetime (k/sec): (4260436/3368)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

Section 7 Introduction

This section utilizes Firewall Services

Lab Topics

The following topics are high-level topics of the CCIE Security Lab Exam Blueprint:

- IKE (V1/V2)
- IPsec LAN-to-LAN (Cisco IOS/ASA)
- Dynamic Multipoint VPN (DMVPN)
- Flex VPN
- Group Encrypted Transport (GET) VPN
- Remote Access VPN
  - Easy VPN Server (Cisco IOS/ASA)
  - VPN Client 5.X
  - Clientless WebVPN
  - AnyConnect VPN
  - Easy VPN Remote
  - SSL VPN Gateway
- VPN High Availability
- QoS for VPN
- VRF-aware VPN
- MacSec
- Digital Certificates (Enrollment and Policy Matching)
- Wireless Access
  - EAP methods
  - WPA/WPA-2
  - WIPS

These topics are described in the following checklist:

https://t.me/learningnets

Technology	Mastered?
Understanding Cryptographic Protocols (ISAKMP, IKEv1 and IKEv2, ESP, Authentication Header, CA)	
IPsec VPN Architecture on Cisco IOS Software and Cisco ASA Security Appliance	
Configuring VPNs Using ISAKMP Profiles	
Configuring VPNs Using IPsec Profiles	
GRE over IPsec Using IPsec Profiles	
Router-to-Router Site-to-Site IPsec Using the Classical Command Set (Using Preshared Keys and Certificates)	
Router-to-Router Site-to-Site IPsec Using the New VTI Command Set (Using Preshared Keys and Certificates)	
Router-to-ASA Site-to-Site IPsec (Using Preshared Keys and Certificates)	
Understanding DMVPN architecture (NHRP, mGRE, IPsec, Routing)	
DMVPN Using NHRP and mGRE (Hub-and-Spoke)	
DMVPN Using NHRP and mGRE (Full-Mesh)	
DMVPN Through Firewalls and NAT Devices	

https://t.me/learningnets

Technology	Mastered?
Understanding GETVPN Architecture (GDOI, Key Server, Group Member, Header Preservation, Policy, Rekey, KEK, TEK, and COOP)	
Implementing GETVPN (Using Preshared Keys and Certificates)	
GETVPN Unicast Rekey	
GETVPN Multicast Rekey	
GETVPN Group Member Authorization List	
GETVPN Key Server Redundancy	
GETVPN Through Firewalls and NAT Devices	
Integrating GET VPN with a DMVPN Solution	
Basic VRF-Aware IPsec	
Enabling the CA (PKI) Server (on the Router and Cisco ASA Security Appliance)	
CA Enrollment Process on a Router Client	
CA Enrollment Process on a Cisco ASA Security Appliance Client	
CA Enrollment Process on a PC Client	

https://t.me/learningnets

Technology	Mastered?
Clientless SSL VPN (Cisco IOS WebVPN) on the Cisco ASA Security Appliance (URLs)	
AnyConnect VPN Client on Cisco IOS Software	
AnyConnect VPN Client on the Cisco ASA Security Appliance	
Remote Access Using a Traditional Cisco VPN Client – on a Cisco IOS Router	
Remote Access Using a Traditional Cisco VPN Client – on a Cisco ASA Security Appliance	
Cisco Easy VPN – Router Server and Router Client (Using DVTI)	
Cisco Easy VPN – Router Server and Router Client (Using Classical Style)	
Cisco Easy VPN – Cisco ASA Server and Router Client	
Cisco Easy VPN Remote Connection Modes (Client, Network, Network+)	
Enabling Extended Authentication (XAUTH) on Cisco IOS Software and the Cisco ASA Security Appliance	
Enabling Split Tunneling on Cisco IOS Software and the Cisco ASA Security Appliance	

https://t.me/learningnets

Technology	Mastered?
Enabling Reverse Route Injection (RRI) on Cisco IOS Software and the Cisco ASA Security Appliance	
Enabling NAT-T on Cisco IOS Software and the Cisco ASA Security Appliance	
High-Availability Stateful Failover for IPsec with Stateful Switchover (SSO) and Hot Standby Router Protocol (HSRP)	
High Availability Using Link Resiliency (with Loopback Interface for Peering)	
High Availability Using HSRP and RRI	
High Availability Using IPsec Backup Peers	
High Availability Using GRE over IPsec (Dynamic Routing)	
Basic QoS Features for VPN Traffic on Cisco IOS Software and the Cisco ASA Security Appliance	
Identifying Injected Errors in Troubleshooting Scenarios (for Site-to-Site, DMVPN, GET VPN, and Cisco Easy VPN)	
Understanding and Interpreting the show crypto Commands	
Understanding and Interpreting the debug crypto Commands	

https://t.me/learningnets

Technology	Mastered?
Anyconnect VPN including DAP support	
MacSec (switch-switch, Host-switch)	
Wireless Security on AP and WLC	
EAP methods	
WPA/WPA-2	
WIPS	

Configurations

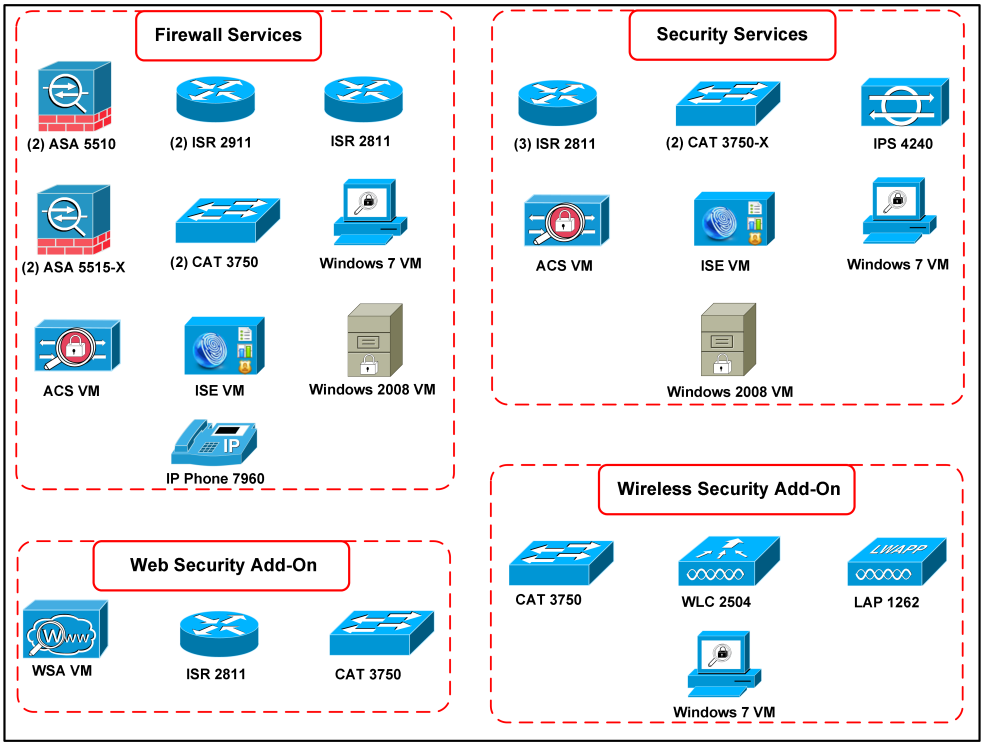
Initial configuration files are provided at the top of this page. Each lab has an initial configuration that is assumed. The initial configuration in some cases is the configuration that the previous lab ended with. If there is only one configuration file, the other devices did not change in the previous lab.

Load the Section 7 Initial Configuration Files to initialize your rack. Pay attention to the instructions in each task, because you may be required to load a different set of initial configuration files.

Section 7 General Lab Topology

The INE CCIE Security Racks are divided into modules as shown in the figure below. This section utilizes Firewall Services.

https://t.me/learningnets



The topology diagram for each task can be found by clicking the **Resources** button.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124
- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VPN High Availability with Backup Peers

Task

- Configure an IKEv1 IPsec tunnel between R1 and R3 to protect traffic between their Loopback0 interfaces:
  - Use 3DES/SHA for both IPsec phases.
  - Use a 1024-bit DH algorithm.
  - Use a pre-shared-key of **CISCO**.
- Configure the tunnel to accomplish the following:
  - R1 should always prefer R3's VLAN38 interface to terminate the tunnel.
  - R3's VLAN 83 interface should be used as backup.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **HA IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.

Overview

This is a classic crypto-map-based configuration using multiple `set peer <IP>` crypto-map-level commands to provide tunnel redundancy. Assuming that you have three or more `set peer` commands, it functions as follows:

- The router tries to negotiate with the first peer. If it doesn't receive a response, it tries the second peer, and so on.
- If the router successfully negotiates with the second peer and detects that this peer failed, by default it tries to negotiate with the next, the third peer.
- The router can be forced to always fall back and try negotiating with a certain peer if the current peer is detected as failed by using the `set peer <IP> default` command; this must be the first peer declared in the crypto-map to be functional.

https://t.me/learningnets

- If the default peer fails, it will try the other peer but continue to try to reach the default at the same time. If at any moment the default peer revives, it drops the current connection and reactivates the primary tunnel.

After IPsec phases are successfully negotiated, they are UP for the configured lifetime, which can be based on seconds for Phase1 and seconds/kilobytes for Phase2. After SAs are up, the router cannot by default detect that the current remote VPN peer failed and it's no longer reachable. Considering this, for a crypto-map entry with multiple `set peer` commands configured to be functional, the router must be configured for Dead Peer Detection (DPD) by using the command `crypto isakmp keepalive <INTERVAL> <RETRY_INTERVAL> [on-demand|periodic]` . Note the following about DPD:

- It is not enabled by default.
- Because it is negotiated in the IKEv1 Phase1 process, it must be configured before the tunnel is UP for it to be functional.
- It is required to detect VPN gateway failure because ESP is not a reliable protocol; as with UDP, there are no acknowledgments.
- **INTERVAL** represents the frequency in seconds for sending DPD messages, assuming that messages are echoed back (remote peer responds).
- **RETRY_INTERVAL** represents frequency in seconds for sending DPD messages when one DPD messages is not echoed back (remote peer does not respond).
- When one DPD message is unanswered, the router sends another four DPD messages at the **RETRY_INTERVAL**, trying to recover the remote peer. If all five DPD messages are unanswered, the peer is declared down, but Phase1 remains UP and Phase2 goes into IDLE state. From the IDLE state, if interesting traffic is detected, the router tries to recover Phase2 once every 10 seconds for five attempts; if the remote peer does not respond, it deletes both Phase1 and Phase2 SAs.
- **On-demand** is the default setting if DPD is configured and tells the router to initiate DPD messages only if needed: If the router sends traffic over the tunnel but no packets are received over the tunnel, it determines that the remote peer may be no longer reachable and initiates the DPD process.
- **Periodic** tells the router to send DPD messages unconditionally at every configured **INTERVAL** value, even if there is traffic flowing bidirectionally over the tunnel or there is no traffic at all. The periodic option consumes more resources but allows for faster dead peer detection. When both VPN gateways are reachable, a message sequence like this should be seen if using a `debug crypto isakmp` 1:

```
ISAKMP:(1022): processing NOTIFY DPD/R_U_THERE_ACK protocol 1
        spi 0, message ID = 4216297052, sa = 0x31480BEC
```

https://t.me/learningnets

```
    ISAKMP:(1022): DPD/R_U_THERE_ACK received from peer 136.1.38.3, sequence 0x5420A026
```

Configuration

ISAKMP profiles are not strictly required but were used in the configuration of this task. DPD periodic mode was configured for faster detection, but on-demand could have been used as well.

R1:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash sha
 group 2
!
crypto keyring R3
 pre-shared-key address 136.1.38.3 255.255.255.255 key CISCO
 pre-shared-key address 136.1.83.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3
 keyring R3
 match identity address 136.1.38.3 255.255.255.255
 match identity address 136.1.83.3 255.255.255.255
!
ip access-list extended LOR1_LOR3
 permit ip host 150.1.1.1 host 150.1.3.3
!
crypto isakmp keepalive 10 3 periodic
crypto ipsec transform-set 3DES_SHA esp-3des esp-sha-hmac
!
crypto map VPN 10 ipsec-isakmp
 set peer 136.1.38.3 default
 set peer 136.1.83.3
 set isakmp-profile R3
 set transform-set 3DES_SHA
 match address LOR1_LOR3
!
interface GigabitEthernet0/0
 crypto map VPN
```

R3:

```
crypto isakmp policy 100
```

https://t.me/learningnets

```
 authentication pre-share
 encryption 3des
 hash sha
 group 2
!
crypto keyring R1
 pre-shared-key address 136.1.18.1 255.255.255.255 key CISCO
!
crypto isakmp profile R1
 keyring R1
 match identity address 136.1.18.1 255.255.255.255
!
ip access-list extended LOR3_LOR1
 permit ip host 150.1.3.3 host 150.1.1.1
!
crypto isakmp keepalive 10 3 periodic
crypto ipsec transform-set 3DES_SHA esp-3des esp-sha-hmac
!
crypto map VPN 10 ipsec-isakmp
 set peer 136.1.18.1
 set isakmp-profile R1
 set transform-set 3DES_SHA
 match address LOR3_LOR1
!
interface FastEthernet0/0.38
 crypto map VPN
!
interface FastEthernet0/0.83
 crypto map VPN
```

Verification

Generate interesting traffic to trigger the IPsec process and verify that both IPsec phases are UP:

```
R1#ping 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
136.1.38.3      136.1.18.1      QM_IDLE           1004 ACTIVE
```

https://t.me/learningnets

```
IPv6 Crypto ISAKMP SA
!
!R1#show crypto session remote 136.1.38.3
Crypto session current status

Interface: GigabitEthernet0/0 Profile: R3
Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Inactive
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3 Active SAs: 2, origin: crypto map
```

Simulate VLAN 38 failure on R3 and issue a **`debug crypto isakmp`** on R1 to observe the unanswered DPD messages and Phase2 becoming IDLE or DOWN depending on IOS version (timesptamps were configured for debug logs to observe the interval for DPD messages). Based on the IOS version running, the actual debug output may differ, but the end result is the same:

```
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#interface fastEthernet0/0.38
R3(config-subif)#shutdown
%DUAL-5-NBRCHANGE: EIGRP-IPv4 111: Neighbor 136.1.38.8 (FastEthernet0/0.38) is down: interface down
R1#debug crypto isakmp
!
!
*May  9 06:13:41.441: ISAKMP: set new node -70244370 to QM_IDLE
*May  9 06:13:41.441: ISAKMP:(1008):Sending NOTIFY DPD/R_U_THERE protocol 1
        spi 746837200, message ID = 4224722926
*May  9 06:13:41.441: ISAKMP:(1008): seq. no 0x542097EC
*May  9 06:13:41.441: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:13:41.441: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:13:41.441: ISAKMP:(1008):purging node -70244370
*May  9 06:13:41.441: ISAKMP:(1008):Input = IKE_MESG_FROM_TIMER, IKE_TIMER_IM_ALIVE
*May  9 06:13:41.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE

*May  9 06:13:42.285: ISAKMP:(1008):purging node 118517802
*May  9 06:13:44.441: ISAKMP:(1008):DPD incrementing error counter (1/5)
*May  9 06:13:44.441: ISAKMP: set new node -83795609 to QM_IDLE
*May  9 06:13:44.441: ISAKMP:(1008):Sending NOTIFY DPD/R_U_THERE protocol 1
        spi 746837200, message ID = 4211171687
*May  9 06:13:44.441: ISAKMP:(1008): seq. no 0x542097ED
*May  9 06:13:44.441: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:13:44.441: ISAKMP:(1008):Sending an IKE IPv4 Packet.
```

https://t.me/learningnets

```
*May  9 06:13:44.441: ISAKMP:(1008):purging node -83795609
*May  9 06:13:44.441: ISAKMP:(1008):Input = IKE_MESG_FROM_TIMER, IKE_TIMER_PEERS_ALIVE
*May  9 06:13:44.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE

*May  9 06:13:47.441: ISAKMP:(1008):DPD incrementing error counter (2/5)
*May  9 06:13:47.441: ISAKMP: set new node 267570671 to QM_IDLE
*May  9 06:13:47.441: ISAKMP:(1008):Sending NOTIFY DPD/R_U_THERE protocol 1
        spi 746837200, message ID = 267570671
*May  9 06:13:47.441: ISAKMP:(1008): seq. no 0x542097EE
*May  9 06:13:47.441: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:13:47.441: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:13:47.441: ISAKMP:(1008):purging node 267570671
*May  9 06:13:47.441: ISAKMP:(1008):Input = IKE_MESG_FROM_TIMER, IKE_TIMER_PEERS_ALIVE
*May  9 06:13:47.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE

*May  9 06:13:50.441: ISAKMP:(1008):DPD incrementing error counter (3/5)
*May  9 06:13:50.441: ISAKMP: set new node 1601369122 to QM_IDLE
*May  9 06:13:50.441: ISAKMP:(1008):Sending NOTIFY DPD/R_U_THERE protocol 1
        spi 746837200, message ID = 1601369122
*May  9 06:13:50.441: ISAKMP:(1008): seq. no 0x542097EF
*May  9 06:13:50.441: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:13:50.441: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:13:50.441: ISAKMP:(1008):purging node 1601369122
*May  9 06:13:50.441: ISAKMP:(1008):Input = IKE_MESG_FROM_TIMER, IKE_TIMER_PEERS_ALIVE
*May  9 06:13:50.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE

*May  9 06:13:51.897: ISAKMP:(1008):purging node 115908208
*May  9 06:13:53.441: ISAKMP:(1008):DPD incrementing error counter (4/5)
*May  9 06:13:53.441: ISAKMP: set new node -1073302555 to QM_IDLE
*May  9 06:13:53.441: ISAKMP:(1008):Sending NOTIFY DPD/R_U_THERE protocol 1
        spi 746837200, message ID = 3221664741
*May  9 06:13:53.441: ISAKMP:(1008): seq. no 0x542097F0
*May  9 06:13:53.441: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:13:53.441: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:13:53.441: ISAKMP:(1008):purging node -1073302555
*May  9 06:13:53.441: ISAKMP:(1008):Input = IKE_MESG_FROM_TIMER, IKE_TIMER_PEERS_ALIVE
*May  9 06:13:53.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE
*May  9 06:13:56.441: ISAKMP:(1008):DPD incrementing error counter (5/5)
*May  9 06:13:56.441: ISAKMP:(1008):peer 136.1.38.3 not responding!
*May  9 06:13:56.441: ISAKMP:(1008):Input = IKE_MESG_FROM_TIMER, IKE_TIMER_PEERS_ALIVE
*May  9 06:13:56.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE

*May  9 06:13:56.441:  ISAKMP: Failed to find peer index node to update peer_info_list
*May  9 06:13:56.441: ISAKMP: set new node 1404063624 to QM_IDLE
*May  9 06:13:56.441: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:13:56.441: ISAKMP:(1008):Sending an IKE IPv4 Packet.
```

https://t.me/learningnets

```
*May  9 06:13:56.441: ISAKMP:(1008):purging node 1404063624
*May  9 06:13:56.441: ISAKMP:(1008):Input = IKE_MESG_FROM_IPSEC, IKE_PHASE2_DEL

*May  9 06:13:56.441: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE
```

While DPD messages are not responded by the remote peer, Phase2 will be in a negotiating state:

```
R1#show crypto session remote 136.1.38.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3 Session status: DOWN-NEGOTIATING
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Inactive
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3  Active SAs: 0, origin: crypto map
```

Verify that Phase2 is in the IDLE or DOWN state, depending on IOS version:

```
R1#show crypto session remote 136.1.38.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3 Session status: UP-IDLE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3  Active SAs: 0, origin: crypto map

!
!R1#show crypto session remote 136.1.38.3
Crypto session current status

Interface: GigabitEthernet0/0 Session status: DOWN

Peer: 136.1.38.3 port 500
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3
        Active SAs: 0, origin: crypto map
```

Generate interesting traffic, keeping the debug active on R1 to observe the router trying five times to recover the Phase2, after which it takes the next peer in the crypto-map and starts the negotiation process:

https://t.me/learningnets

```
R1#ping 150.1.3.3 source loopback0 repeat 50

Type escape sequence to abort.
Sending 50, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 ..............................!!!!!!!!!!!!!!!!!!!!
Success rate is 38 percent (19/50), round-trip min/avg/max = 1/2/4 ms
!
!
*May  9 06:16:23.277: ISAKMP: set new node 0 to QM_IDLE
*May  9 06:16:23.277: SA has outstanding requests  (local 49.76.32.84 port 500, remote 49.76.32.56 port 500)
*May  9 06:16:23.277: ISAKMP:(1008): sitting IDLE. Starting QM immediately (QM_IDLE      )
*May  9 06:16:23.277: ISAKMP:(1008):beginning Quick Mode exchange, M-ID of 488336096
*May  9 06:16:23.277: ISAKMP:(1008):QM Initiator gets spi
*May  9 06:16:23.277: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:16:23.277: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:16:23.277: ISAKMP:(1008):Node 488336096, Input = IKE_MESG_INTERNAL, IKE_INIT_QM
*May  9 06:16:23.277: ISAKMP:(1008):Old State = IKE_QM_READY  New State = IKE_QM_I_QM1
*May  9 06:16:33.277: ISAKMP:(1008): retransmitting phase 2 QM_IDLE       488336096 ...
*May  9 06:16:33.277: ISAKMP (1008): incrementing error counter on node, attempt 1 of 5: retransmit phase 2
*May  9 06:16:33.277: ISAKMP (1008): incrementing error counter on sa, attempt 1 of 5: retransmit phase 2
*May  9 06:16:33.277: ISAKMP:(1008): retransmitting phase 2 488336096 QM_IDLE
*May  9 06:16:33.277: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:16:33.277: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:16:43.277: ISAKMP:(1008): retransmitting phase 2 QM_IDLE       488336096 ...
*May  9 06:16:43.277: ISAKMP (1008): incrementing error counter on node, attempt 2 of 5: retransmit phase 2
*May  9 06:16:43.277: ISAKMP (1008): incrementing error counter on sa, attempt 2 of 5: retransmit phase 2
*May  9 06:16:43.277: ISAKMP:(1008): retransmitting phase 2 488336096 QM_IDLE
*May  9 06:16:43.277: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:16:43.277: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:16:53.273: ISAKMP: set new node 0 to QM_IDLE
*May  9 06:16:53.273: SA has outstanding requests  (local 49.76.32.84 port 500, remote 49.76.32.56 port 500)
*May  9 06:16:53.273: ISAKMP:(1008): sitting IDLE. Starting QM immediately (QM_IDLE      )
*May  9 06:16:53.273: ISAKMP:(1008):beginning Quick Mode exchange, M-ID of 954236601
*May  9 06:16:53.273: ISAKMP:(1008):QM Initiator gets spi
*May  9 06:16:53.273: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:16:53.273: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:16:53.273: ISAKMP:(1008):Node 954236601, Input = IKE_MESG_INTERNAL, IKE_INIT_QM
*May  9 06:16:53.273: ISAKMP:(1008):Old State = IKE_QM_READY  New State = IKE_QM_I_QM1
*May  9 06:16:53.277: ISAKMP:(1008): retransmitting phase 2 QM_IDLE       488336096 ...
*May  9 06:16:53.277: ISAKMP (1008): incrementing error counter on node, attempt 3 of 5: retransmit phase 2
*May  9 06:16:53.277: ISAKMP (1008): incrementing error counter on sa, attempt 3 of 5: retransmit phase 2
*May  9 06:16:53.277: ISAKMP:(1008): retransmitting phase 2 488336096 QM_IDLE
*May  9 06:16:53.277: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:16:53.277: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:17:03.273: ISAKMP:(1008): retransmitting phase 2 QM_IDLE       954236601 ...
*May  9 06:17:03.273: ISAKMP (1008): incrementing error counter on node, attempt 1 of 5: retransmit phase 2
*May  9 06:17:03.273: ISAKMP (1008): incrementing error counter on sa, attempt 4 of 5: retransmit phase 2
```

https://t.me/learningnets

```
*May  9 06:17:03.273: ISAKMP:(1008): retransmitting phase 2 954236601 QM_IDLE
*May  9 06:17:03.273: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:17:03.273: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:17:03.277: ISAKMP:(1008): retransmitting phase 2 QM_IDLE       488336096 ...
*May  9 06:17:03.277: ISAKMP (1008): incrementing error counter on node, attempt 4 of 5: retransmit phase 2
*May  9 06:17:03.277: ISAKMP (1008): incrementing error counter on sa, attempt 5 of 5: retransmit phase 2
*May  9 06:17:03.277: ISAKMP:(1008): retransmitting phase 2 488336096 QM_IDLE
*May  9 06:17:03.277: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:17:03.277: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:17:13.273: ISAKMP:(1008): retransmitting phase 2 QM_IDLE       954236601 ...
*May  9 06:17:13.273: ISAKMP:(1008):peer does not do paranoid keepalives.

*May  9 06:17:13.273: ISAKMP:(1008):deleting SA reason "Death by retransmission P2" state (I) QM_IDLE       (peer 13
*May  9 06:17:13.273: ISAKMP: set new node -557291832 to QM_IDLE
*May  9 06:17:13.273: ISAKMP:(1008): sending packet to 136.1.38.3 my_port 500 peer_port 500 (I) QM_IDLE
*May  9 06:17:13.273: ISAKMP:(1008):Sending an IKE IPv4 Packet.
*May  9 06:17:13.273: ISAKMP:(1008):purging node -557291832
*May  9 06:17:13.273: ISAKMP:(1008):Input = IKE_MESG_INTERNAL, IKE_PHASE1_DEL
*May  9 06:17:13.273: ISAKMP:(1008):Old State = IKE_P1_COMPLETE  New State = IKE_DEST_SA

*May  9 06:17:13.273: ISAKMP:(1008):deleting SA reason "Death by retransmission P2" state (I) QM_IDLE       (peer 13
*May  9 06:17:13.273: ISAKMP: Unlocking peer struct 0x314BBA10 for isadb_mark_sa_deleted(), count 0
*May  9 06:17:13.273: ISAKMP:(1008):deleting node 488336096 error FALSE reason "IKE deleted"
*May  9 06:17:13.273: ISAKMP:(1008):deleting node 954236601 error FALSE reason "IKE deleted"
*May  9 06:17:13.273: ISAKMP:(1008):Input = IKE_MESG_FROM_PEER, IKE_MM_EXCH
*May  9 06:17:13.273: ISAKMP:(1008):Old State = IKE_DEST_SA  New State = IKE_DEST_SA

*May  9 06:17:23.273: ISAKMP:(1008):peer does not do paranoid keepalives.

*May  9 06:17:23.273: ISAKMP:(0): SA request profile is (NULL)
*May  9 06:17:23.273: ISAKMP: Created a peer struct for 136.1.83.3, peer port 500
```

Verify that the tunnel was successfully negotiated with R3's VLAN 83 interface:

```
R1#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
136.1.83.3      136.1.18.1      QM_IDLE           1013 ACTIVE R3

IPv6 Crypto ISAKMP SA
!
!R1#show crypto session remote 136.1.83.3
Crypto session current status
```

https://t.me/learningnets

```
Interface: GigabitEthernet0/0 Profile: R3
Session status: UP-ACTIVE

Peer: 136.1.83.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.83.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3 Active SAs: 2, origin: crypto map
```

Bring R3's VLAN 38 interface back up and verify that the tunnel is automatically renegotiated on VLAN 38:

```
R3#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R3(config)#interface fastEthernet0/0.38
R3(config-subif)#no shutdown
%DUAL-5-NBRCHANGE: EIGRP-IPv4 111: Neighbor 136.1.38.8 (FastEthernet0/0.38) is up: new adjacency
!
!R1#ping 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 ..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/3/4 ms
!
!R1#show crypto session remote 136.1.38.3
Crypto session current status

Interface: GigabitEthernet0/0 Profile: R3
Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.3.3 Active SAs: 2, origin: crypto map

!
!R1#show crypto session remote 136.1.83.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3 Session status: DOWN-NEGOTIATING
Peer: 136.1.83.3 port 500   IKEv1 SA: local 136.1.18.1/500 remote 136.1.83.3/500 Inactive
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VPN High Availability with Loopback Peers

Task

- Configure an IKEv1 IPsec tunnel between R1 and R3 to protect traffic between R1 and SW2 Loopback0 interfaces:
  - Use 3DES/MD5 for both IPsec phases.
  - Use a 1536-bit DH algorithm.
  - Use a pre-shared-key of **CISCO**.
- Configure the tunnel to accomplish the following:
  - Ensure tunnel redundancy.
  - Use a single `set peer` command on R1.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **HA IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is an alternative to configuring multiple crypto-map peers for providing IPsec tunnel redundancy. Because the Loopback0 of the R3 interface is reachable via both VLAN 38 and VLAN 83, the tunnel between R1 and R3 is configured to be initiated/terminated using R3's Loopback0 interface with the command `crypto map <NAME> local-address <INTERFACE>`. With this configuration, redundancy is achieved at routing level (from/to the Loopback0 of R3), not at IPsec level. So the tunnel will remain UP as long as R3 is reachable via any of thw two interfaces (VLAN 38 or

https://t.me/learningnets

VLAN 83).

Configuration

R1:

```
crypto isakmp policy 50
 authentication pre-share
 encryption 3des
 hash md5
 group 5
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
ip access-list extended LOR1_LOSW2
 permit ip host 150.1.1.1 host 150.1.8.8
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto map VPN 5 ipsec-isakmp
 set peer 150.1.3.3
 set isakmp-profile R3_LO
 set transform-set 3DES_MD5
 match address LOR1_LOSW2
!
interface GigabitEthernet0/0
 crypto map VPN
```

R3:

```
crypto isakmp policy 50
 authentication pre-share
 encryption 3des
 hash md5
 group 5
!
crypto keyring R1
 pre-shared-key address 136.1.18.1 255.255.255.255 key CISCO
!
```

https://t.me/learningnets

```
crypto isakmp profile R1
 keyring R1
 match identity address 136.1.18.1 255.255.255.255
!
ip access-list extended LOSW2_LOR1
 permit ip host 150.1.8.8 host 150.1.1.1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto map VPN local-address Loopback0
crypto map VPN 5 ipsec-isakmp
 set peer 136.1.18.1
 set isakmp-profile R1
 set transform-set 3DES_MD5
 match address LOSW2_LOR1
!
interface FastEthernet0/0.38
 crypto map VPN
!
interface FastEthernet0/0.83
 crypto map VPN
```

Verification

Generate interesting traffic and verify tunnel establishment:

```
R1#ping 150.1.8.8 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.8.8, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1001  136.1.18.1      150.1.3.3 ACTIVE 3des md5     psk  5
  23:59:28 D
      Engine-id:Conn-id =  SW:1
```

https://t.me/learningnets

```
IPv6 Crypto ISAKMP SA
!
!R1#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0 Profile: R3_LO
Session status: UP-ACTIVE
Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit ip host 150.1.1.1 host 150.1.8.8  Active SAs: 2, origin: crypto map
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VPN High Availability with Crypto Map and RRI

Task

- Configure an IKEv1 IPsec tunnel between R3 and R1/R2 to protect traffic between SW1 and SW2 Loopback0 interfaces:
  - Use AES-128/MD5 for both IPsec phases.
  - Use an 1024-bit DH algorithm.
  - Use a pre-shared-key of **CISCO**.
- Configure the tunnel to accomplish the following:
  - Ensure tunnel redundancy.
  - Use a single `set peer` command on R1/R2.
  - R3 should prefer R1 to reach SW1's Loopback0.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **HA IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is a crypto-map redundancy scenario that relies on Reverse Route Injection (RRI). As explained in earlier tasks, RRI basically injects static routes in the routing table for remote protected networks negotiated by IPsec (remote networks from local proxy-ACL). In our scenario, because SW1 has two possible exit points (R1 or R2), the VPN gateway that terminates the IPsec tunnel with R3 will redistribute the

https://t.me/learningnets

RRI routes into OSPF. Note that Dead Peer Detection (DPD) was also configured; it is required for a crypto-map entry with multiple `set peer` commands to function.

Configuration

R1:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
ip access-list extended LOSW1_LOSW2
 permit ip host 150.1.7.7 host 150.1.8.8
!
crypto isakmp keepalive 10 3 periodic
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto map VPN 4 ipsec-isakmp
 set peer 150.1.3.3
 set isakmp-profile R3_LO
 set transform-set AES128_MD5
 match address LOSW1_LOSW2
 reverse-route
!
router ospf 111
 redistribute static subnets
!
interface GigabitEthernet0/0
 crypto map VPN
```

R2:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
```

https://t.me/learningnets

```
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
ip access-list extended LOSW1_LOSW2
 permit ip host 150.1.7.7 host 150.1.8.8
!
crypto isakmp keepalive 10 3 periodic
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto map VPN 4 ipsec-isakmp
 set peer 150.1.3.3
 set isakmp-profile R3_LO
 set transform-set AES128_MD5
 match address LOSW1_LOSW2
 reverse-route
!
router ospf 111
 redistribute static subnets
!
interface GigabitEthernet0/0
 crypto map VPN
```

R3:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R1_R2
 pre-shared-key address 136.1.18.1 255.255.255.255 key CISCO
 pre-shared-key address 136.1.18.2 255.255.255.255 key CISCO
!
crypto isakmp profile R1_R2
 keyring R1_R2
 match identity address 136.1.18.1 255.255.255.255
 match identity address 136.1.18.2 255.255.255.255
!
```

https://t.me/learningnets

```
ip access-list extended LOSW2_LOSW1
 permit ip host 150.1.8.8 host 150.1.7.7
!
crypto isakmp keepalive 10 3 periodic
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto map VPN local-address Loopback0
crypto map VPN 4 ipsec-isakmp
 set peer 136.1.18.1 default
 set peer 136.1.18.2
 set isakmp-profile R1_R2
 set transform-set AES128_MD5
 match address LOSW2_LOSW1
!
interface FastEthernet0/0.38
 crypto map VPN
!
interface FastEthernet0/0.83
 crypto map VPN
```

Verification

Verify the routing table on SW1 before the IPsec tunnel is established, generate interesting traffic from SW2, and verify the routing table on SW1 again:

```
SW1#show ip route ospf
 SW1#
!
!SW2#ping 150.1.7.7 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/4/8 ms
!
!SW1#show ip route ospf
     150.1.0.0/32 is subnetted, 2 subnetsO E2    150.1.8.8 [110/20] via 136.1.12.1, 00:01:13, Vlan12
```

Determine whether the required algorithms are used for the tunnel establishment. Note that if configuration from previous tasks was not removed, IPsec Phase2 will negotiate all SAs for all tunnels created between R1 and R3.

https://t.me/learningnets

```
R3#show crypto isakmp sa detail

Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF    Status Encr Hash Auth DH Lifetime Cap.
1003  150.1.3.3       136.1.18.1ACTIVE aes  md5  psk  2
  23:57:39 D
      Engine-id:Conn-id =  SW:3

IPv6 Crypto ISAKMP SA
!
!R3#show crypto session remote 136.1.18.1
Crypto session current status

Interface: FastEthernet0/0.38 FastEthernet0/0.83 Profile: R1_R2
Session status: UP-ACTIVE
Peer: 136.1.18.1 port 500
  IKEv1 SA: local 150.1.3.3/500 remote 136.1.18.1/500 Active
  IPSEC FLOW: permit ip host 150.1.8.8 host 150.1.7.7 Active SAs: 2, origin: crypto map
```

Bring down R1's interface from VLAN 18 and generate interesting traffic again; the tunnel should be negotiated with R2, and R2 should now perform the RRI:

```
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#interface gigabitEthernet0/0
R1(config-if)#shutdown
%DUAL-5-NBRCHANGE: EIGRP-IPv4 111: Neighbor 136.1.18.2 (GigabitEthernet0/0) is down: holding time expired
%DUAL-5-NBRCHANGE: EIGRP-IPv4 111: Neighbor 136.1.18.8 (GigabitEthernet0/0) is down: holding time expired
!
!SW2#ping 150.1.7.7 source loopback0 repeat 10

Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 ..!!!!!!!!
Success rate is 80 percent (8/10), round-trip min/avg/max = 1/4/9 ms
!
!SW1#show ip route ospf
     150.1.0.0/32 is subnetted, 2 subnetsO E2    150.1.8.8 [110/20] via 136.1.12.2, 00:00:06, Vlan12
```

https://t.me/learningnets

```
!
!R2#show crypto route

VPN Routing Table: Shows RRI and VTI created routes
Codes: RRI - Reverse-Route, VTI- Virtual Tunnel Interface
        S - Static Map ACLs

Routes created in table GLOBAL DEFAULT
150.1.8.8/255.255.255.255 [1/0] via 150.1.3.3 tag 0 count 1 rtid 2
 on GigabitEthernet0/0 RRI
```

Verify that the tunnel is established with R2 using required algorithms:

```
R2#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1002  136.1.18.2      150.1.3.3 ACTIVE aes  md5    psk  2
  23:58:32 D
       Engine-id:Conn-id =  SW:2

IPv6 Crypto ISAKMP SA
!
!R2#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0 Profile: R3_LO
Session status: UP-ACTIVE
Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.2/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit ip host 150.1.7.7 host 150.1.8.8  Active SAs: 2, origin: crypto map
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VPN High Availability with GRE over IPsec

Task

- Protect traffic between SW1 and SW2 Loopback0 interfaces, and make use of GRE tunnels between R1-R3 and R2-R3:
  - Use AES-128/MD5 for both IPsec phases.
  - Use a 1024-bit DH algorithm.
  - Use a pre-shared-key of **CISCO**.
- Enable OSPF in area 0 over the tunnels.
- Do not use `crypto map` -related commands.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **HA IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

In the previous task, we used classic crypto-maps with multiple `set peer` commands and RRI, relying on DPD to trigger the failover and accomplish redundancy. This works, but the convergence time is not the best. By using GRE over IPsec or SVTI with dynamic routing for the given topology, we can achieve better convergence time because we rely strictly on routing to trigger the failover. Therefore, if dynamic routing is used, we need only to enhance the routing protocol convergence; for EIGRP/OSPF, this can be done by modifying the default hold-down/dead intervals. The task does not allow use of crypto-maps, so IPsec profiles

https://t.me/learningnets

are used. As discussed in the L2L tasks, for GRE over IPsec, crypto maps are valid solutions, whereas for SVTI, only IPsec profiles are supported.

Configuration

R1:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto ipsec profile R1
 set transform-set AES128_MD5
 set isakmp-profile R3_LO
!
interface Tunnel13
 tunnel source GigabitEthernet0/0
 tunnel destination 150.1.3.3
 tunnel mode gre ip
 ip address 136.1.13.1 255.255.255.0
 tunnel protection ipsec profile R1
!
router ospf 111
 network 136.1.13.0 0.0.0.255 area 0
```

R2:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
```

https://t.me/learningnets

```
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto ipsec profile R1
 set transform-set AES128_MD5
 set isakmp-profile R3_LO
!
interface Tunnel32
 tunnel source GigabitEthernet0/0
 tunnel destination 150.1.3.3
 tunnel mode gre ip
 ip address 136.1.32.2 255.255.255.0
 tunnel protection ipsec profile R1
!
router ospf 111
 network 136.1.32.0 0.0.0.255 area 0
```

R3:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R1_R2
 pre-shared-key address 136.1.18.1 255.255.255.255 key CISCO
 pre-shared-key address 136.1.18.2 255.255.255.255 key CISCO
!
crypto isakmp profile R1_R2
 keyring R1_R2
 match identity address 136.1.18.1 255.255.255.255
 match identity address 136.1.18.2 255.255.255.255
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto ipsec profile R1_R2
 set transform-set AES128_MD5
```

https://t.me/learningnets

```
 set isakmp-profile R1_R2
!
interface Tunnel13
 tunnel source Loopback0
 tunnel destination 136.1.18.1
 tunnel mode gre ip
 ip address 136.1.13.3 255.255.255.0
 tunnel protection ipsec profile R1_R2
!
interface Tunnel32
 tunnel source Loopback0
 tunnel destination 136.1.18.2
 tunnel mode gre ip
 ip address 136.1.32.3 255.255.255.0
 tunnel protection ipsec profile R1_R2
!
router ospf 111
 network 136.1.13.0 0.0.0.255 area 0
 network 136.1.32.0 0.0.0.255 area 0
```

Verification

Verify that SW1 and SW2 learn each other's Loopback0 interfaces through OSPF, and verify OSPF adjacency over the tunnels:

```
R3#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.2.2         0   FULL/  -        00:00:37    136.1.32.2      Tunnel32
150.1.1.1         0   FULL/  -        00:00:37    136.1.13.1      Tunnel13
150.1.8.8         1   FULL/DR         00:00:32    136.1.23.8      FastEthernet0/1
!
!SW2#show ip route ospf | i 150.1.7.7
O       150.1.7.7 [110/1003] via 136.1.23.3, 00:22:15, Vlan23
!
!SW1#show ip route ospf | i 150.1.8.8
O       150.1.8.8 [110/1003] via 136.1.12.2, 00:22:36, Vlan12
!
!SW2#ping 150.1.7.7 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 !!!!!
```

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
```

Verify that the IPsec tunnel is established with both R1 and R2:

```
R3#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
136.1.18.2      150.1.3.3       QM_IDLE           1006 ACTIVE
136.1.18.1      150.1.3.3       QM_IDLE           1007 ACTIVE

IPv6 Crypto ISAKMP SA
!
!R3#show crypto session remote 136.1.18.1
Crypto session current status
 Interface: Tunnel13
Profile: R1_R2
Session status: UP-ACTIVE
Peer: 136.1.18.1 port 500
  IKEv1 SA: local 150.1.3.3/500 remote 136.1.18.1/500 Active
  IPSEC FLOW: permit 47 host 150.1.3.3 host 136.1.18.1 Active SAs: 2, origin: crypto map

!
!R3#show crypto session remote 136.1.18.2
Crypto session current status
 Interface: Tunnel32
Profile: R1_R2
Session status: UP-ACTIVE
Peer: 136.1.18.2 port 500
  IKEv1 SA: local 150.1.3.3/500 remote 136.1.18.2/500 Active
  IPSEC FLOW: permit 47 host 150.1.3.3 host 136.1.18.2 Active SAs: 2, origin: crypto map
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VPN High Availability with Crypto Map and SSO

Task

- Configure an IKEv1 IPsec tunnel between R3 and R1/R2 to protect traffic between SW1 and SW2 Loopback0 interfaces:
  - Use AES-128/MD5 for both IPsec phases.
  - Use a 1024-bit DH algorithm.
  - Use a pre-shared-key of **CISCO**.
- Configure the tunnel to accomplish the following:
  - Ensure tunnel redundancy.
  - All three routers should use a single `set peer` command.
  - R3 should prefer R1 to reach SW1's Loopback0.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **HA IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

Because we are forced to use a single `set peer` command on R3 but have two possible VPN gateways to terminate the tunnel (R1 and R2), we must use the HSRP-based High Availability function of IPsec. Basically we enable HSRP on VLAN 18 between R1 and R2 and configure R3 to terminate the tunnel on the VIP of HSRP; the tunnel will always be established with the ACTIVE HSRP device. There are two deployment options:

- **Stateless IPsec redundancy**. We must configure HSRP and IPsec, but the IPsec Phase1/Phase2 session states will not be replicated between the ACTIVE and STANDBY HSRP routers, so when the ACTIVE role changes, the tunnel must be renegotiated.
- **Stateful IPsec redundancy**. We must configure HSRP, IPsec, and SSO. In this case, IPsec Phase1/Phase2 session states are synchronized between ACTIVE and STANDBY routers, so when the ACTIVE role changes, the tunnel does not need to be renegotiated.

For this task, stateful failover was configured; it is more complex to deploy, although it is not strictly required by the task. Stateful Switch-Over (SSO) inter-device redundancy was briefly explained in earlier tasks; it relies on two protocols, HSRP and SCTP:

- HSRP determines the router roles: ACTIVE and STANDBY.
- SCTP ensures automatic synchronization between ACTIVE and STANDBY; for IPsec, the Phase1/Phase2 session states are replicated, not the crypto configurations.

Because the order of operations may cause your configuration to be non-functional, the following steps are recommended, in order, for configuring SSO with IPsec:

1. Configure and verify HSRP.
2. Configure and verify SSO.
3. Configure IPsec.
4. Configure and verify IPsec for SSO.

From the IPsec configuration perspective, only two things change from a basic crypto-map scenario:

- The `set peer` is using the VIP of the HSRP.
- IPsec redundancy is activated when crypto-map is applied at the interface level using the command `crypto map <NAME> redundancy <HSRP_NAME> [stateful] .`
  - Without the `stateful` keyword, you enable stateless IPsec, so there is no need to configure SSO.

https://t.me/learningnets

- With the `stateful` keyword, you enable stateful IPsec, so you must configure SSO.

Using ISAKMP keepalives and DPD messages is not recommended; rely on the SPI recovery mechanism described in the verification section to recover the tunnel. ISAKMP keepalives and DPD messages may cause the Phase1/Phase2 to be renegotiated and lose the stateful failover functionality.

Configuration

HSRP preemption was configured to allow R1 to always be the preferred peer, as required by the task. Interface tracking was enabled on R1 to ease the verification phase, although in real deployments it is mandatory to automatically move the HSRP ACTIVE role based on LAN/WAN interface availability.

First, configure HSRP and SSO on R1 and R2:

```
R1:
interface GigabitEthernet0/0
 standby ip 136.1.18.12
 standby priority 150
 standby preempt
 standby name IPSEC
 standby track GigabitEthernet0/1 60
!
ipc zone default
 association 1
  protocol sctp
   local-port 5000
   local-ip 136.1.18.1
   remote-port 5000
   remote-ip 136.1.18.2
!
redundancy inter-device
 scheme standby IPSEC


R2:

interface GigabitEthernet0/0
 standby ip 136.1.18.12
 standby preempt
 standby name IPSEC
!
ipc zone default
```

https://t.me/learningnets

```
 association 1
  protocol sctp
   local-port 5000
   local-ip 136.1.18.2
   remote-port 5000
   remote-ip 136.1.18.1
!
redundancy inter-device
 scheme standby IPSEC
```

At this point, as with any other High Availability scenario where HSRP/SSO is used, we must reload the HSRP active and standby devices for SSO to get activated:

```
R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_INIT
Pending Scheme: Standby (Will not take effect until next reload)
 Pending Groupname: IPSEC

  Scheme: <NOT CONFIGURED>
  Peer present: UNKNOWN
  Security: Not configured
```

After reload, if SSO is functional (check both ACTIVE and STANDBY routers), start configuring the IPsec process:

```
R1#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_ACT Scheme: Standby
 Groupname: IPSEC Group State: Active
  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
!R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_STDBY Scheme: Standby
 Groupname: IPSEC Group State: Standby

  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
```

Configure all three routers for IPsec:

```
R1:

crypto isakmp policy 40
 authentication pre-share
 encryption aes
```

https://t.me/learningnets

```
 hash md5
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
ip access-list extended LOSW1_LOSW2
 permit ip host 150.1.7.7 host 150.1.8.8
!
crypto isakmp invalid-spi-recovery
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto map VPN 3 ipsec-isakmp
 set peer 150.1.3.3
 set isakmp-profile R3_LO
 set transform-set AES128_MD5
 match address LOSW1_LOSW2
 reverse-route
!
router ospf 111
 redistribute static subnets
!
interface GigabitEthernet0/0
 crypto map VPN redundancy IPSEC stateful
```

R2:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
ip access-list extended LOSW1_LOSW2
```

https://t.me/learningnets

```
 permit ip host 150.1.7.7 host 150.1.8.8
!
crypto isakmp invalid-spi-recovery
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto map VPN 3 ipsec-isakmp
 set peer 150.1.3.3
 set isakmp-profile R3_LO
 set transform-set AES128_MD5
 match address LOSW1_LOSW2
 reverse-route
!
router ospf 111
 redistribute static subnets
!
interface GigabitEthernet0/0
 crypto map VPN redundancy IPSEC stateful
```

R3:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R1_R2_HSRP
 pre-shared-key address 136.1.18.12 255.255.255.255 key CISCO
!
crypto isakmp profile R1_R2_HSRP
 keyring R1_R2_HSRP
 match identity address 136.1.18.12 255.255.255.255
!
ip access-list extended LOSW2_LOSW1
 permit ip host 150.1.8.8 host 150.1.7.7
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto map VPN local-address Loopback0
crypto map VPN 3 ipsec-isakmp
 set peer 136.1.18.12
 set isakmp-profile R1_R2_HSRP
 set transform-set AES128_MD5
 match address LOSW2_LOSW1
!
```

https://t.me/learningnets

```
interface FastEthernet0/0.38
 crypto map VPN
!
interface FastEthernet0/0.83
 crypto map VPN
```

When stateful IPsec is enabled on R1, as it is the ACTIVE router, it will log the following messages to indicate that it will take over all SAs established on the VIP HSRP address:

```
%CRYPTO-5-IKE_SA_HA_STATUS: IKE sa's if any, for vip  136.1.18.12 will change from STANDBY to ACTIVE
```

Verification

Verify SSO and HSRP:

```
R1#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi0/0       0    150 P Active  local           136.1.18.2      136.1.18.12
!
!R1#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_ACT Scheme: Standby
 Groupname: IPSEC Group State: Active

  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
```

Generate interesting traffic, and verify that crypto-map is configured for redundancy and stateful failover:

```
SW2#ping 150.1.7.7 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/4/9 ms
!
!R1#show crypto map tag VPN
Crypto Map IPv4 "VPN" 3 ipsec-isakmp
        Peer = 150.1.3.3
```

https://t.me/learningnets

```
        ISAKMP Profile: R3_LO
        Extended IP access list LOSW1_LOSW2
            access-list LOSW1_LOSW2 permit ip host 150.1.7.7 host 150.1.8.8 ISAKMP profile R3_LO
 Current peer: 150.1.3.3
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={
                AES128_MD5:  { esp-aes esp-md5-hmac  } ,
        } Reverse Route Injection Enabled
        Interfaces using crypto map VPN:
                GigabitEthernet0/0
 Redundancy Status:
Group: IPSEC,  Type: Stateful HA,  VIP: 136.1.18.12

                Replay-interval: inbound:1000  outbound:100000
```

Verify that both IPsec phases are established on the HSRP ACTIVE router:

```
R1#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1004  136.1.18.12     150.1.3.3              ACTIVE aes  md5    psk  2  23:56:50
       Engine-id:Conn-id =  SW:4

IPv6 Crypto ISAKMP SA
!
!R1#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3_LO Session status: UP-ACTIVE
Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.12/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit ip host 150.1.7.7 host 150.1.8.8  Active SAs: 2, origin: crypto map
```

https://t.me/learningnets

Verify that both IPsec phases are synchronized to the HSRP STANDBY router:

```
R2#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1001  136.1.18.12     150.1.3.3 STDBY
  aes  md5    psk  2  23:57:49
       Engine-id:Conn-id =  SW:1

IPv6 Crypto ISAKMP SA
!
!R2#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3_LO Session status: UP-STANDBY

Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.12/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit ip host 150.1.7.7 host 150.1.8.8
        Active SAs: 2, origin: crypto map
```

Verify that SW1 learned about SW2's Loopback0 interface through R1, which is the ACTIVE VPN gateway:

https://t.me/learningnets

```
SW1#show ip route | i 150.1.8.8
O E2    150.1.8.8 [110/20] via 136.1.12.1, 00:04:59, Vlan12

!
!R1#show crypto route


VPN Routing Table: Shows RRI and VTI created routes
Codes: RRI - Reverse-Route, VTI- Virtual Tunnel Interface
        S - Static Map ACLs

Routes created in table GLOBAL DEFAULT
150.1.8.8/255.255.255.255 [1/0] via 150.1.3.3 tag 0 count 1 rtid 3
 on GigabitEthernet0/0 RRI
```

Simulate a LAN failure on R1 (toward SW1) to force R2 to become the ACTIVE HSRP router. Because R1 will become STANDBY, SSO will force a reload:

```
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#interface gigabitEthernet0/1
R1(config-if)#shutdown
```

Generate traffic from SW2 to verify that protected traffic recovers automatically, without the need for IPsec renegotiations (note that convergence time is mainly given by SW1 detecting that it lost R1 as OSPF neighbor):

```
SW2#ping 150.1.7.7 source loopback0 repeat 50

Type escape sequence to abort.
Sending 50, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 ..........................................!!!!!!!!!!

Success rate is 22 percent (11/50), round-trip min/avg/max = 1/3/9 ms
```

The following messages will appear on R2's console, confirming that it takes over the ACTIVE role and also transitions the synchronized Phase1/Phase2 SAs from STANDBY to ACTIVE:

https://t.me/learningnets

```
%HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 0 state Standby -> Active
%CRYPTO-5-IKE_SA_HA_STATUS: IKE sa's if any, for vip  136.1.18.12 will change from STANDBY to ACTIVE
%CRYPTO-5-IPSEC_SA_HA_STATUS: IPSec sa's if any, for vip  136.1.18.12 will change from STANDBY to ACTIVE
```

Verify that R2 is now the ACTIVE router for the IPsec tunnel:

```
R2#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1001  136.1.18.12     150.1.3.3 ACTIVE
aes  md5    psk  2  23:55:57
       Engine-id:Conn-id =  SW:1

IPv6 Crypto ISAKMP SA
!
!R2#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3_LO Session status: UP-ACTIVE
Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.12/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit ip host 150.1.7.7 host 150.1.8.8  Active SAs: 2, origin: crypto map
!
!R2#show crypto route

VPN Routing Table: Shows RRI and VTI created routes
Codes: RRI - Reverse-Route, VTI- Virtual Tunnel Interface
        S - Static Map ACLs

Routes created in table GLOBAL DEFAULT
150.1.8.8/255.255.255.255 [1/0] via 150.1.3.3 tag 0 count 1 rtid 1
 on GigabitEthernet0/0 RRI
```

https://t.me/learningnets

Note that because HSRP preemption has been configured, after R1 boots up and has its interface toward SW1 functional, it takes over the ACTIVE role, triggering a forced reload on R2, the new STANDBY router. This presents a potential design issue in which the IPsec tunnel may take a very long time to converge (its negotiated lifetime):

- Because HSRP preemption is enabled, when R1 recovers and takes over the ACTIVE role it forces a reload on R2 (the new STANDBY), and Phase1/Phase2 will not have enough time to be synchronized between R1 and R2; R3 has its Phase1/Phase2 UP and will send traffic over the tunnel, but it will be black-holed because R1 does not have Phase1/Phase2 established.
- If traffic is generated from SW2, you will see messages similar to the following on R1: `ISAKMP: ignoring request to send delete notify (no ISAKMP sa) src 136.1.18.12 dst 150.1.3.3 for SPI 0xF0C196A6` . Essentially, R1 says that it received encrypted traffic for an SA/SPI that is not established/negotiated. To fix this problem, configure SPI recovery on R1/R2 by using the command `crypto isakmp invalid-spi-recovery` ; this tells the router to start IKE negotiations whenever it receives encrypted traffic for unknown SA/SPI negotiated.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 VPN High Availability with GRE over IPsec and SSO

Task

- Configure IKEv1 IPsec on R1, R2, and R3 to protect traffic between SW1 and SW2 Loopback0 interfaces using a single GRE tunnel on R3:
  - Use AES-128/MD5 for both IPsec phases.
  - Use a 1024-bit DH algorithm.
  - Use a pre-shared-key of **CISCO**.
- Do not use `crypto map` -related commands.
- Use static routing as needed.
- R3 should prefer R1 to reach SW1's Loopback0.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **HA IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is standard IPsec configuration using IPsec and ISAKMP profiles. To enable High Availability with HSRP for the tunnel, consider the following:

- Tunnel source on the HSRP-enabled routers will no longer use the interface IP address, but the VIP HSRP address.
- Tunnel redundancy is enabled at IPsec profile level by using the command `redundancy <HSRP_NAME> [stateful]`

https://t.me/learningnets

:

- Without the `stateful` keyword, you enable stateless IPsec, so there is no need to configure SSO.
- With the `stateful` keyword, you enable stateful IPsec, so you must configure SSO.

Note that although the documentation says that SVTI is not supported with IPsec stateful failover, it actually is with the IOS code used, 15.1(3)T2. So both GRE and SVTI are supported. IPsec stateful mode was configured, although not strictly required by the task.

Configuration

First, configure HSRP and SSO on R1 and R2:

```
R1:
interface GigabitEthernet0/0
 standby ip 136.1.18.12
 standby priority 150
 standby preempt
 standby name IPSEC
 standby track GigabitEthernet0/1 60
!
ipc zone default
 association 1
  protocol sctp
   local-port 5000
   local-ip 136.1.18.1
   remote-port 5000
   remote-ip 136.1.18.2
!
redundancy inter-device
 scheme standby IPSEC


R2:

interface GigabitEthernet0/0
 standby ip 136.1.18.12
 standby preempt
 standby name IPSEC
!
ipc zone default
 association 1
```

https://t.me/learningnets

```
  protocol sctp
   local-port 5000
   local-ip 136.1.18.2
   remote-port 5000
   remote-ip 136.1.18.1
!
redundancy inter-device
 scheme standby IPSEC
```

At this point, as with any other High Availability scenario where HSRP/SSO is used, we must reload the HSRP active and standby devices for SSO to get activated:

```
R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_INIT
Pending Scheme: Standby (Will not take effect until next reload)
 Pending Groupname: IPSEC

  Scheme: <NOT CONFIGURED>
  Peer present: UNKNOWN
  Security: Not configured
```

After reload, if SSO is functional (check both ACTIVE and STANDBY routers), start configuring the IPsec process:

```
R1#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_ACT Scheme: Standby
 Groupname: IPSEC Group State: Active
  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
!R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_STDBY Scheme: Standby
 Groupname: IPSEC Group State: Standby

  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
```

Configure all three routers for IPsec:

```
R1:
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
```

https://t.me/learningnets

```
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto ipsec profile R1
 set transform-set AES128_MD5
 set isakmp-profile R3_LO
 redundancy IPSEC stateful
!
interface Tunnel123
 tunnel source 136.1.18.12
 tunnel destination 150.1.3.3
 tunnel mode gre ip
 ip address 136.1.123.1 255.255.255.0
 tunnel protection ipsec profile R1
!
ip route 150.1.8.8 255.255.255.255 Tunnel123
!
router ospf 111
 redistribute static subnets metric 10
```

R2:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R3_LO
 pre-shared-key address 150.1.3.3 255.255.255.255 key CISCO
!
crypto isakmp profile R3_LO
 keyring R3_LO
 match identity address 150.1.3.3 255.255.255.255
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
```

https://t.me/learningnets

```
crypto ipsec profile R1
 set transform-set AES128_MD5
 set isakmp-profile R3_LO
 redundancy IPSEC stateful
!
interface Tunnel123
 tunnel source 136.1.18.12
 tunnel destination 150.1.3.3
 tunnel mode gre ip
 ip address 136.1.123.2 255.255.255.0
 tunnel protection ipsec profile R1
!
ip route 150.1.8.8 255.255.255.255 Tunnel123
!
router ospf 111
 redistribute static subnets metric 20
```

R3:

```
crypto isakmp policy 40
 authentication pre-share
 encryption aes
 hash md5
 group 2
!
crypto keyring R1_R2_HSRP
 pre-shared-key address 136.1.18.12 255.255.255.255 key CISCO
!
crypto isakmp profile R1_R2_HSRP
 keyring R1_R2_HSRP
 match identity address 136.1.18.12 255.255.255.255
!
crypto ipsec transform-set AES128_MD5 esp-aes esp-md5-hmac
!
crypto ipsec profile R1_R2_HSRP
 set transform-set AES128_MD5
 set isakmp-profile R1_R2_HSRP
!
interface Tunnel123
 tunnel source Loopback0
 tunnel destination 136.1.18.12
 tunnel mode gre ip
 ip address 136.1.123.3 255.255.255.0
 tunnel protection ipsec profile R1_R2_HSRP
```

https://t.me/learningnets

```
!
ip route 150.1.7.7 255.255.255.255 Tunnel123
!
router ospf 111
 redistribute static subnets
```

Verification

Verify that there is connectivity between SW1 and SW2 Loopback0 interfaces, and between R1 and R3 over the tunnel:

```
SW2#ping 150.1.7.7 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!R3#ping 136.1.123.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.123.1, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Verify that IPsec Phase1/Phase2 are established and synchronized between R1 and R2:

```
R1#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
136.1.18.12     150.1.3.3       QM_IDLE           1002 ACTIVE

IPv6 Crypto ISAKMP SA
!
!R1#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3_LOSession status: UP-ACTIVE
Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.12/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit 47 host 136.1.18.12 host 150.1.3.3 Active SAs: 2, origin: crypto map
```

https://t.me/learningnets

```
!
!R2#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
136.1.18.12     150.1.3.3       QM_IDLE           1002 STDBY

IPv6 Crypto ISAKMP SA
!
!R2#show crypto session remote 150.1.3.3
Crypto session current status

Interface: GigabitEthernet0/0
Profile: R3_LO Session status: UP-STANDBY
Peer: 150.1.3.3 port 500
  IKEv1 SA: local 136.1.18.12/500 remote 150.1.3.3/500 Active
  IPSEC FLOW: permit 47 host 136.1.18.12 host 150.1.3.3  Active SAs: 2, origin: crypto map
```

Further verify that IPsec phases are synchronized between R1 and R2:

```
R1#show crypto ha
IKE VIP: 136.1.18.12 stamp: 16 86 FB 04 4D 15 D3 1F 4D EF 1C 3B E9 98 87 75
IPSec VIP: 136.1.18.12
!
!R2#show crypto ha
IKE VIP: 136.1.18.12 stamp: 16 86 FB 04 4D 15 D3 1F 4D EF 1C 3B E9 98 87 75

IPSec VIP: 136.1.18.12
```

Simulate a LAN failure on R1 (toward SW1) to force R2 become the ACTIVE HSRP router. Because R1 will become STANDBY, SSO will force a reload:

```
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#interface gigabitEthernet0/1
R1(config-if)#shutdown
```

Generate traffic from SW2 to verify that protected traffic recovers automatically, without the need for IPsec re-negotiations (note that convergence time is mainly given by SW1 detecting that it lost R1 as OSPF neighbor):

```
SW2#ping 150.1.7.7 source loopback0 repeat 70

Type escape sequence to abort.
```

https://t.me/learningnets

```
Sending 70, 100-byte ICMP Echos to 150.1.7.7, timeout is 2 seconds:
Packet sent with a source address of 150.1.8.8
.................!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

Success rate is 75 percent (53/70), round-trip min/avg/max = 1/5/9 ms
```

Verify SA state on R2:

```
R2#show crypto isakmp sa
IPv4 Crypto ISAKMP SA
dst             src             state          conn-id status
136.1.18.12     150.1.3.3       QM_IDLE           1005 ACTIVE

IPv6 Crypto ISAKMP SA
!
!R2#show crypto session remote 150.1.3.3 detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation
 Interface: Tunnel123
Profile: R3_LOSession status: UP-ACTIVE
Peer: 150.1.3.3 port 500 fvrf: (none) ivrf: (none)
      Desc: (none)
      Phase1_id: (none)
  IKEv1 SA: local 136.1.18.12/500 remote 150.1.3.3/500 Active
         Capabilities:(none) connid:1005 lifetime:23:53:58
  IPSEC FLOW: permit 47 host 136.1.18.12 host 150.1.3.3
       Active SAs: 2, origin: crypto map Inbound:  #pkts dec'ed 70 drop 0 life (KB/Sec) 3815763/3239
 Outbound: #pkts enc'ed 54 drop 0 life (KB/Sec) 3815765/3239
```

The tunnel will immediately be activated on R2, confirmed by the log messages and tunnel connectivity between R2 and R3:

```
R2#
%HSRP-5-STATECHANGE: GigabitEthernet0/0 Grp 0 state Standby -> Active
%CRYPTO-5-IKE_SA_HA_STATUS: IKE sa's if any, for vip  136.1.18.12 will change from STANDBY to ACTIVE
%CRYPTO-5-IPSEC_SA_HA_STATUS: IPSec sa's if any, for vip  136.1.18.12 will change from STANDBY to ACTIVE
!
!R3#ping 136.1.123.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.123.2, timeout is 2 seconds:!!!!!
```

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Server with Crypto Map and PSK

Task

- Configure R1 as Easy VPN server according to the following requirements:
  - Use 3DES/MD5 as cipher/hash for both IPSec phases.
  - Authenticate remote users identified by the group name **EZVPN** and the password **CISCO**.
  - Use the address pool 20.0.0.1-20.0.0.254 for remote users.
  - Enable Xauth against the local user database, and create a new user named **CISCOUSER1** with the password **CISCO1234**.
  - Only encrypt user traffic destined to R1's Loopback0 subnet.
- Configure **TEST-PC-A **to verify your configuration.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.

Overview

Cisco Easy VPN (EzVPN) has two components: an EzVPN Server and an EzVPN Remote (client). It is basically an IKEv1 IPsec tunnel between two peers (one server and one remote); however, by design it is scoped to minimize the configuration required on the remote side so that most of the tunnel settings are pushed dynamically from the server side. Depending on the authentication method used (pre-shared-key or digital certificates), from the IPsec perspective, little information needs to be configured on the remote side, such as server IP address and group/password or digital certificate. The technical benefits to this include:

- Easing the troubleshooting process on the remote side.
- Allowing for faster change in tunnel policy settings, because these are configured

https://t.me/learningnets

only on the server side.

Easy VPN Server is supported on the following platforms: ISR G1 and ISR G2 routers and ASA firewalls. Easy VPN Remote (Client) is supported on: ISR G1 and ISR G2 routers, ASA 5505 firewall, and Cisco IPsec VPN Client. Cisco extended the IKE/ISAKMP functionality to allow the VPN server to push configuration information to the remote client. A special new phase of ISAKMP negotiations called **Phase 1.5** has been introduced. This phase implements the special ISAKMP (IKE) configuration mode (Mode Config), which allows the initiator and responder to perform a series of transactions after the ISAKMP SA has been established. Generally, the initiator (usually the VPN client) requests certain attribute values from the server (IP address, DNS server, split-tunnel access-list), and the server returns the corresponding values. Additionally, the server may forcefully set some values without any request from the client (for example, configuring proxy settings for Internet Explorer). Notice that ISAKMP configuration mode occurs before the IPSec SA has been negotiated, and this is why it is often called Phase 1.5.

ISAKMP configuration mode allows for Extended Authentication (Xauth) after the ISAKMP SA has been established. This authentication mechanism allows additional validation of end-user identity by using several supported protocols. For example, the server may ask for a username and password and validate them against the local database. The framework is flexible enough to allow for simple password authentication, challenge-response, or one-time passwords. EzVPN requires additional information to be configured for Phase 1.5, and this is accomplished by configuring an EzVPN group with the command `crypto isakmp client configuration group <GROUP>`. When the remote VPN client connects, it must provide a group name by some means so that when the server performs a local database search, it finds the matching group to apply the settings. The group name could be provided by using two common methods:

- **IKE pre-shared-keys authentication**. In this case, the client must initiate Aggressive mode negotiations and present the group name as its IKE ID in the initial packet. This allows the server to find the matching pre-shared-key. EzVPN does not support Main mode with pre-shared-keys authentication. The pre-shared-key is configured under the EzVPN group settings using the command `key <STRING>`.
- **IKE digital signatures authentication and ISAKMP Main Mode.** In this case, the group name is taken from the OU field of the certificate's DN. This is a more secure and scalable approach for ISAKMP authentication.

After the ISAKMP SA has been authenticated using a pre-shared-key or digital signatures, the client may request additional information from the server using ISAKMP configuration mode. The information is extracted from the respective group

https://t.me/learningnets

setting. Here are the most important and commonly used settings:

- `access-restrict <INTERFACE>` : Only accept EzVPN connections terminating on the particular interface.
- `dns <PRIMARY_IP> <SECONDARY_IP>` and `domain <DOMAIN_NAME>` : Specify the list of DNS servers to be pushed to the client and the default domain name.
- `split-dns <DOMAIN_NAME>` : Instructs the VPN client to use the provided DNS servers only for names ending with this domain. All other names are resolved using the remote PC's default DNS servers. This allows using enterprise private DNS servers for accessing VPN resources while allowing Internet access using the client ISP's DNS servers.
- `pool <ADDRESS_POOL_NAME>` : Specifies the name of the address pool configured using the command `ip local pool <NAME> <SATRT_IP> <END_IP> .` This address pool is used for allocating IP addresses to remote clients. This address pool must be referenced in IKE using the command `crypto isakmp client configuration address-pool local <Address-Pool-Name> .`
- `acl [<100-199>|<EXTEDNED_ACL_NAME>]` : Set the split-tunnel access-list for the EzVPN client. Every entry in the split-tunnel access list has the format `[permit|deny] <subnet> <mask> [any|<remote-address-pool-subnet> <remote-address-pool-mask] .` If the entry permits traffic, the client will route traffic to the configured subnet across the VPN tunnel and encrypt it. If the entry denies traffic, the client will not tunnel the traffic to the specified subnet. With no split-tunnel access-list configured, the client will tunnel all traffic across the VPN tunnel.
- `save-password` : Allows the remote client to save the password prompted for Xauth. If this is not configured, the remote client will not save the password locally. If configured, it applies only after the second attempted connection because on first connection it only gets activated.
- `include-local-lan` : Useful when you don't use split-tunneling, this feature allows the remote PC to access the network on the local LAN subnet, bypassing the VPN tunnel. This is helpful if you have a local network printer connected to the LAN segment or any other local resources needed when your VPN connection is up.
- `backup-gateway <IP_ADDRESS>` : After the first connection is successfully established, it pushes a backup EzVPN server IP address to the remote client to be used if the current one is non-responsive. This setting can also be configured manually on the remote client, but it is easier to manage if pushed from the EzVPN server.

When you're done configuring ezVPN group settings, you must configure other regular IPSec VPN settings, such as ISAKMP policy and IPSec transform set. Remember that IOS routers use DH group 1 by default, and EzVPN client requires

https://t.me/learningnets

DH group 2 for 3DES and DH group 5 for AES-256. Also, you must create a dynamic crypto-map to accept connections from any remote peer. When configuring a regular crypto-map, you need the following statements:

- `crypto map VPN client configuration address respond` : Instructs the router to respond to ISAKMP Mode Config messages (specifically, the IP address requests).
- `crypto map VPN isakmp authorization list <AAA_LIST_NAME>` : Required to allow the router to authorize ISAKMP Mode Config attributes using the configured AAA list (local or from a remote RADIUS server).
- `crypto map VPN client authentication list <AAA_LIST_NAME>` : Enables XAUTH for the crypto-map. The AAA list should reference the source of authentication information: local user database or a remote server. If you don't provide this command, XAUTH is not used and the remote session is authenticated solely based on ISAKMP authentication parameters (pre-shared-key or digital certificates).

The remaining question is how will the router (EzVPN server) route traffic to the remote endpoint (EzVPN client) assigned IP addresses. If the router has a default route pointing to the interface where the remote sessions are terminated, it will implicitly work, because it will match on the crypto-map and get encrypted. However, if default routing is not in use or points to a different interface, the router needs a mechanism to route for the allocated IP addresses of remote client. This can be achieved by using the procedure known as Reverse Route Injection (RRI). For every new IP address allocated via ISAKMP configuration mode, the router installs a local static /32 route that may also be redistributed into an IGP to let other network devices learn about it. To enable RRI, use the command `reverse-route` under dynamic crypto-map configuration mode. All RRI options have already been addressed in the L2L tasks.

Cisco VPN Client has DPD in periodic mode enabled by default, and this is not configurable. Note that for DPD to be active for an IPsec session, both peers must support it, but it is sufficient for only one peer to have it enabled and start negotiation. DPD is negotiated in the IPsec tunnel establishment phase.

Configuration

R1:

```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCOUSER1 password CISCO1234
```

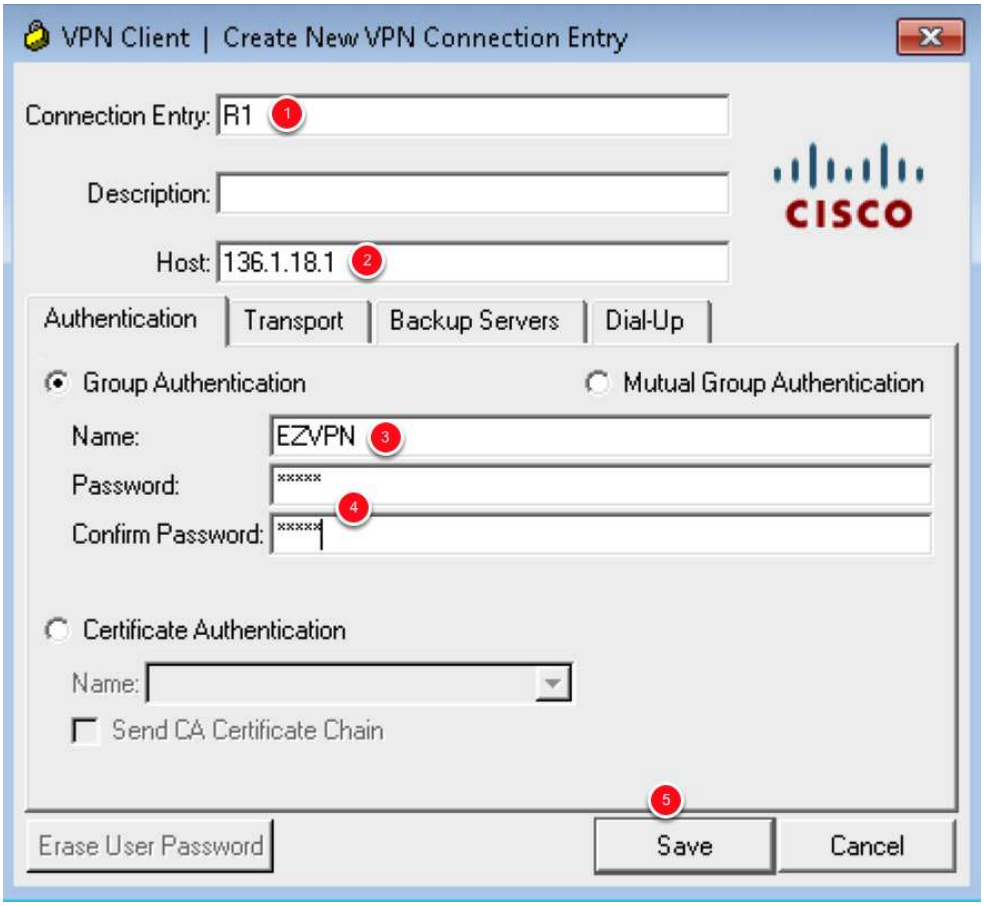
https://t.me/learningnets

```
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration  address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto dynamic-map EZVPN_DYNAMIC 100
 set transform-set 3DES_MD5
 reverse-route
!
crypto map VPN isakmp authorization list EZVPN
crypto map VPN client authentication list EZVPN
crypto map VPN client configuration address respond
crypto map VPN 100 ipsec-isakmp dynamic EZVPN_DYNAMIC
!
interface GigabitEthernet0/0
 crypto map VPN
```

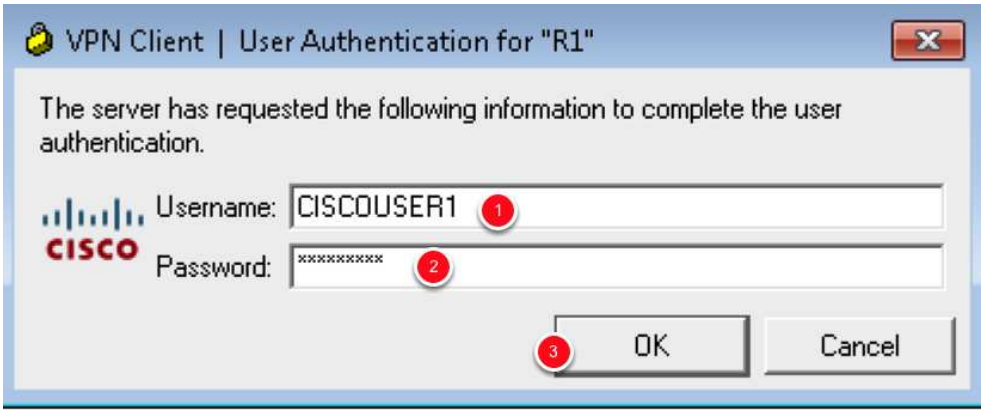
Verification

TEST-PC-A will be automatically assigned an IP address from the DHCP server configured on R3. Configure a new Cisco VPN Client connection entry by clicking the **New** button:

https://t.me/learningnets

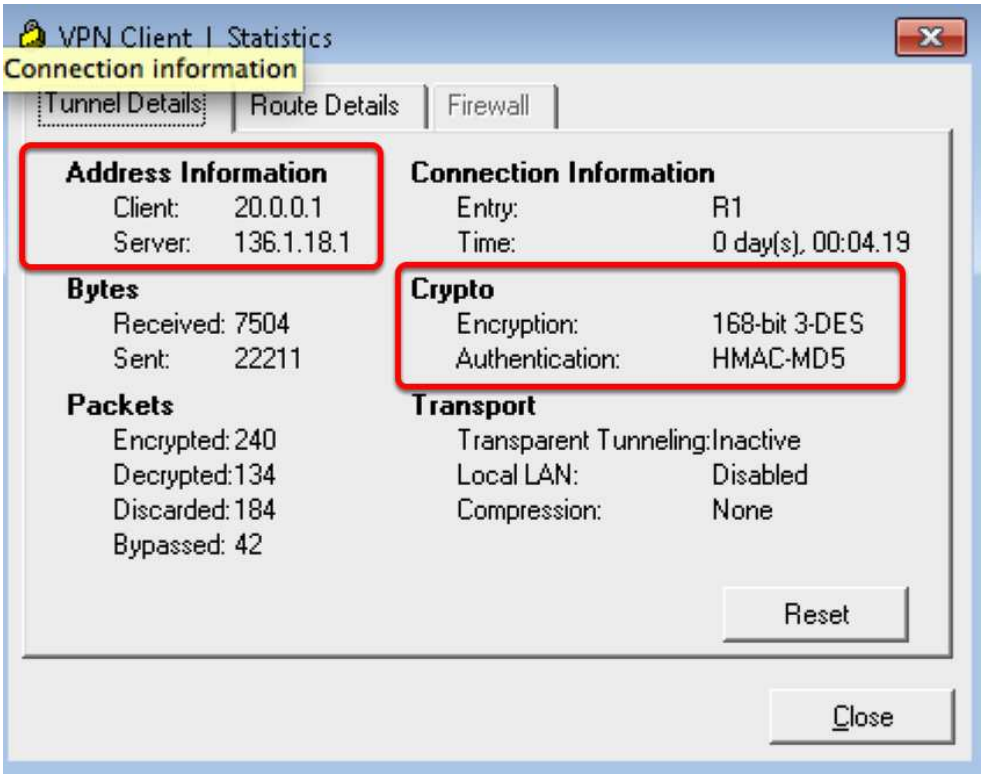


Initiate the connection by clicking the **Connect** button. An authentication window prompts you for a username and password:

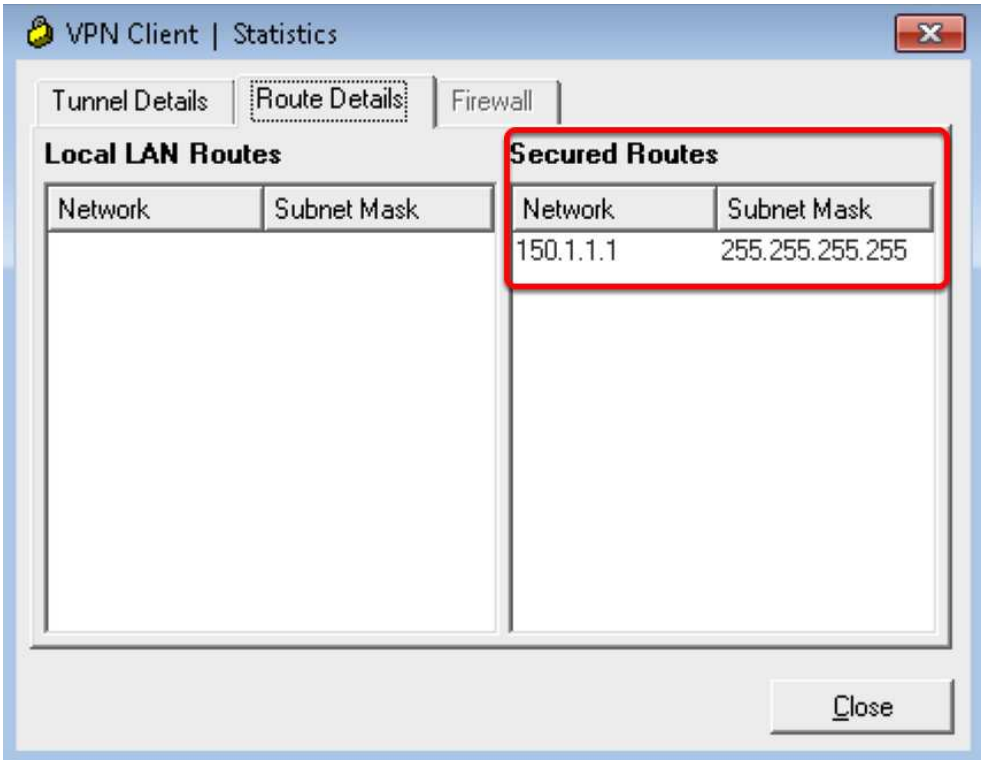


After the connection is established, right-click the VPN Client icon on the toolbar and open the **Statistics** window to verify negotiated algorithms, assigned IP address, and traffic statistics:

https://t.me/learningnets



On the **Route Details** tab, verify for which destinations traffic is forwarded through the tunnel; this is given by the split-tunneling access-list configured on the EzVPN server:



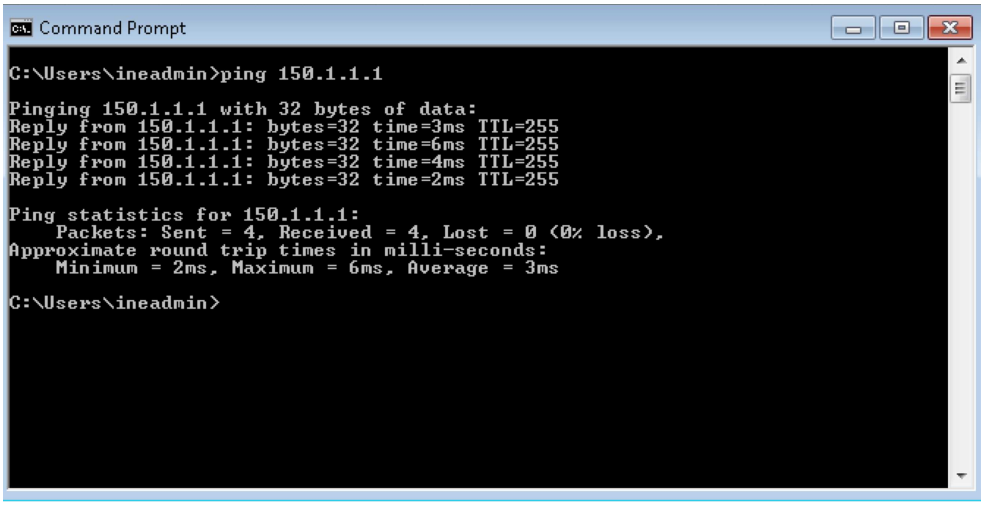
Verify that traffic from **TEST-PC-A** to R1's Loopback0 subnet is sent over the tunnel:

```
R1#show crypto ipsec  sa peer 136.1.100.100 | i pkts
```

https://t.me/learningnets

```
#pkts encaps: 0, #pkts encrypt: 0, #pkts digest: 0
 #pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0

    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
```




```
R1#show crypto ipsec  sa peer 136.1.100.100 | i pkts
 #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
 #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4

    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
```

Verify IPsec Phase1 and Phase2:

```
R1#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF  Status Encr Hash   Auth DH Lifetime Cap.
1004  136.1.18.1      136.1.100.100          ACTIVE 3des md5         2  23:38:43 CX
       Engine-id:Conn-id =  SW:4

IPv6 Crypto ISAKMP SA
!
!R1#show crypto session username CISCOUSER1
```

https://t.me/learningnets

```
Crypto session current status

Interface: GigabitEthernet0/0Username: CISCOUSER1
Group: EZVPN
Assigned address: 20.0.0.1
Session status: UP-ACTIVE

Peer: 136.1.100.100 port 58438
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.100.100/58438 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 host 20.0.0.4
        Active SAs: 2, origin: dynamic crypto map
```

The following additional commands are useful for verifying which users are connected to which EzVPN groups:

```
R1#show crypto isakmp peers
Peer: 136.1.100.100 Port: 58438 Local: 136.1.18.1 Phase1 id: EZVPN
!
!R1#show crypto isakmp peers config
 Client-Public-Addr=136.1.100.100:58438; Client-Assigned-Addr=20.0.0.1;
Client-OS=Cisco Systems VPN Client 5.0.07.0410:WinNT; Client-Group=EZVPN; Client-User=CISCOUSER1;
```

Before establishing the VPN connection, enable the following two debugs on R1 to identify the IPsec tunnel negotiation phases:

```
R1#debug crypto isakmp
Crypto ISAKMP debugging is on
!R1#debug crypto ipsec
Crypto IPSEC debugging is on
```

At this point, the EzVPN server receives the initial IKEv1 packet from the client, which is the first packet in IKEv1 Aggressive mode exchange. This packet contains the IKE ID and the SA payload in addition to the other ISAKMP payloads. The IKE ID contains the EzVPN group name, which the server will use to find the locally specified pre-shared-key (note that no ISAKMP profiles are being matched):

```
ISAKMP (0): received packet from 136.1.100.100 dport 500 sport 58520 Global (N) NEW SA
ISAKMP: Created a peer struct for 136.1.100.100, peer port 58520
ISAKMP: New peer created peer = 0x2CC17F98 peer_handle = 0x80000008
ISAKMP: Locking peer struct 0x2CC17F98, refcount 1 for crypto_isakmp_process_block
ISAKMP:(0):Setting client config settings 310871CC
ISAKMP:(0):(Re)Setting client xauth list  and state
```

https://t.me/learningnets

```
ISAKMP/xauth: initializing AAA request
ISAKMP: local port 500, remote port 58520
ISAKMP:(0):insert sa successfully sa = 319A0D48
ISAKMP:(0): processing SA payload. message ID = 0
ISAKMP:(0): processing ID payload. message ID = 0
ISAKMP (0): ID payload
        next-payload : 13
        type         : 11  group id      : EZVPN
        protocol     : 17
        port         : 500
        length       : 13ISAKMP:(0):: peer matches *none* of the profiles

ISAKMP:(0): processing vendor id payload
ISAKMP:(0): vendor ID seems Unity/DPD but major 215 mismatch
ISAKMP:(0): vendor ID is XAUTH
ISAKMP:(0): processing vendor id payload
ISAKMP:(0): vendor ID is DPD
ISAKMP:(0): processing vendor id payload
ISAKMP:(0): processing IKE frag vendor id payload
ISAKMP:(0):Support for IKE Fragmentation not enabled
ISAKMP:(0): processing vendor id payload
ISAKMP:(0): vendor ID seems Unity/DPD but major 123 mismatch
ISAKMP:(0): vendor ID is NAT-T v2
ISAKMP:(0): processing vendor id payload
ISAKMP:(0): vendor ID is Unity
```

The EzVPN server tries to find a match in the list of locally configured ISAKMP policies for the proposal offered by the client, and finally finds one:

```
ISAKMP:(0): Authentication by xauth preshared
Rack1R1#0):Checking ISAKMP transform 1 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash SHA
ISAKMP:      default group 2
ISAKMP:      auth XAUTHInitPreShared
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 256ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 2 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash MD5
ISAKMP:      default group 2
ISAKMP:      auth XAUTHInitPreShared
ISAKMP:      life type in seconds
```

https://t.me/learningnets

```
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 256ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 3 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash SHA
ISAKMP:      default group 2
ISAKMP:      auth pre-share
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 256ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 4 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash MD5
ISAKMP:      default group 2
ISAKMP:      auth pre-share
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 256ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 5 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash SHA
ISAKMP:      default group 2
ISAKMP:      auth XAUTHInitPreShared
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 128
ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 6 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash MD5
ISAKMP:      default group 2
ISAKMP:      auth XAUTHInitPreShared
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 128ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 7 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash SHA
ISAKMP:      default group 2
ISAKMP:      auth pre-share
ISAKMP:      life type in seconds
```

https://t.me/learningnets

```
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 128ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 8 against priority 100 policy
ISAKMP:      encryption AES-CBC
ISAKMP:      hash MD5
ISAKMP:      default group 2
ISAKMP:      auth pre-share
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:      keylength of 128ISAKMP:(0):Encryption algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 9 against priority 100 policy
ISAKMP:      encryption 3DES-CBC
ISAKMP:      hash SHA
ISAKMP:      default group 2
ISAKMP:      auth XAUTHInitPreShared
ISAKMP      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:(0):Hash algorithm offered does not match policy!
ISAKMP:(0):atts are not acceptable. Next payload is 3
ISAKMP:(0):Checking ISAKMP transform 10 against priority 100 policyISAKMP:      encryption 3DES-CBC
ISAKMP:      hash MD5
ISAKMP:      default group 2
ISAKMP:      auth XAUTHInitPreShared
ISAKMP:      life type in seconds
ISAKMP:      life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:(0):atts are acceptable. Next payload is 3

ISAKMP:(0):Acceptable atts:actual life: 86400
ISAKMP:(0):Acceptable atts:life: 0
ISAKMP:(0):Fill atts in sa vpi_length:4
ISAKMP:(0):Fill atts in sa life_in_seconds:2147483
ISAKMP:(0):Returning Actual lifetime: 86400
ISAKMP:(0)::Started lifetime timer: 86400.
```

Additionally, the initial IKEv1 packet contains the key exchange payload for DH negotiations. After this, the EzVPN server constructs a response packet to the client. The IKE ID used by the server is the server's IPv4 address. This is the second packet in IKE AM exchange; it contains the response key exchange payload as well. The response packet uses the **SKEY_IDs** generated based on the pre-shared-key for the group and DH key. Additionally, the response packet contains the authentication information based on the pre-shared-key:

https://t.me/learningnets

```
ISAKMP:(0): processing KE payload. message ID = 0
ISAKMP:(0): processing NONCE payload. message ID = 0
ISAKMP:(0): vendor ID is NAT-T v2 ISAKMP:(0):Input = IKE_MESG_FROM_PEER, IKE_AM_EXCH
ISAKMP:(0):Old State = IKE_READY  New State = IKE_R_AM_AAA_AWAIT

ISAKMP:(1007): constructed NAT-T vendor-02 ID
 ISAKMP:(1007):SA is doing pre-shared key authentication plus XAUTH using id type ID_IPV4_ADDR
ISAKMP (1007): ID payload
        next-payload : 10
        type         : 1
        address      : 136.1.18.1
        protocol     : 0
        port         : 0
        length       : 12
ISAKMP:(1007):Total payload length: 12
ISAKMP:(1007): sending packet to 136.1.100.100 my_port 500 peer_port 58520 (R) AG_INIT_EXCH
ISAKMP:(1007):Sending an IKE IPv4 Packet. ISAKMP:(1007):Input = IKE_MESG_FROM_AAA, PRESHARED_KEY_REPLY

ISAKMP:(1007):Old State = IKE_R_AM_AAA_AWAIT  New State = IKE_R_AM2
```

The server receives the third packet from the client. This is the last packet in IKE AM exchange. It contains the authentication information (HASH payload), and now the server assumes the SA to be authenticated:

```
ISAKMP (1007): received packet from 136.1.100.100 dport 500 sport 58520 Global (R) AG_INIT_EXCH
ISAKMP:(1007): processing HASH payload. message ID = 0
ISAKMP:(1007): processing NOTIFY INITIAL_CONTACT protocol 1
        spi 0, message ID = 0, sa = 0x319A0D48
ISAKMP:received payload type 20
ISAKMP (1007): His hash no match - this node outside NAT
ISAKMP:received payload type 20
ISAKMP (1007): No NAT Found for self or peer
ISAKMP:(1007):SA authentication status:
        authenticated
ISAKMP:(1007):SA has been authenticated with 136.1.100.100 ISAKMP:(1007):SA authentication status:
 authenticated
```

The client sends special Initial Contact payload information in the final packet. This information instructs the server to bring down any previous SAs associated with this client. This step is needed to ensure that no old SAs are lingering because of disrupted communications:

https://t.me/learningnets

```
ISAKMP:(1007): Process initial contact,

bring down existing phase 1 and 2 SA's with local 136.1.18.1 remote 136.1.100.100 remote port 58520
ISAKMP:(1007):returning IP addr to the address pool
ISAKMP: Trying to insert a peer 136.1.18.1/136.1.100.100/58520/,  and inserted successfully 2CC17F98.
ISAKMP:(1007):Returning Actual lifetime: 86400
ISAKMP: set new node 319025020 to CONF_XAUTH
ISAKMP:(1007):Sending NOTIFY RESPONDER_LIFETIME protocol 1
        spi 824065064, message ID = 319025020
ISAKMP:(1007): sending packet to 136.1.100.100 my_port 500 peer_port 58520 (R) QM_IDLE
ISAKMP:(1007):Sending an IKE IPv4 Packet.
ISAKMP:(1007):purging node 319025020
ISAKMP: Sending phase 1 responder lifetime 86400

ISAKMP:(1007):Input = IKE_MESG_FROM_PEER, IKE_AM_EXCH
ISAKMP:(1007):Old State = IKE_R_AM2  New State = IKE_P1_COMPLETE

IPSEC(key_engine): got a queue event with 1 KMI message(s)ISAKMP:(1007):Need XAUTH
ISAKMP: set new node -163585383 to CONF_XAUTH
```

The server now starts ISAKMP configuration mode. The first configuration procedure is XAUTH. The server requests **USERNAME** and **PASSWORD** attributes from the client:

```
ISAKMP/xauth: request attribute XAUTH_USER_NAME_V2
ISAKMP/xauth: request attribute XAUTH_USER_PASSWORD_V2
ISAKMP:(1007): initiating peer config to 136.1.100.100. ID = 4131381913
ISAKMP:(1007): sending packet to 136.1.100.100 my_port 500 peer_port 58520 (R) CONF_XAUTH
ISAKMP:(1007):Sending an IKE IPv4 Packet.
ISAKMP:(1007):Input = IKE_MESG_INTERNAL, IKE_PHASE1_COMPLETE
ISAKMP:(1007):Old State = IKE_P1_COMPLETE  New State = IKE_XAUTH_REQ_SENT
```

The client responds with a username/password pair:

https://t.me/learningnets

```
ISAKMP (1007): received packet from 136.1.100.100 dport 500 sport 58520 Global (R) CONF_XAUTH
ISAKMP:(1007):processing transaction payload from 136.1.100.100. message ID = -163585383
ISAKMP: Config payload REPLY ISAKMP/xauth: reply attribute XAUTH_USER_NAME_V2
ISAKMP/xauth: reply attribute XAUTH_USER_PASSWORD_V2

ISAKMP:(1007):deleting node -163585383 error FALSE reason "Done with xauth request/reply exchange"
ISAKMP:(1007):Input = IKE_MESG_FROM_PEER, IKE_CFG_REPLY

ISAKMP:(1007):Old State = IKE_XAUTH_REQ_SENT  New State = IKE_XAUTH_AAA_CONT_LOGIN_AWAIT
```

The server responds with the XAUTH OK packet to the client, confirming successful authentication. The client responds with the XAUTH ACK packet to this final message. Essentially, IKEv1 Phase1 is now established (Phase 1.5, which involves the XAUTH, is actually an extension of Phase1):

```
ISAKMP: set new node -1371087872 to CONF_XAUTH
ISAKMP:(1007): initiating peer config to 136.1.100.100. ID = 2923879424
ISAKMP:(1007): sending packet to 136.1.100.100 my_port 500 peer_port 58520 (R) CONF_XAUTH
ISAKMP:(1007):Sending an IKE IPv4 Packet.
ISAKMP:(1007):Input = IKE_MESG_FROM_AAA, IKE_AAA_CONT_LOGIN
ISAKMP:(1007):Old State = IKE_XAUTH_AAA_CONT_LOGIN_AWAIT  New State = IKE_XAUTH_SET_SENT

ISAKMP (1007): received packet from 136.1.100.100 dport 500 sport 58520 Global (R) CONF_XAUTH
ISAKMP:(1007):processing transaction payload from 136.1.100.100. message ID = -1371087872
ISAKMP: Config payload ACK ISAKMP:(1007):        (blank) XAUTH ACK Processed
ISAKMP:(1007):deleting node -1371087872 error FALSE reason "Transaction mode done"
ISAKMP:(1007):Talking to a Unity Client
ISAKMP:(1007):Input = IKE_MESG_FROM_PEER, IKE_CFG_ACK
ISAKMP:(1007):Old State = IKE_XAUTH_SET_SENT  New State = IKE_P1_COMPLETE

ISAKMP:(1007):Input = IKE_MESG_INTERNAL, IKE_PHASE1_COMPLETE
ISAKMP:(1007):Old State = IKE_P1_COMPLETE  New State = IKE_P1_COMPLETE
```

XAUTH is finished, and the client requests several configuration attributes using ISAKMP configuration mode transactions. The server uses AAA subsystem to extract the configuration information from the locally configured EzVPN group:

```
ISAKMP (1007): received packet from 136.1.100.100 dport 500 sport 58520 Global (R) QM_IDLE
ISAKMP: set new node 583863136 to QM_IDLE
ISAKMP:(1007):processing transaction payload from 136.1.100.100. message ID = 583863136
ISAKMP: Config payload REQUEST
```

https://t.me/learningnets

```
ISAKMP:(1007):checking request:
ISAKMP:    IP4_ADDRESS
ISAKMP:    IP4_NETMASK
ISAKMP:    IP4_DNS
ISAKMP:    IP4_NBNS
ISAKMP:    ADDRESS_EXPIRY
ISAKMP:    MODECFG_BANNER
ISAKMP:    MODECFG_SAVEPWD
ISAKMP:    DEFAULT_DOMAIN
ISAKMP:    SPLIT_INCLUDE
ISAKMP:    SPLIT_DNS
ISAKMP:    PFS
ISAKMP:    MODECFG_BROWSER_PROXY
ISAKMP:    BACKUP_SERVER
ISAKMP:    MODECFG_SMARTCARD_REMOVAL_DISCONNECT
ISAKMP:    APPLICATION_VERSION
ISAKMP: Client Version is : Cisco Systems VPN Client 5.0.07.0410:WinNTp
ISAKMP:    FW_RECORD
ISAKMP:    MODECFG_HOSTNAME
ISAKMP:    CONFIG_MODE_UNKNOWN Unknown Attr: 0x7005
ISAKMP/author: Author request for group EZVPNsuccessfully sent to AAA

ISAKMP:(1007):Input = IKE_MESG_FROM_PEER, IKE_CFG_REQUEST
ISAKMP:(1007):Old State = IKE_P1_COMPLETE  New State = IKE_CONFIG_AUTHOR_AAA_AWAIT
```

At this point, configured group attributes are pushed to the client. Most important and mandatory are the VPN IP address and the split-tunnel access list. Some other optional attributes are also mentioned in the debug output:

```
ISAKMP:(1007):attributes sent in message:
        Address: 0.2.0.0ISAKMP:(1007):allocating address 20.0.0.7
ISAKMP: Sending private address: 20.0.0.7
ISAKMP: Sending ADDRESS_EXPIRY seconds left to use the address: 86397
ISAKMP: Sending save password reply value 0
ISAKMP: Sending split include name EZVPN_TUNNEL network 150.1.1.1 mask 255.255.255.255 protocol 0, src port 0, dst p

ISAKMP: Sending smartcard_removal_disconnect reply
                  value 0
ISAKMP: Sending APPLICATION_VERSION string: Cisco IOS Software, C2900 Software (C2900-UNIVERSALK9-M), Version 15.2(3
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2013 by Cisco Systems, Inc.
Compiled Wed 13-Mar-13 05:26 by prod_rel_team
ISAKMP (1007): Unknown Attr: MODECFG_HOSTNAME (0x700A)
ISAKMP (1007): Unknown Attr: CONFIG_MODE_UNKNOWN (0x7005)
ISAKMP:(1007): responding to peer config from 136.1.100.100. ID = 583863136
```

https://t.me/learningnets

```
ISAKMP: Marking node 583863136 for late deletion
ISAKMP:(1007): sending packet to 136.1.100.100 my_port 500 peer_port 58520 (R) CONF_ADDR
ISAKMP:(1007):Sending an IKE IPv4 Packet.
ISAKMP:(1007):Talking to a Unity Client
ISAKMP:(1007):Input = IKE_MESG_FROM_AAA, IKE_AAA_GROUP_ATTR
ISAKMP:(1007):Old State = IKE_CONFIG_AUTHOR_AAA_AWAIT  New State = IKE_P1_COMPLETE
```

The client now has all configuration information and proceeds to IPSec Phase 2 using Quick mode. IPSec SA is now negotiated and SPIs are created (note that only negotiation from the first and the final accepted proposal are presented):

```
ISAKMP (1007): received packet from 136.1.100.100 dport 500 sport 58520 Global (R) QM_IDLE
ISAKMP: set new node 1928434778 to QM_IDLE
ISAKMP:(1007): processing HASH payload. message ID = 1928434778
ISAKMP:(1007): processing SA payload. message ID = 1928434778
ISAKMP:(1007):Checking IPSec proposal 1
ISAKMP: transform 1, ESP_AES
ISAKMP:   attributes in transform:
ISAKMP:      authenticator is HMAC-MD5
ISAKMP:      key length is 256
ISAKMP:      encaps is 1 (Tunnel)
ISAKMP:      SA life type in seconds
ISAKMP:      SA life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:(1007):atts are acceptable.
ISAKMP:(1007):Checking IPSec proposal 1
ISAKMP:(1007):transform 1, IPPCP LZS
ISAKMP:   attributes in transform:
ISAKMP:      encaps is 1 (Tunnel)
ISAKMP:      SA life type in seconds
ISAKMP:      SA life duration (VPI) of  0x0 0x20 0xC4 0x9B
ISAKMP:(1007):atts are acceptable.
IPSEC(validate_proposal_request): proposal part #1
IPSEC(validate_proposal_request): proposal part #1,
  (key eng. msg.) INBOUND local= 136.1.18.1:0, remote= 136.1.100.100:0,
    local_proxy= 0.0.0.0/0.0.0.0/256/0,
    remote_proxy= 20.0.0.7/255.255.255.255/256/0,
    protocol= ESP, transform= NONE  (Tunnel),
    lifedur= 0s and 0kb,
    spi= 0x0(0), conn_id= 0, keysize= 256, flags= 0x0
IPSEC(validate_proposal_request): proposal part #2
IPSEC(validate_proposal_request): proposal part #2,
  (key eng. msg.) INBOUND local= 136.1.18.1:0, remote= 136.1.100.100:0,
    local_proxy= 0.0.0.0/0.0.0.0/256/0,
    remote_proxy= 20.0.0.7/255.255.255.255/256/0,
```

https://t.me/learningnets

```
    protocol= PCP, transform= NONE  (Tunnel),
    lifedur= 0s and 0kb,
    spi= 0x0(0), conn_id= 0, keysize= 0, flags= 0x0
IPSEC(ipsec_process_proposal): transform proposal not supported for identity:
 {esp-aes 256 esp-md5-hmac comp-lzs }
ISAKMP:(1007): IPSec policy invalidated proposal with error 256
!
!ISAKMP:(1007):Checking IPSec proposal 11
ISAKMP: transform 1, ESP_3DES
ISAKMP:   attributes in transform:ISAKMP:      authenticator is HMAC-MD5
ISAKMP:      encaps is 1 (Tunnel)
ISAKMP:      SA life type in seconds
ISAKMP:      SA life duration (VPI) of  0x0 0x20 0xC4 0x9B ISAKMP:(1007):atts are acceptable.
IPSEC(validate_proposal_request): proposal part #1
IPSEC(validate_proposal_request): proposal part #1,
  (key eng. msg.) INBOUND local= 136.1.18.1:0, remote= 136.1.100.100:0,
    local_proxy= 0.0.0.0/0.0.0.0/256/0,
    remote_proxy= 20.0.0.7/255.255.255.255/256/0,
    protocol= ESP, transform= NONE  (Tunnel),
    lifedur= 0s and 0kb,
    spi= 0x0(0), conn_id= 0, keysize= 0, flags= 0x0
ISAKMP:(1007): processing NONCE payload. message ID = 1928434778
ISAKMP:(1007): processing ID payload. message ID = 1928434778
ISAKMP:(1007): processing ID payload. message ID = 1928434778
ISAKMP:(1007):QM Responder gets spi
ISAKMP:(1007):Node 1928434778, Input = IKE_MESG_FROM_PEER, IKE_QM_EXCH
ISAKMP:(1007):Old State = IKE_QM_READY  New State = IKE_QM_SPI_STARVE
ISAKMP:(1007):deleting node 583863136 error FALSE reason "No Error"
ISAKMP:(1007):Node 1928434778, Input = IKE_MESG_INTERNAL, IKE_GOT_SPI
ISAKMP:(1007):Old State = IKE_QM_SPI_STARVE  New State = IKE_QM_IPSEC_INSTALL_AWAIT
IPSEC(key_engine): got a queue event with 1 KMI message(s)
IPSEC(crypto_ipsec_create_ipsec_sas): Map found EZVPN_DYNAMICIPSEC(create_sa): sa created,

  (sa) sa_dest= 136.1.18.1, sa_proto= 50,
    sa_spi= 0xCE106699(3457181337),
    sa_trans= esp-3des esp-md5-hmac , sa_conn_id= 2013
    sa_lifetime(k/sec)= (4608000/3600)
IPSEC(create_sa): sa created,
  (sa) sa_dest= 136.1.100.100, sa_proto= 50,
    sa_spi= 0xEB400227(3946840615),
    sa_trans= esp-3des esp-md5-hmac , sa_conn_id= 2014
    sa_lifetime(k/sec)= (4608000/3600)
 ISAKMP: Failed to find peer index node to update peer_info_list
```

Finally, the EzVPN server installs a static reverse route into the local routing table,

https://t.me/learningnets

and Phase2 is completed:

```
IPSEC(rte_mgr): VPN Route Event Install new outbound sa: Static keyword or dynamic SA create for 136.1.100.100
IPSEC(rte_mgr): VPN Route Refcount 1 GigabitEthernet0/0
IPSEC(rte_mgr): VPN Route Added 20.0.0.7 255.255.255.255 via 136.1.100.100 in IP DEFAULT TABLE with tag 0 distance 1
ISAKMP:(1007):Received IPSec Install callback... proceeding with the negotiation
ISAKMP:(1007): sending packet to 136.1.100.100 my_port 500 peer_port 58520 (R) QM_IDLE
ISAKMP:(1007):Sending an IKE IPv4 Packet.
ISAKMP:(1007):Node 1928434778, Input = IKE_MESG_FROM_IPSEC, IPSEC_INSTALL_DONE
ISAKMP:(1007):Old State = IKE_QM_IPSEC_INSTALL_AWAIT  New State = IKE_QM_R_QM2
ISAKMP (1007): received packet from 136.1.100.100 dport 500 sport 58520 Global (R) QM_IDLE
ISAKMP:(1007):deleting node 1928434778 error FALSE reason "QM done (await)"
ISAKMP:(1007):Node 1928434778, Input = IKE_MESG_FROM_PEER, IKE_QM_EXCH
ISAKMP:(1007):Old State = IKE_QM_R_QM2  New State = IKE_QM_PHASE2_COMPLETE
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Server with ISAKMP Profiles

Task

- Configure R1 as Easy VPN server according to the following requirements:
  - Use 3DES/MD5 as cipher/hash for both IPSec phases.
  - Authenticate remote users identified by the group name **EZVPN** and the password **CISCO**.
  - Use the address pool 20.0.0.1-20.0.0.254 for remote users.
  - Enable Xauth against the local user database, and create a new user named **CISCOUSER1** with the password **CISCO1234**.
  - Only encrypt user traffic destined to R1's Loopback0 subnet.
  - Make use of ISAKMP profiles.
- Configure **TEST-PC-A** to verify your configuration.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

As seen from the configuration of the previous task, EzVPN authentication and authorization settings are globally applied at the static crypto-map level. What if we have a different set of users for which we want different authentication and authorization settings (local or from remote RADIUS servers)? In this situation, ISAKMP profiles allow you to obtain this configuration granularity, along with all of

https://t.me/learningnets

its features, as described in the L2L tasks. You should understand the following:

- The ISAKMP profile is applied at the dynamic crypto-map level (technically, it doesn't need to be applied because the router will match incoming connections against the configured ISAKMP profiles, which will then force all incoming EzVPN sessions to match on the first dynamic crypto-map entry.
- Globally configured settings at the static crypto-map level take precedence over settings configured at the ISAKMP profile level, so removing the settings at the static crypto-map level is recommended.

For our task, we re-use the same AAA authentication and authorization lists as in the previous task, as well as the EzVPN configuration group. We create an ISAKMP profile that matches the remote group ID and associates the ISAKMP settings with it.

Configuration

R1:

```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCOUSER1 password CISCO1234
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto isakmp profile CRYPTO_MAP
 match identity group EZVPN
```

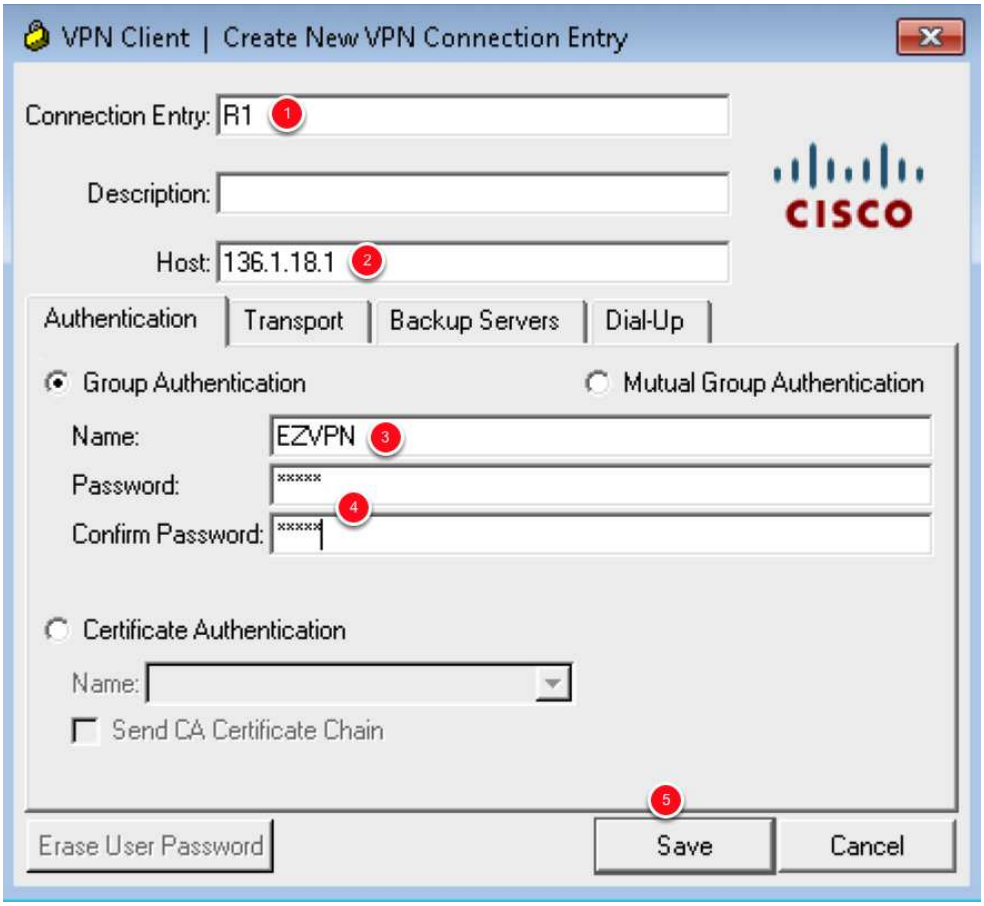
https://t.me/learningnets

```
 client authentication list EZVPN
 isakmp authorization list EZVPN
 client configuration address respond
 client configuration group EZVPN
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto dynamic-map EZVPN_DYNAMIC 100
 set isakmp-profile CRYPTO_MAP
 set transform-set 3DES_MD5
 reverse-route
!
crypto map VPN 100 ipsec-isakmp dynamic EZVPN_DYNAMIC
!
interface GigabitEthernet0/0
 crypto map VPN
```

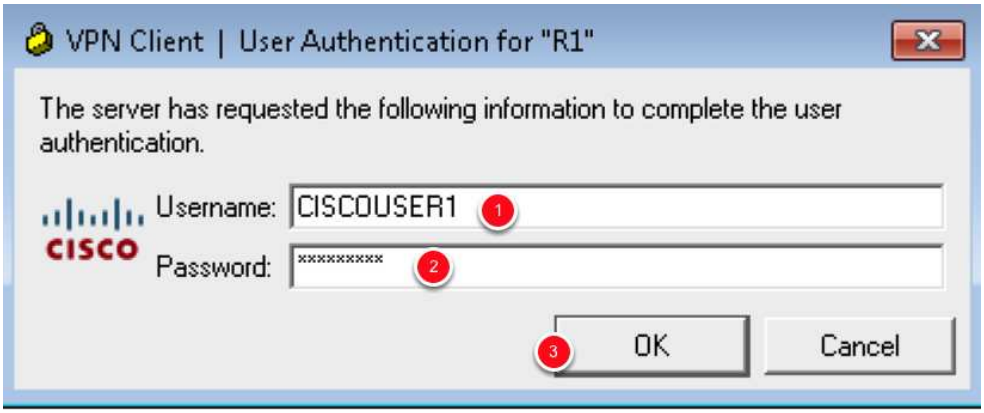
Verification

TEST-PC-A will be automatically assigned an IP address from the DHCP server configured on R3. Configure a new Cisco VPN Client connection entry by clicking the **New** button:

https://t.me/learningnets

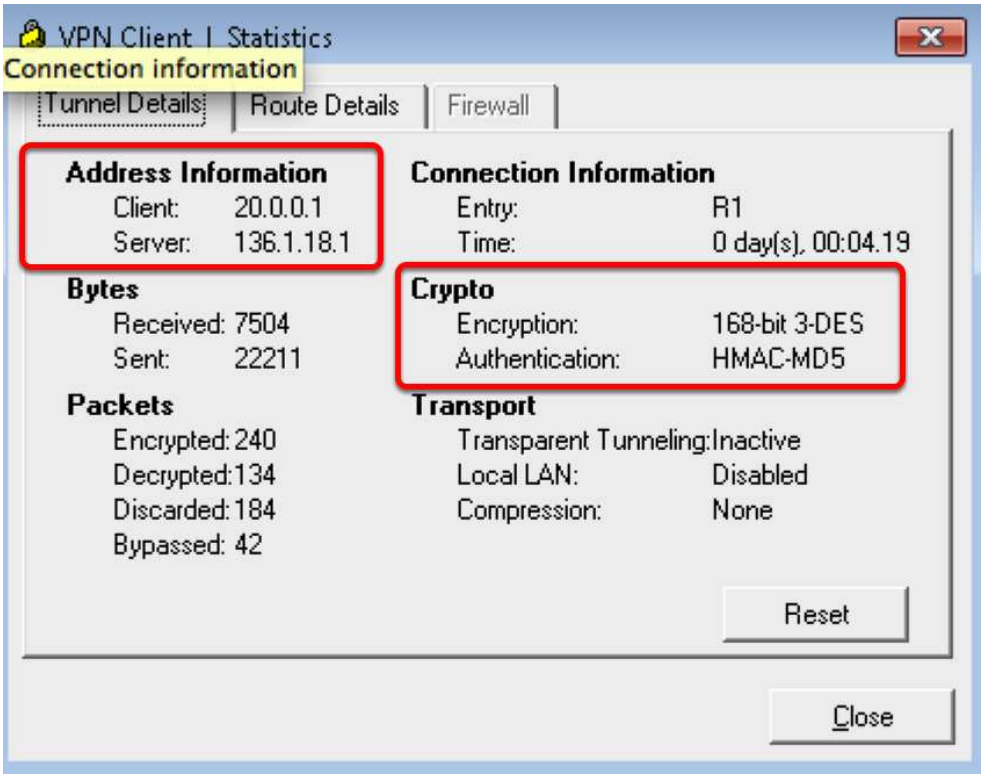


Initiate the connection by clicking the **Connect** button. An authentication window prompts you for a username and password:

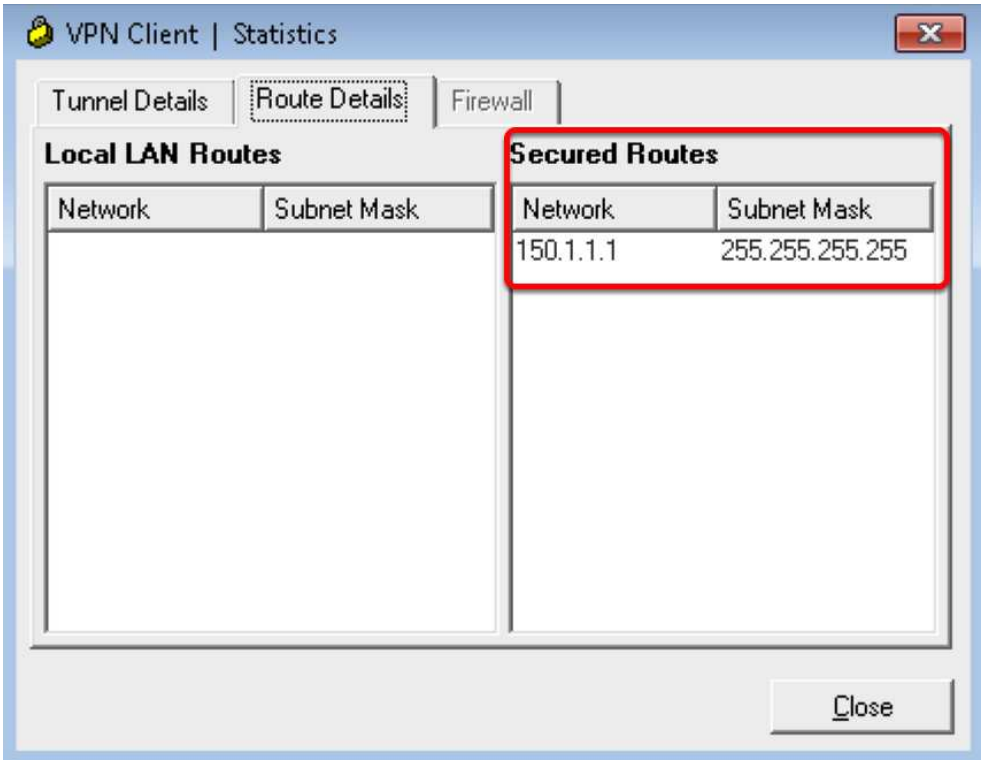


After the connection is established, right-click the VPN Client icon on the toolbar and open the **Statistics** window to verify negotiated algorithms, assigned IP address, and traffic statistics:

https://t.me/learningnets



On the **Route Details** tab, verify for which destinations traffic is forwarded through the tunnel; this is given by the split-tunneling access-list configured on the EzVPN server:

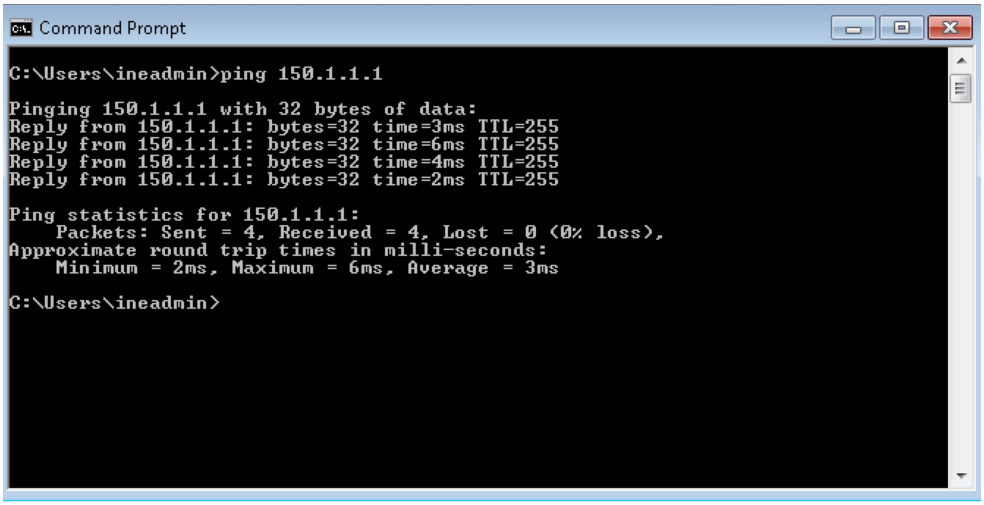


https://t.me/learningnets

Issue a debug on R1 to see the difference from the previous task output, basically based on the received remote IKE ID (group name); the IKE connection matches the configured ISAKMP profile and is forced to inherit all of its settings:

```
R1#debug crypto isakmp
Crypto ISAKMP debugging is on
!
!
ISAKMP (0): received packet from 136.1.100.100 dport 500 sport 52768 Global (N) NEW SA
ISAKMP: Created a peer struct for 136.1.100.100, peer port 52768
ISAKMP: New peer created peer = 0x2CC17F98 peer_handle = 0x8000000B
ISAKMP: Locking peer struct 0x2CC17F98, refcount 1 for crypto_isakmp_process_block
ISAKMP: local port 500, remote port 52768
ISAKMP: Find a dup sa in the avl tree during calling isadb_insert sa = 2D392530
ISAKMP:(0): processing SA payload. message ID = 0
ISAKMP:(0): processing ID p
Rack1R1#ayload. message ID = 0
ISAKMP (0): ID payload
        next-payload : 13
        type         : 11
        group id     : EZVPN
        protocol     : 17
        port         : 500
        length       : 13ISAKMP:(0):: peer matches CRYPTO_MAP profile
ISAKMP:(0):Setting client config settings 319A17ACISAKMP:(0):(Re)Setting client xauth list  and state
ISAKMP/xauth: initializing AAA request
ISAKMP:(0): Profile CRYPTO_MAP assigned peer the group named EZVPN
```

From TEST-PC-A, check connectivity with R1's Loopback0 subnet:



Verify that the IPsec session is established and matches the configured ISAKMP

https://t.me/learningnets

profile:

```
R1#show crypto isakmp profile

IKEv1 PROFILE CRYPTO_MAP Ref Count = 2
   Identities matched are:
    group EZVPN
   Certificate maps matched are:
   keyring(s): <none>
   trustpoint(s): <all>
!
!R1#show crypto session username CISCOUSER1
Crypto session current status

Interface: GigabitEthernet0/0Username: CISCOUSER1
Profile: CRYPTO_MAP
Group: EZVPN
Assigned address: 20.0.0.1Session status: UP-ACTIVE

Peer: 136.1.100.100 port 50068
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.100.100/50068 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 host 20.0.0.16
        Active SAs: 2, origin: dynamic crypto map
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Server with DVTI and PSK

Task

- Configure R1 as Easy VPN server according to the following requirements:
  - Use 3DES/MD5 as cipher/hash for both IPSec phases.
  - Authenticate remote users identified by the group name **EZVPN** and the password **CISCO**.
  - Use the address pool 20.0.0.1-20.0.0.254 for remote users.
  - Enable Xauth against the local user database, and create a new user named **CISCOUSER1** with the password **CISCO1234**.
  - Only encrypt user traffic destined to R1's Loopback0 subnet.
- Configure **TEST-PC-A** to verify your configuration, and do not use any `crypto-map` -related commands.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

As described in the L2L section, Virtual Tunnel Interfaces (VTIs) allow for significantly simpler configuration of IPSec VPNs and offers additional functionalities. Recall that a special IPSec SA is associated with the VTI interface, encrypting all traffic routed across the tunnel. To implement an L2L VPN, point-to-point tunnel interface (SVTI) is used. For remote access VPN implementation, IOS

https://t.me/learningnets

introduces a special type of virtual tunnel interfaces, namely Dynamic VTI (DVTI) configured using the command **`interface virtual-template <NUMBER> type tunnel`**. The DVTI acts as a template for all incoming sessions, and for each established session a dynamic Virtual-Access interface is created, allowing application of per-session IPsec settings as with SVTI.

ISAKMP profiles were also discussed in the L2L tasks; to make the binding between the profile and the DVTI, the ISAKMP profile-level command **`virtual-template <NUMBER>`** is used. As usual, we create an IPSec profile and apply it to the virtual-template interface. Notice that the profile applies the reverse-route settings to the virtual-template interface. No crypto-maps are used for this scenario, because all IPsec Phase 2 policy is now encapsulated within the VTI and IPsec profile. Applying the ISAKMP profile at the IPsec profile level is not technically required; it was configured just to make the configuration clear.

Configuration

R1:

```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCOUSER1 password CISCO1234
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto isakmp profile CRYPTO_MAP
 match identity group EZVPN
```

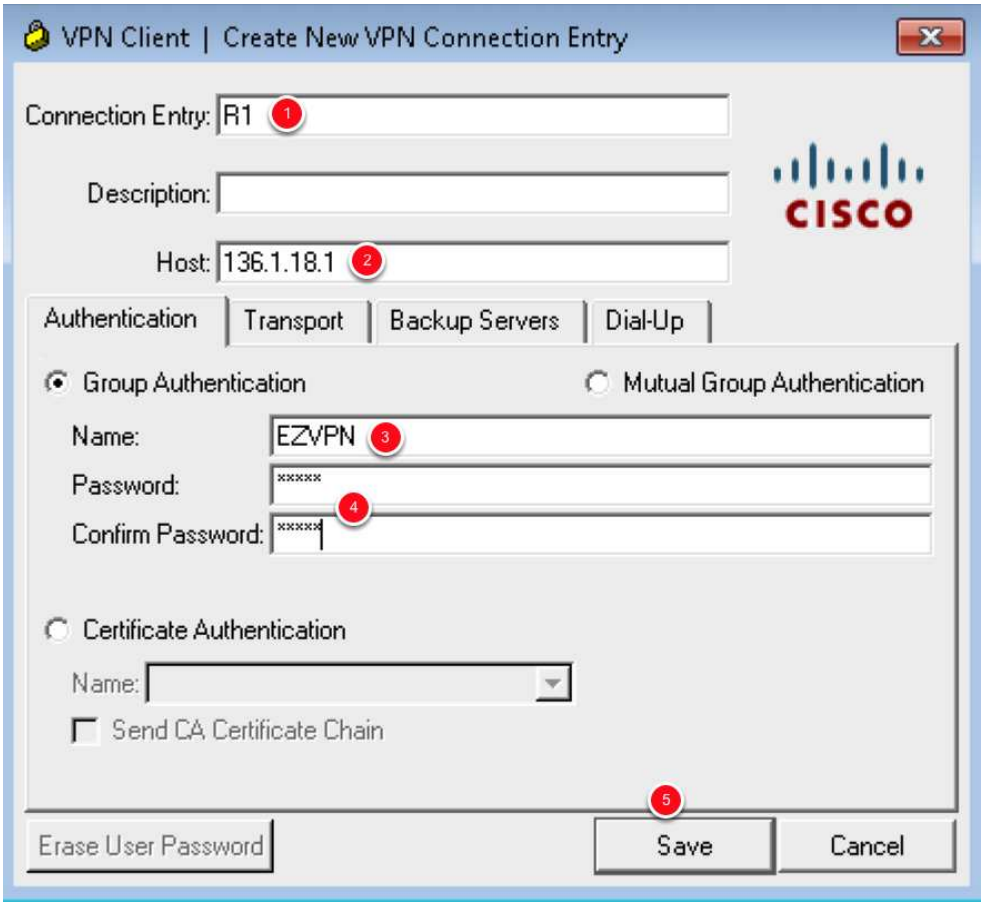
https://t.me/learningnets

```
 client authentication list EZVPN
 isakmp authorization list EZVPN
 client configuration address respond
 client configuration group EZVPN
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto ipsec profile EZVPN
 set isakmp-profile CRYPTO_MAP
 set transform-set 3DES_MD5
 set reverse-route tag 1
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile EZVPN
```

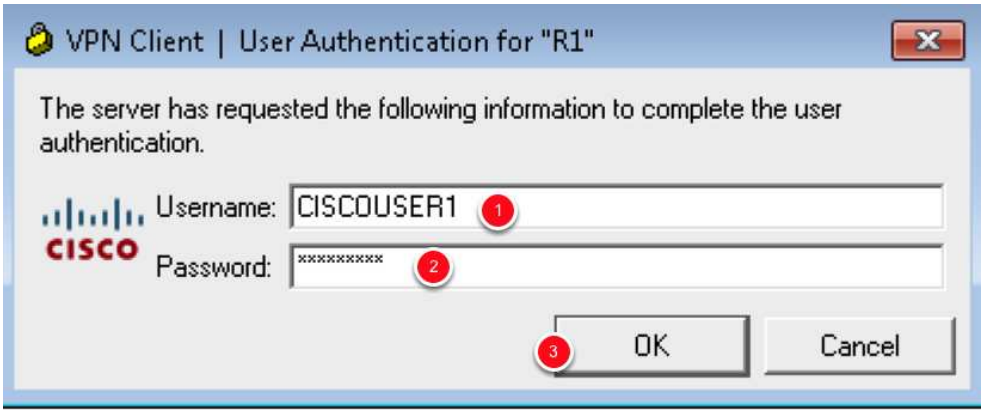
Verification

TEST-PC-A will be automatically assigned an IP address from the DHCP server configured on R3. Configure a new Cisco VPN Client connection entry by clicking the **New** button:

https://t.me/learningnets

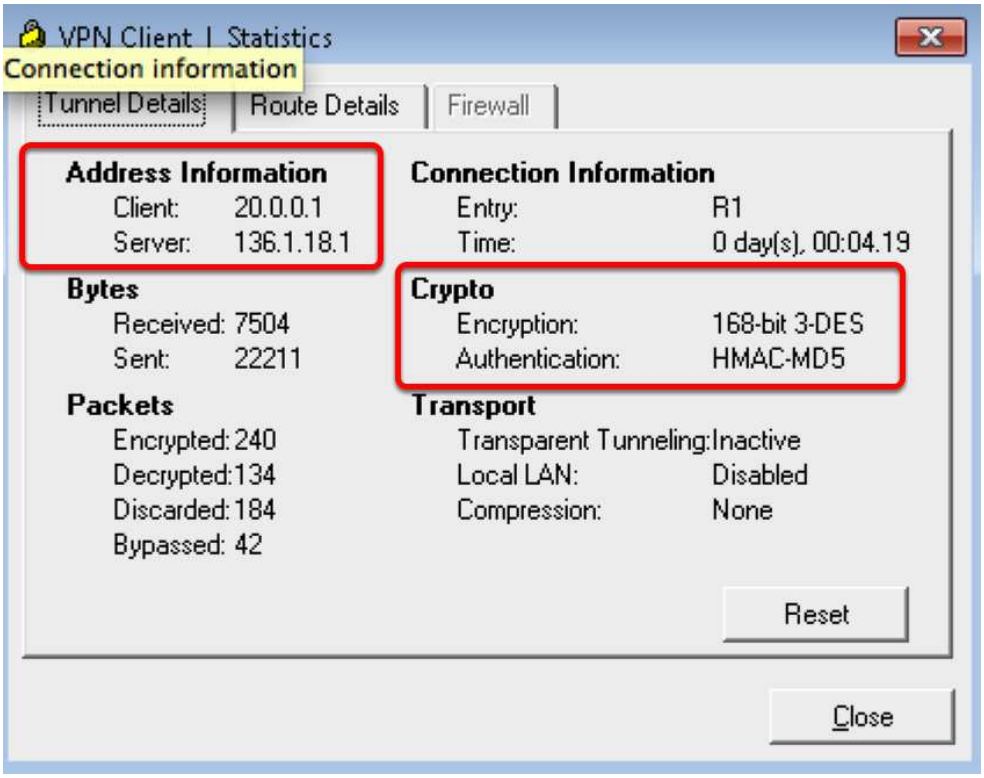


Initiate the connection by clicking the **Connect** button. An authentication window prompts you for a username and password:

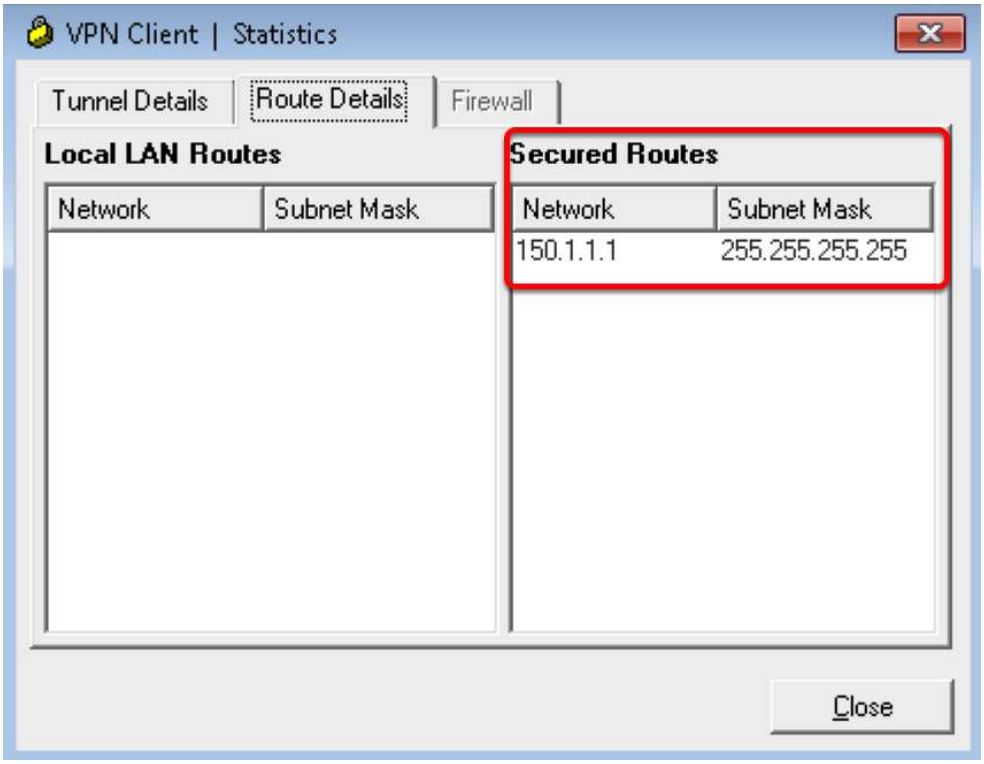


After the connection is established, right-click the VPN Client icon on the toolbar and open the **Statistics** window to verify negotiated algorithms, assigned IP address, and traffic statistics:

https://t.me/learningnets

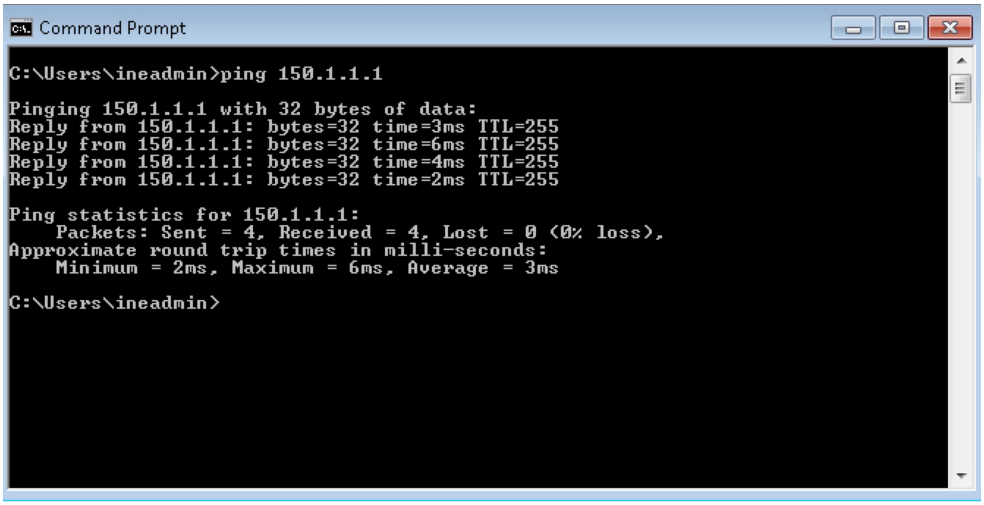


On the **Route Details** tab, verify for which destinations traffic is forwarded through the tunnel; this is given by the split-tunneling access-list configured on the EzVPN server:



Check connectivity with R1's Loopback0 subnet:

https://t.me/learningnets



Verify IPsec Phase1/Phase2:

```
R1#show crypto isakmp peers config
 Client-Public-Addr=136.1.100.100:52360; Client-Assigned-Addr=20.0.0.1
Client-OS=Cisco Systems VPN Client 5.0.07.0410:WinNT; Client-Group=EZVPN; Client-User=CISCOUSER1;
!
!R1#show crypto ipsec sa peer 136.1.100.100
 interface: Virtual-Access1
    Crypto map tag: Virtual-Access1-head-0, local addr 136.1.18.1

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
   remote ident (addr/mask/prot/port): (20.0.0.1/255.255.255.255/0/0)
   current_peer 136.1.100.100 port 52360
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
 #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.1, remote crypto endpt.: 136.1.100.100
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0
     current outbound spi: 0x2952E23C(693297724)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0xDF05BA71(3741694577)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2035, flow_id: Onboard VPN:35, sibling_flags 80000040, crypto map: Virtual-Access1-head-0
        sa timing: remaining key lifetime (k/sec): (4362160/3530)
        IV size: 8 bytes
        replay detection support: Y
```

https://t.me/learningnets

```
Status: ACTIVE(ACTIVE)


     inbound ah sas:

     inbound pcp sas:
 outbound esp sas:
      spi: 0x2952E23C(693297724)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2036, flow_id: Onboard VPN:36, sibling_flags 80000040, crypto map: Virtual-Access1-head-0
        sa timing: remaining key lifetime (k/sec): (4362160/3530)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

As shown in the following output, the dynamically created Virtual-Access interface acts like an SVTI, where both tunnel source and destination have been dynamically configured, and the other configuration is cloned from the Virtual-Template interface:

```
R1#show running-config interface virtual-access1
Building configuration...

Current configuration : 193 bytes
!
interface Virtual-Access1 tunnel source 136.1.18.1
 tunnel mode ipsec ipv4 tunnel destination 136.1.100.100
 tunnel protection ipsec profile EZVPN
 no tunnel protection ipsec initiate
!
!R1#show interfaces virtual-access1
Virtual-Access1 is up, line protocol is up
  Hardware is Virtual Access interface Interface is unnumbered. Using address of Loopback0 (150.1.1.1)
  MTU 17886 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL Tunnel vaccess, cloned from Virtual-Template1
  Vaccess status 0x0, loopback not set
  Keepalive not set Tunnel source 136.1.18.1, destination 136.1.100.100
 Tunnel protocol/transport IPSEC/IP
  Tunnel TTL 255
  Tunnel transport MTU 1446 bytes
```

https://t.me/learningnets

```
  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps) Tunnel protection via IPSec (profile "EZVPN")
```

Verify that the reverse-route is injected through the new DVTI, check the route tag:

```
R1#show ip route static | b Gateway
Gateway of last resort is not set

      20.0.0.0/32 is subnetted, 1 subnets S        20.0.0.1 is directly connected, Virtual-Access1

!
!R1#show ip route 20.0.0.1
Routing entry for 20.0.0.1/32
  Known via "static", distance 1, metric 0 (connected) Tag 1
  Routing Descriptor Blocks:
  * directly connected, via Virtual-Access1
      Route metric is 0, traffic share count is 1 Route tag 1
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Server with Group Lock

Task

- Modify the EzVPN server configuration from the previous task, so that user **CISCO** is only allowed to remotely connect to group **EZVPN**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The purpose of the Group Lock feature is to make sure a user is only allowed to log in under a certain group. Without this feature, a user who knows another group's password may use the group's name to authenticate with the server and gain access to features not allowed for its native group. When the Group Lock feature is enabled for a group, all XAUTH usernames must be in one of the following formats, where **GROUP** is the EzVPN group name:

- username@GROUP
- username\GROUP
- username/GROUP
- username%GROUP

The server will extract the group name from the XAUTH username and compare it with the EzVPN group name. If they don't match, the connection is rejected. This

https://t.me/learningnets

automatically means that all names stored in the local database must follow the syntax specified above. Notice that this Group Lock functionality differs from the Group Lock feature found in ASA firewalls, where users are allowed to log in using a regular username/password combination. In addition, IOS Group Lock only works with pre-shared-keys authentication, not with digital signatures.

Configuration

R1:

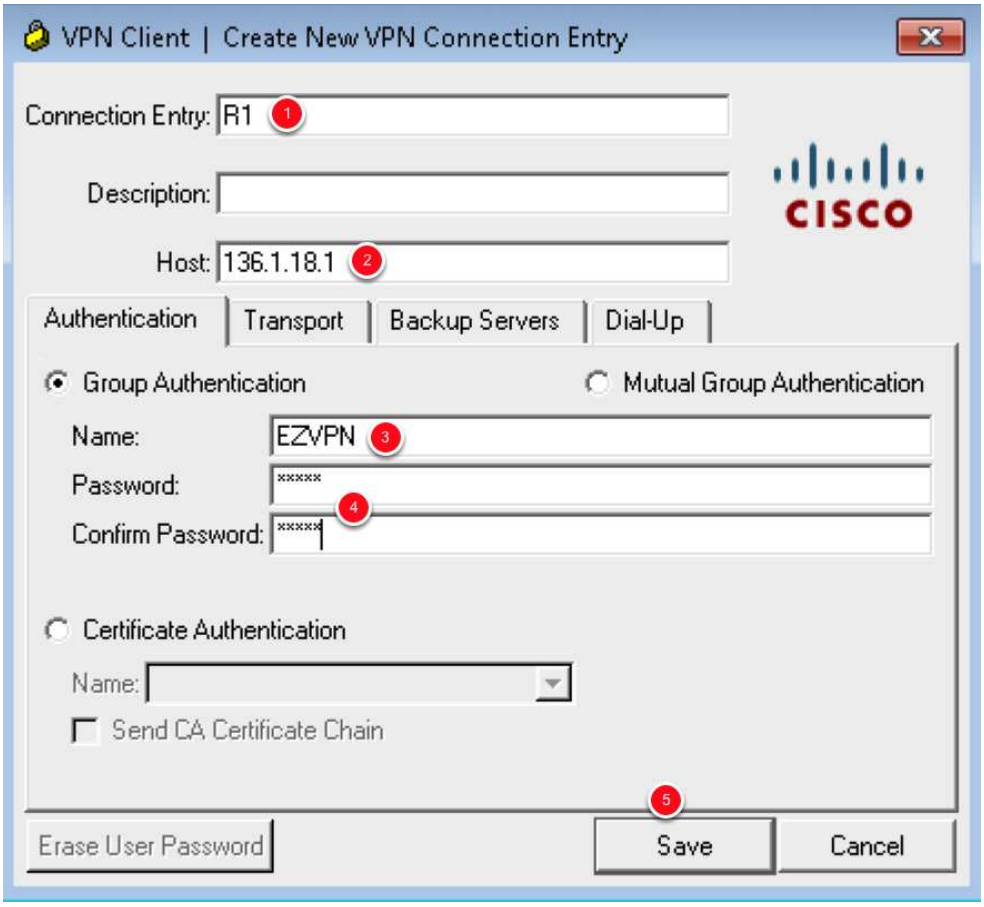
```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCO@EZVPN password CISCO
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 group-lock
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto isakmp profile CRYPTO_MAP
 match identity group EZVPN
 client authentication list EZVPN
 isakmp authorization list EZVPN
 client configuration address respond
 client configuration group EZVPN
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
```

https://t.me/learningnets

```
crypto ipsec profile EZVPN
 set isakmp-profile CRYPTO_MAP
 set transform-set 3DES_MD5
 set reverse-route tag 1
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile EZVPN
```

Verification

TEST-PC-A will be automatically assigned an IP address from the DHCP server configured on R3. Configure a new Cisco VPN Client connection entry by clicking the **New** button:

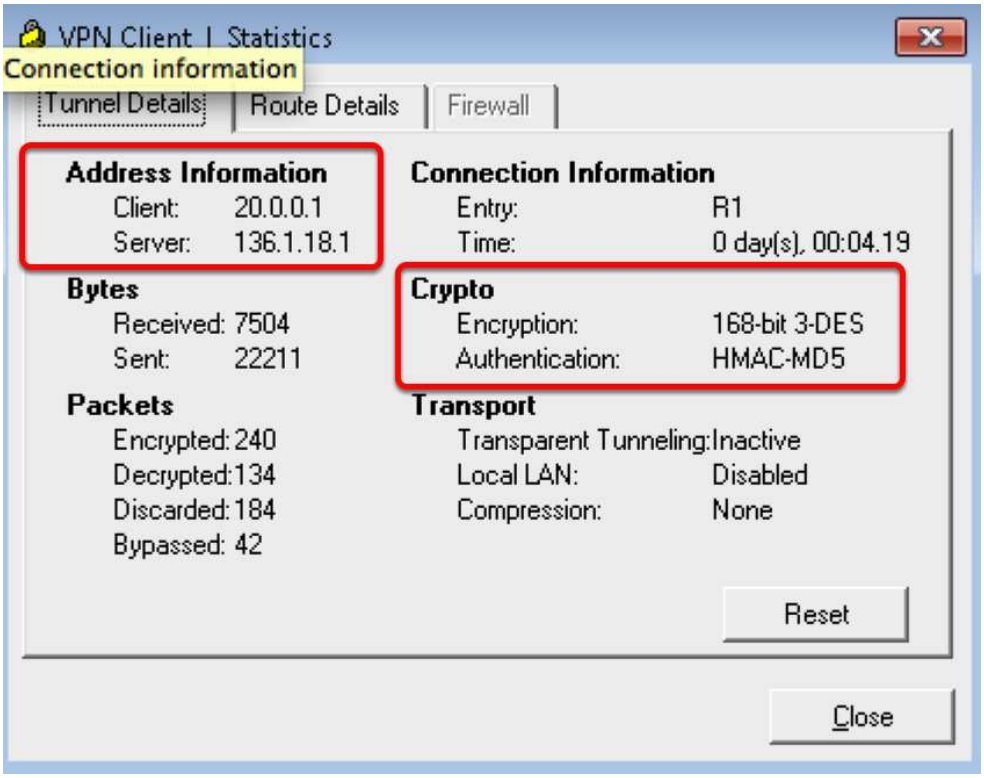


Initiate the connection by clicking the **Connect** button. An authentication window prompts you for a username and password; specify both username and group:

https://t.me/learningnets

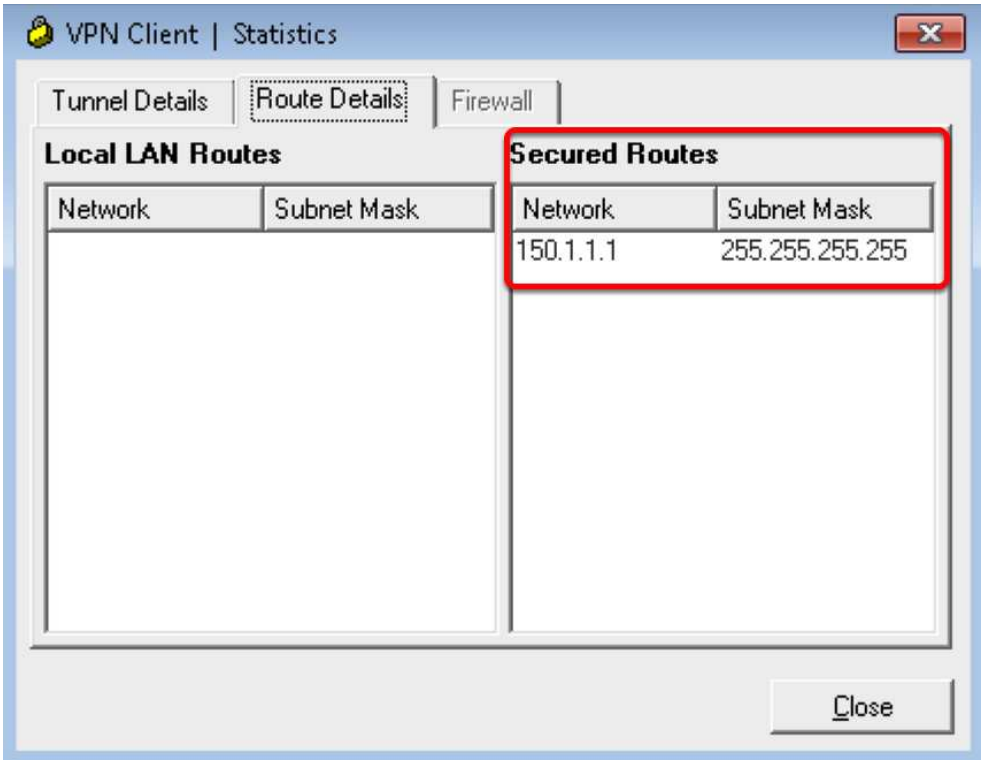


After the connection is established, right-click the VPN Client icon on the toolbar and open the **Statistics** window to verify negotiated algorithms, assigned IP address, and traffic statistics:

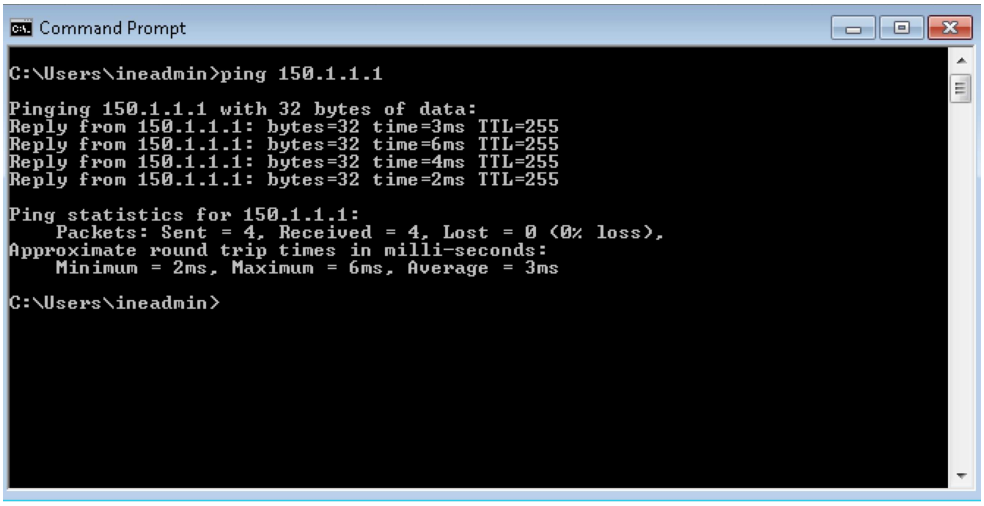


On the **Route Details** tab, verify for which destinations traffic is forwarded through the tunnel; this is given by the split-tunneling access-list configured on the EzVPN server:

https://t.me/learningnets



Check connectivity with R1's Loopback0 subnet:



Verify on the EzVPN server that both the username and group have been entered:

```
R1#show crypto isakmp peers config

Client-Public-Addr=136.1.100.100:62979; Client-Assigned-Addr=20.0.0.1;
Client-OS=Cisco Systems VPN Client 5.0.07.0410:WinNT; Client-Group=EZVPN;Client-User=CISCO@EZVPN;

!
!R1#show crypto session
Crypto session current status

Interface: Virtual-Access1Username: CISCO@EZVPN
```

https://t.me/learningnets

```
Profile: CRYPTO_MAP
Group: EZVPN
Assigned address: 20.0.0.1 Session status: UP-ACTIVE

Peer: 136.1.100.100 port 62979
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.100.100/62979 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 host 20.0.0.1
        Active SAs: 2, origin: crypto map
```

Test the Group Lock feature by connecting with the user of **CISCO** instead of **CISCO@EZVPN**. On the XAUTH phase, if a **`debug crypto isakmp`** command is used, you'll see a message saying, **`User Authentication in this group failed.`**

```
ISAKMP:(1024):Input = IKE_MESG_FROM_PEER, IKE_AM_EXCH
ISAKMP:(1024):Old State = IKE_R_AM2  New State = IKE_P1_COMPLETE

ISAKMP:(1024):Need XAUTH
ISAKMP: set new node -216238437 to CONF_XAUTH   ISAKMP/xauth: request attribute XAUTH_USER_NAME_V2
ISAKMP/xauth: request attribute XAUTH_USER_PASSWORD_V2
ISAKMP:(1024): initiating peer config to 136.1.100.100. ID = 4078728859
ISAKMP:(1024): sending packet to 136.1.100.100 my_port 500 peer_port 53993 (R) CONF_XAUTH
ISAKMP:(1024):Sending an IKE IPv4 Packet.
ISAKMP:(1024):Input = IKE_MESG_INTERNAL, IKE_PHASE1_COMPLETE
ISAKMP:(1024):Old State = IKE_P1_COMPLETE  New State = IKE_XAUTH_REQ_SENT

ISAKMP (1024): received packet from 136.1.100.100 dport 500 sport 53993 Global (R) CONF_XAUTH
ISAKMP:(1024):processing transaction payload from 136.1.100.100. message ID = -216238437
ISAKMP: Config payload REPLY ISAKMP/xauth: reply attribute XAUTH_USER_NAME_V2
ISAKMP/xauth: reply attribute XAUTH_USER_PASSWORD_V2
ISAKMP:(1024):User Authentication in this group failed

ISAKMP: set new node -1260806757 to CONF_XAUTH
ISAKMP:(1024): initiating peer config to 136.1.100.100. ID = 3034160539
ISAKMP:(1024): sending packet to 136.1.100.100 my_port 500 peer_port 53993 (R) CONF_XAUTH
ISAKMP:(1024):Sending an IKE IPv4 Packet.
ISAKMP:(1024):Input = IKE_MESG_FROM_PEER, IKE_CFG_REPLY
ISAKMP:(1024):Old State = IKE_XAUTH_REQ_SENT  New State = IKE_XAUTH_SET_SENT

ISAKMP (1024): received packet from 136.1.100.100 dport 500 sport 53993 Global (R) CONF_XAUTH
ISAKMP:(1024):processing transaction payload from 136.1.100.100. message ID = -1260806757
ISAKMP: Config payload ACK
ISAKMP:(1024):peer does not do paranoid keepalives.
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Remote in Client Mode with PSK

Task

- Configure R1 as Easy VPN server according to the following requirements:
  - Use 3DES/MD5 as cipher/hash for both IPSec phases.
  - Authenticate remote clients identified by the group name **EZVPN** and the password **CISCO**.
  - Use the address pool 20.0.0.1-20.0.0.254 for remote clients.
  - Enable Xauth against the local user database using the group-lock functionality.
- Configure R3 as EzVPN Remote in Client mode, and ensure that only traffic to R1's Loopback0 subnet brings the EzVPN tunnel UP.
- Users should enter XAUTH credentials by starting an HTTP session across the client router.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Since IOS 12.2T, Cisco has included support for EzVPN client functionality in its routers. This feature allows the use of small office/home routers as hardware EzVPN clients connecting to a central EzVPN server. A typical small router has one outside interface connected to the ISP and one inside interface connected to the

https://t.me/learningnets

internal network. Using the router as EzVPN client allows for an encrypted tunnel to the central site, encrypting only relevant traffic by means of split-tunneling while allowing all Internet traffic to go directly to the ISP. The EzVPN Remote feature supports three modes of operation:

- **Client mode** requires the router requesting an IP address using ISAKMP Mode Config (just like Cisco VPN Client). This IP address is used to create a /32 PAT pool in the router and translate all traffic sourced off the inside subnet destined to the split-tunnel protected corporate networks using this pool. Thus, all inside office traffic is PAT translated and then encrypted using the IPsec SA. The central office never sees the actual remote office IP addressing. EzVPN allows the use of custom NAT configuration in the router with Client mode, but only if split-tunneling is not in use.
- **Network Extension mode** does not require a new IP address assigned to the client. Instead, the client router presents the server with its inside network subnet during IPsec SA negotiation, by means of Proxy IDs. The central router creates the respective Phase 2 SA and injects the corresponding reverse route into the routing table. The server propagates the prefix across the central network through IGP, allowing all routers to learn the IP subnet of the remote office. With NEM, address translation is not required on the client router, and all traffic sourced off the client inside network to the split-tunnel protected subnets is IPsec encapsulated. The client router may still request other ISAKMP Mode Config information, such as internal DNS servers or configuration URL.
- **Network Plus mode** is identical to Network Extension mode except that it allows the client router to request an IP address from the server. The client router does not use this IP address for PAT, but automatically assigns it to a loopback interface. This allows for verification and troubleshooting using the well-known IP address, in situations where the remote inside network is not reachable for some reason.

Like the regular EzVPN client, EzVPN Remote supports a two-step authentication process:

- The first is the group authentication, where both parties establish the ISAKMP SA; pre-shared-keys or digital certificates can be used for authentication.
-

https://t.me/learningnets

The second is XAUTH, or extended authentication, implemented using ISAKMP Mode Config. This requires input of the username and the password by the end user interactively. It is possible to configure the router to store the XAUTH credentials in the configuration, if the server allows this option. In case of interactive authentication, the user may enter the username/password via the router's CLI using the `crypto ipsec client ezvpn xauth <PROFILE_NAME>` command or by using the SDM Web-interface, if SDM is installed on the router. Alternatively, you may configure the router to intercept transit HTTP sessions (using the HTTP authentication proxy function) initiated across the router and prompt the internal user for XAUTH credentials interactively.

With the regular Cisco IPsec VPN client, you activate the VPN tunnel manually. With the EzVPN Remote feature, you have three options:

- The default option is automatic activation, which brings the tunnel up as soon as the EzVPN profile is applied to the outside interface. Note that this option requires you to store XAUTH credentials in the router or not use XAUTH at all.
- The second option is manual tunnel activation, which is triggered through CLI with the `crypto ipsec client ezvpn connect <PROFILE_NAME>` command or through web-based interface. The router will prompt the user for XAUTH credentials if needed.
- The last option is traffic-triggered activation, which brings up the tunnel only when traffic matching the configured access-list crosses the outside interface. This option is similar to automatic tunnel activation, so it requires the XAUTH credentials to be stored locally in the router.

To configure EzVPN Remote, perform the following steps:

1. Configure an EzVPN profile using the command `crypto ipsec client ezvpn <GROUP_NAME>` ; further configuration includes setting various sub-commands under the profile. There are mandatory and optional settings. Here are the mandatory options:
    - `mode [client|network-extension|network-plus]` defines the EzVPN remote mode; the default is client.
    - `connect [auto|manual|acl]` configures the tunnel trigger mode.
    - `group <EZVPN_GROUP> key <PSK>` defines the group name used to authenticate with the server and the pre-shared-key used for authentication, which must match those configured on the EzVPN server.
    - `peer [<IP_ADDRESS>|<HOSTNAME>] [default]` defines the EzVPN server IP address or DNS name. If a name is specified, it is resolved every time a

https://t.me/learningnets

tunnel starts, allowing for dynamic IP address change via DNS. There could be multiple peer statements defined under the profile, and they are tried in sequence until at least one responds. As with L2L connections, you may designate the first peer as preferred, for reasons like latency or link costs, by using the `default` keyword. The EzVPN client will try to connect to the primary peer first. If the primary peer fails, it will try the other peer but continue to try to reach the primary at the same time. If at any moment the primary peer revives, the EzVPN client drops the current connection and reactivates the primary tunnel.

2. Configure any non-mandatory settings. Here are just a few of them. Other important settings will be discussed in separate tasks.
    - `username <XAUTH_NAME> password <XAUTH_PASSWORD>` defines Phase 1.5 XAUTH credentials. This in only possible if the server allows the save-password option; otherwise, the router will not use the stored credentials. If you are using the automatic tunnel connection method, you may need this command to allow the router establishing the tunnel without user intervention.
    - `local-address [<INTERFACE>]` specifies the local router interface to be used for ISAKMP/IPSec SA endpoint. You may want to use a loopback interface if you have multiple uplink interfaces with EzVPN profile applied.
    - `idletime <SECONDS>` specifies the amount of time the IPSec SA should remain idle to bring down the EzVPN tunnel. This is particularly useful for traffic-triggered tunnels because of their connect-on-demand behavior. The default is no idle time defined.
    - `xauth userid mode [http-intercept|interactive]` specifies the way to obtain XAUTH credentials interactively. If the credentials are not configured under the EzVPN profile, they could be prompted in three different ways: via CLI using the command `crypto ipsec client ezvpn xauth [<PROFILE_NAME>]` , by using HTTP intercept, or by using SDM router interface.
3. Apply the profile to the inside and outside interface, defining their roles. The interface command is `crypto ipsec client ezvpn <NAME> [inside|outside]` , with `outside` being the default setting. You need at least one inside and one outside interface to get a working EzVPN configuration, similar to how classic NAT with inside/outside interfaces works.

https://t.me/learningnets

Note that no IKE Phase1 or Phase2 algorithms are specified on the EzVPN remote; these are automatically pushed from the EzVPN server. This is in contrast to a scenario in which Cisco IPsec VPN Client is used, in which the client has built-in Phase1/Phase2 settings that are negotiated with the EzVPN server.

Configuration

R1:

```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCO@EZVPN password CISCO
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 group-lock
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto isakmp profile CRYPTO_MAP
 match identity group EZVPN
 client authentication list EZVPN
 isakmp authorization list EZVPN
 client configuration address respond
 client configuration group EZVPN
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto ipsec profile EZVPN
 set isakmp-profile CRYPTO_MAP
```

https://t.me/learningnets

```
 set transform-set 3DES_MD5
 set reverse-route tag 1
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile EZVPN
```

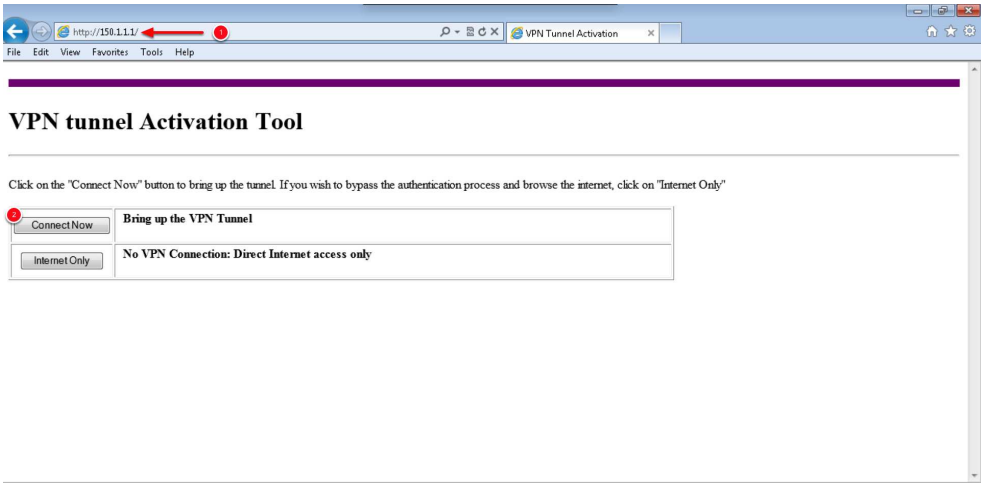
R3:

```
ip http server
!
ip access-list extended EZVPN_TRAFFIC
 permit ip 136.1.100.0 0.0.0.255 host 150.1.1.1
!
crypto ipsec client ezvpn EZVPN_CLIENT
 group EZVPN key CISCO
 connect acl EZVPN_TRAFFIC
 mode client
 peer 136.1.18.1
 xauth userid mode http-intercept
!
interface FastEthernet0/1
 crypto ipsec client ezvpn EZVPN_CLIENT inside
!
interface FastEthernet0/0.38
 crypto ipsec client ezvpn EZVPN_CLIENT outside
```

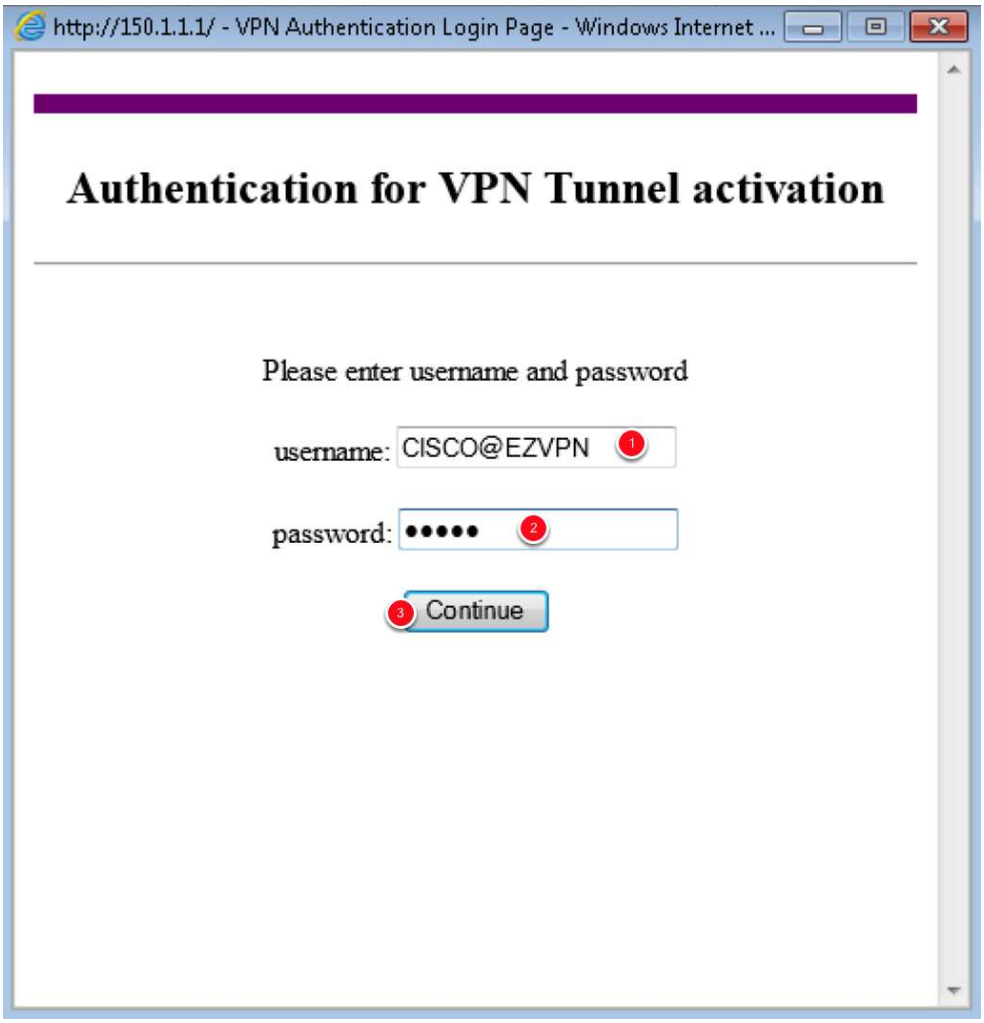
Verification

Verify configuration by trying to bring the tunnel up. Initiate an HTTP session from **TEST-PC-A** across the router, toward a VPN tunnel–protected destination (in our case, it is R1's Loopback0 subnet). The EzVPN client router will intercept the session and present you with the VPN activation tool screen. You may activate the tunnel or leave it as is. If you activate it, the router will prompt you for XAUTH credentials:

https://t.me/learningnets

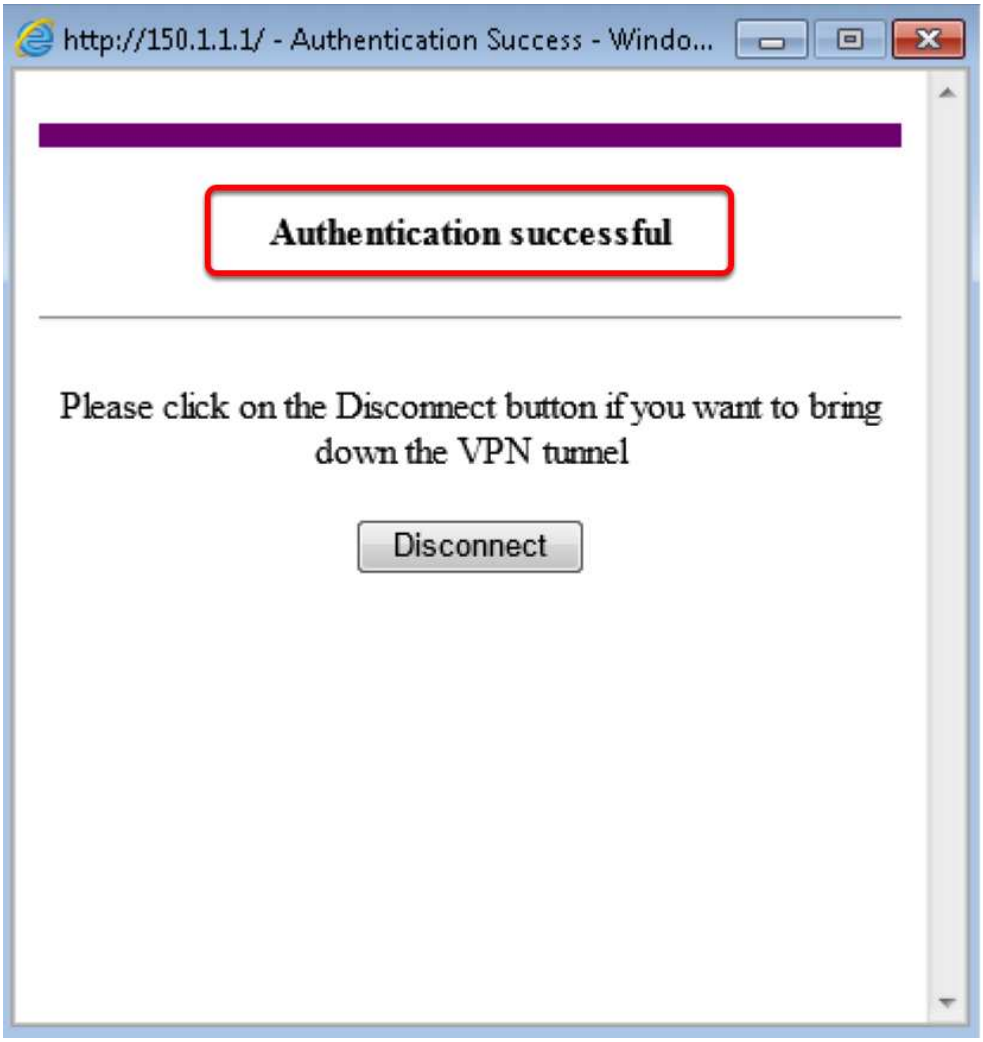


Clicking **Connect Now** will prompt you to enter a username and password (if the group lock configuration is still active on the EzVPN server, you must enter the username in the required format):



If the username and password are correct and your browser does not have a pop-up blocker enabled, an Authentication Successful message appears:

https://t.me/learningnets



Verify that traffic from VLAN 100 toward R1's Loopback0 subnet is sent over the tunnel:

```
R3#clear crypto sa counters
!R3#show crypto ipsec sa peer 136.1.18.1 | i pkts
#pkts encaps: 0, #pkts encrypt: 0, #pkts digest: 0
#pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
 !
!R3#ping 150.1.1.1 source fastEthernet0/1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 136.1.100.3 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
!
!R3#show crypto ipsec sa peer 136.1.18.1 | i pkts
```

https://t.me/learningnets

```
#pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
 #pkts decaps: 5, #pkts decrypt: 5, #pkts verify: 5

    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
```

After traffic is generated over the tunnel, verify the NAT translation table and statistics:

```
R3#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 20.0.0.22:1       136.1.100.3:1      150.1.1.1:1        150.1.1.1:1
!
!R3#show ip nat statistics
Total active translations: 1 (0 static, 1 dynamic; 1 extended)
Peak translations: 12, occurred 00:35:39 agoOutside interfaces:
 FastEthernet0/0.38
Inside interfaces:
 FastEthernet0/1
Hits: 335  Misses: 0
CEF Translated packets: 325, CEF Punted packets: 0
Expired translations: 40
Dynamic mappings:
-- Inside Source
[Id: 8] access-list EZVPN_CLIENT_internet-list interface FastEthernet0/0.38 refcount 0
[Id: 7] access-list EZVPN_CLIENT_enterprise-list pool EZVPN_CLIENT refcount 1
pool EZVPN_CLIENT: netmask 255.255.255.0
 start 20.0.0.22 end 20.0.0.22

        type generic, total addresses 1, allocated 1 (100%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

Verify the tunnel status and settings on both server and client:

```
R1#show crypto session remote 136.1.38.3
Crypto session current status
 Interface: Virtual-Access1
Username: CISCO@EZVPN
Profile: CRYPTO_MAP
Group: EZVPNAssigned address: 20.0.0.20
Session status: UP-ACTIVE
```

https://t.me/learningnets

```
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 host 20.0.0.20
        Active SAs: 2, origin: crypto map
!
!R3#show crypto ipsec client ezvpn
Easy VPN Remote Phase: 8

Tunnel name : EZVPN_CLIENT
Inside interface list: FastEthernet0/1
Outside interface: FastEthernet0/0.38
Connect : ACL based with access-list EZVPN_TRAFFIC
Current State: IPSEC_ACTIVE
Last Event: MTU_CHANGEDAddress: 20.0.0.20 (applied on Loopback10000)
Mask: 255.255.255.255
Save Password: Disallowed
Split Tunnel List: 1
       Address    : 150.1.1.1
       Mask       : 255.255.255.255
       Protocol   : 0x0
       Source Port: 0
       Dest Port  : 0XAuth credentials: HTTP intercepted

       HTTP return code : 200
       IP addr being prompted: 0.0.0.0
Current EzVPN Peer: 136.1.18.1
```

The next command demonstrates that auth-proxy is used for HTTP traffic interception and web-based tunnel management. The auth-proxy is not currently applied to any interface because the tunnel is now up:

```
R3#show ip auth-proxy configuration

Authentication Proxy Banner not configured
Consent Banner is not configured
Authentication global cache time is 60 minutes
Authentication global absolute time is 0 minutes
Authentication global init state time is 2 minutes
Authentication Proxy Session ratelimit is 100
Authentication Proxy Watch-list is disabled

Authentication Proxy Max HTTP process is 7
Authentication Proxy Auditing is disabled
Max Login attempts per user is 30
Authentication Proxy Rule Configuration
```

https://t.me/learningnets

```
Auth-proxy name ezvpnweb*** (EzVPN defined internal rule)

    http list not specified inactivity-timer 60 minutes
```

If we debug the EzVPN connection on the client side, we can identify the steps involved with the tunnel creation:

```
R3#debug crypto ipsec client ezvpn
EzVPN debugging is on
```

For traffic received on the EzVPN inside interface, the router checks to see whether it matches the protected traffic to trigger the EzVPN connection, so the following messages will be continuously seen on debug:

```
EZVPN: Checking Connect ACL
EZVPN: Checking Connect ACL
```

When the client clicks **Connect Now**, the router receives a signal from the web interceptor to start the EzVPN tunnel:

```
EZVPN(EZVPN_CLIENT): Received CONNECT from 136.1.100.100!
EZVPN(EZVPN_CLIENT): Current State: CONNECT_REQUIREDEZVPN(EZVPN_CLIENT): Event: CONNECT
EZVPN(EZVPN_CLIENT): ezvpn_connect_request
EZVPN(EZVPN_CLIENT): Found valid peer 136.1.18.1
EZVPN(EZVPN_CLIENT): Added PSK for address 136.1.18.1
EzVPN(EZVPN_CLIENT): sleep jitter delay 1575
EZVPN(EZVPN_CLIENT): New State: READY
EZVPN(EZVPN_CLIENT): Current State: READY
EZVPN(EZVPN_CLIENT): Event: IKE_PFS
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): Current State: READY
EZVPN(EZVPN_CLIENT): Event: CONN_UP
EZVPN(EZVPN_CLIENT): ezvpn_conn_up 05D36600 F082930A A848FF64 B5B18EDB
EZVPN(EZVPN_CLIENT): No state changeEZVPN(EZVPN_CLIENT): Current State: READY
```

The EzVPN server asks for XAUTH and the EzVPN client passes the request to the client, so the login page appears to the user:

```
EZVPN(EZVPN_CLIENT): Event: XAUTH_REQUEST
```

https://t.me/learningnets

```
EZVPN(EZVPN_CLIENT): ezvpn_xauth_request
EZVPN(EZVPN_CLIENT): ezvpn_parse_xauth_msg
EZVPN: Attributes sent in xauth request message:
        XAUTH_USER_NAME_V2(EZVPN_CLIENT):
        XAUTH_USER_PASSWORD_V2(EZVPN_CLIENT):EZVPN(EZVPN_CLIENT): Requesting following info for xauth
username:(Null)
password:(Null)
        message:Please enter username and password
EZVPN(EZVPN_CLIENT): New State: XAUTH_REQ
EZVPN: Checking Connect ACL
EZVPN(EZVPN_CLIENT): Current State: XAUTH_REQ
EZVPN(EZVPN_CLIENT): Event: XAUTH_PROMPTINGEZVPN(EZVPN_CLIENT): New State: XAUTH_PROMPT
```

The user responds with its credentials, received by the EzVPN Remote:

```
EZVPN(EZVPN_CLIENT): Communication from Interceptor application.
        Request/Response from 136.1.100.100, via FastEthernet0/1 username:CISCO@EZVPN
password:<omitted>
ok:Continue
```

When the user enters the username/password pair, it sends it over to the EzVPN server for validation:

```
EZVPN(EZVPN_CLIENT): Received usename|password from 136.1.100.100!
EZVPN(EZVPN_CLIENT): Current State: XAUTH_PROMPT
EZVPN(EZVPN_CLIENT): Event: XAUTH_REQ_INFO_READY
EZVPN(EZVPN_CLIENT): ezvpn_xauth_reply
        XAUTH_USER_NAME_V2(EZVPN_CLIENT): CISCO@EZVPN
        XAUTH_USER_PASSWORD_V2(EZVPN_CLIENT): <omitted>
EZVPN(EZVPN_CLIENT): New State: XAUTH_REPLIED
EZVPN(EZVPN_CLIENT): Current State: XAUTH_REPLIED
EZVPN(EZVPN_CLIENT): Event: XAUTH_STATUSEZVPN(EZVPN_CLIENT): xauth status received: Success

EZVPN(EZVPN_CLIENT): New State: READY
EZVPN(EZVPN_CLIENT): Current State: READY
```

After the Authentication Successful message is received from the EzVPN server, the Mode Config starts:

```
EZVPN(EZVPN_CLIENT): Event: MODE_CONFIG_REPLY 05D36600 54A2689A A848FF64 B242FFEC
EZVPN(EZVPN_CLIENT): ezvpn_parse_mode_config_msg
EZVPN: Attributes sent in message:
```

https://t.me/learningnets

```
Address: 20.0.0.22

        Peer has  IPsec Interface support
        Split Tunnel List: 1
              Address    : 150.1.1.1
              Mask       : 255.255.255.255
              Protocol   : 0x0
              Source Port: 0
              Dest Port  : 0
        Savepwd off
EZVPN: Unknown/Unsupported Attr: APPLICATION_VERSION (0x7)EZVPN(EZVPN_CLIENT): ezvpn_mode_config
EZVPN(EZVPN_CLIENT): ezvpn_nat_config
EZVPN(EZVPN_CLIENT): New State: SS_OPEN
EZVPN(EZVPN_CLIENT): Current State: SS_OPEN
EZVPN(EZVPN_CLIENT): Event: SOCKET_READY
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): Current State: SS_OPENEZVPN(EZVPN_CLIENT): Event: SOCKET_UP

ezvpn_socket_up
ezvpn_process_socket_up
%CRYPTO-6-EZVPN_CONNECTION_UP: (Client)  User=  Group=EZVPN  Client_public_addr=136.1.38.3  Server_public_addr=136.1
```

When Mode Config attributes are received from the EzVPN server (such as IP address) and the IPsec tunnel is up, the router disables the authentication proxy on the EzVPN inside interface (it no longer needs to check for interesting traffic); it also displays the Authentication Successful message to the user. The received IP address is assigned to a dynamically created loopback interface:

```
EZVPN(EZVPN_CLIENT): Disabling Connect ACL checking
EzVPN: Local Traffic Feature Deleted
EZVPN(EZVPN_CLIENT): Tunnel UP! Letting user know about it
EZVPN(EZVPN_CLIENT): New State: IPSEC_ACTIVEEZVPN(EZVPN_CLIENT): Current State: IPSEC_ACTIVE
EZVPN(EZVPN_CLIENT): Event: MTU_CHANGED
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): retaining context 9EZVPN(EZVPN_CLIENT): Removing auth_proxy from interface Fa0/1
EzVPN(EZVPN_CLIENT): No URL available to configure
%LINK-3-UPDOWN: Interface Loopback10000, changed state to up

%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback10000, changed state to up
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Remote in NEM Plus Mode with PSK

Task

- Configure R1 as Easy VPN server according to the following requirements:
  - Use 3DES/MD5 as cipher/hash for both IPSec phases.
  - Authenticate remote clients identified by the group name **EZVPN** and the password **CISCO**.
  - Use the address pool 20.0.0.1-20.0.0.254 for remote clients.
  - Enable Xauth against the local user database using the group-lock functionality.
- Configure R3 as EzVPN Remote in Network Extension Mode, and ensure that only traffic to R1's Loopback0 subnet brings the EzVPN tunnel UP.
- Users should enter XAUTH credentials at the router CLI.
- Ensure that R3 still receives an IP address from the EzVPN server.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The configuration is exactly the same as in the previous task; the only required change is the transition of R3 from client mode to network-plus mode.

https://t.me/learningnets

Configuration

R1:

```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCO@EZVPN password CISCO
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 group-lock
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto isakmp profile CRYPTO_MAP
 match identity group EZVPN
 client authentication list EZVPN
 isakmp authorization list EZVPN
 client configuration address respond
 client configuration group EZVPN
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto ipsec profile EZVPN
 set isakmp-profile CRYPTO_MAP
 set transform-set 3DES_MD5
 set reverse-route tag 1
!
interface Virtual-Template1 type tunnel
```

https://t.me/learningnets

```
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile EZVPN
```

R3:

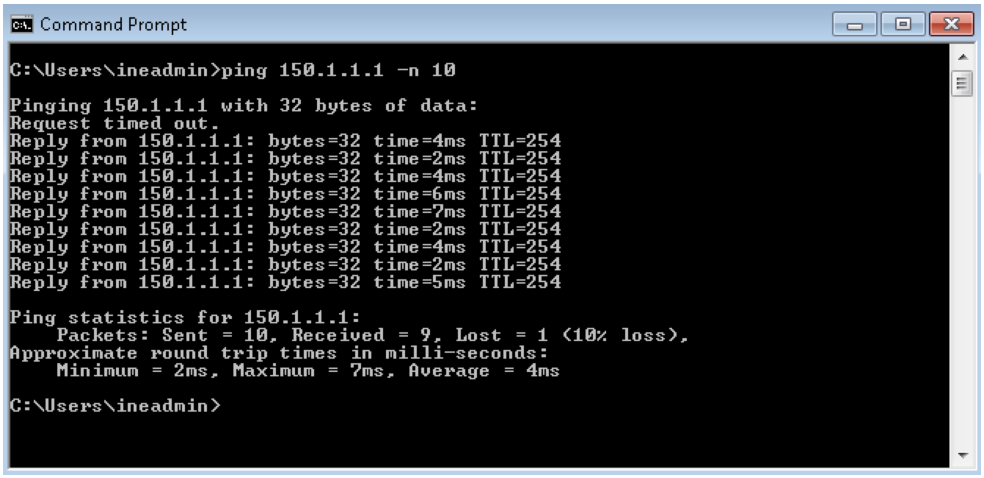
```
ip access-list extended EZVPN_TRAFFIC
 permit ip 136.1.100.0 0.0.0.255 host 150.1.1.1
!
crypto ipsec client ezvpn EZVPN_CLIENT
 group EZVPN key CISCO
 connect acl EZVPN_TRAFFIC
 mode network-plus
 peer 136.1.18.1
 xauth userid mode interactive
!
interface FastEthernet0/1
 crypto ipsec client ezvpn EZVPN_CLIENT inside
!
interface FastEthernet0/0.38
 crypto ipsec client ezvpn EZVPN_CLIENT outside
```

Verification

Start the EzVPN client debugging and generate interesting traffic toward R1's Loopback0 subnet. Note that traffic is allowed also prior to the tunnel being established, and it will be functional as long as routing exists:

```
R3#debug crypto ipsec client ezvpn
EzVPN debugging is on
```

https://t.me/learningnets



Most of the debug outputs are identical to ones from EzVPN client mode. First interesting traffic is detected, then the EzVPN session is initiated, and then the XAUTH request is received from the EzVPN server:

```
EZVPN(EZVPN_CLIENT): Current State: CONNECT_REQUIRED
EZVPN(EZVPN_CLIENT): Event: CONNECT
EZVPN(EZVPN_CLIENT): ezvpn_connect_request
EZVPN(EZVPN_CLIENT): Found valid peer 136.1.18.1
EZVPN(EZVPN_CLIENT): Added PSK for address 136.1.18.1

EzVPN(EZVPN_CLIENT): sleep jitter delay 1151
EZVPN(EZVPN_CLIENT): New State: READY
EZVPN(EZVPN_CLIENT): Current State: READY
EZVPN(EZVPN_CLIENT): Event: IKE_PFS
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): Current State: READYEZVPN(EZVPN_CLIENT): Event: CONN_UP
EZVPN(EZVPN_CLIENT): ezvpn_conn_up 05D36600 EBD232BC A848FF64 B2ED50BC
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): Current State: READYEZVPN(EZVPN_CLIENT): Event: XAUTH_REQUEST

EZVPN(EZVPN_CLIENT): ezvpn_xauth_request
EZVPN(EZVPN_CLIENT): ezvpn_parse_xauth_msg
```

The EzVPN Remote asks for a username and password to be entered:

```
EZVPN: Attributes sent in xauth request message:
        XAUTH_USER_NAME_V2(EZVPN_CLIENT):
        XAUTH_USER_PASSWORD_V2(EZVPN_CLIENT):
EZVPN(EZVPN_CLIENT): Installing auth_proxy in interface Fa0/1
EZVPN(EZVPN_CLIENT): New State: XAUTH_REQ
EZVPN: Checking Connect ACL
EZVPN(EZVPN_CLIENT): Pending XAuth Request, Please enter the following command:
EZVPN: crypto ipsec client ezvpn xauthEZVPN(EZVPN_CLIENT): Requesting following info for xauth
```

https://t.me/learningnets

```
username:(Null)
 password:(Null)

        message:Please enter username and password
EZVPN(EZVPN_CLIENT): Caching Xauth request for future use
```

This time, credentials were entered via CLI using the command **`crypto ipsec client ezvpn xauth`**:

```
R3#crypto ipsec client ezvpn xauth
Username: CISCO@EZVPN
Password: CISCO

R1:
EZVPN(EZVPN_CLIENT): Current State: XAUTH_PROMPT
EZVPN(EZVPN_CLIENT): Event: XAUTH_REQ_INFO_READY
EZVPN(EZVPN_CLIENT): ezvpn_xauth_reply XAUTH_USER_NAME_V2(EZVPN_CLIENT): CISCO@EZVPN
 XAUTH_USER_PASSWORD_V2(EZVPN_CLIENT): <omitted>
EZVPN(EZVPN_CLIENT): New State: XAUTH_REPLIED
EZVPN(EZVPN_CLIENT): Current State: XAUTH_REPLIED
EZVPN(EZVPN_CLIENT): Event: XAUTH_STATUS EZVPN(EZVPN_CLIENT): xauth status received: Success
EZVPN(EZVPN_CLIENT): New State: READY
EZVPN(EZVPN_CLIENT): Current State: READY
EZVPN(EZVPN_CLIENT): Event: MODE_CONFIG_REPLY 05D36600 EBD232BC A848FF64 B2ED50BC 05D36600 EBD232BC A848FF64 B2ED50B
EZVPN(EZVPN_CLIENT): ezvpn_parse_mode_config_msg
EZVPN: Attributes sent in message:
        Address: 20.0.0.25
        Peer has  IPsec Interface support
        Split Tunnel List: 1
              Address    : 150.1.1.1
              Mask       : 255.255.255.255
              Protocol   : 0x0
              Source Port: 0
              Dest Port  : 0
        Savepwd off
EZVPN: Unknown/Unsupported Attr: APPLICATION_VERSION (0x7)
EZVPN(EZVPN_CLIENT): ezvpn_mode_config
EZVPN(EZVPN_CLIENT): ezvpn_nat_config
EZVPN(EZVPN_CLIENT): New State: SS_OPEN
EZVPN(EZVPN_CLIENT): Current State: SS_OPEN
EZVPN(EZVPN_CLIENT): Event: SOCKET_READY
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): Current State: SS_OPEN
EZVPN(EZVPN_CLIENT): Event: SOCKET_READY
EZVPN(EZVPN_CLIENT): No state change
```

https://t.me/learningnets

```
EZVPN(EZVPN_CLIENT): Current State: SS_OPENEZVPN(EZVPN_CLIENT): Event: SOCKET_UP
```

The last step indicates that the IPsec SA is successfully established, but most importantly, the NEM subnet has been advertised to the server and the server assigned an IP address to the EzVPN Remote, so Network Mode Plus is being used:

```
%CRYPTO-6-EZVPN_CONNECTION_UP: (Client)  User=  Group=EZVPN  Client_public_addr=136.1.38.3  Server_public_addr=136.1
EZVPN(EZVPN_CLIENT): Disabling Connect ACL checking
EzVPN: Local Traffic Feature Deleted
EZVPN(EZVPN_CLIENT): Tunnel UP! Letting user know about itEZVPN(EZVPN_CLIENT): New State: IPSEC_ACTIVE
1
EZVPN(EZVPN_CLIENT): Current State: IPSEC_ACTIVE
EZVPN(EZVPN_CLIENT): Event: MTU_CHANGED
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): retaining context 14
EZVPN(EZVPN_CLIENT): Current State: IPSEC_ACTIVEEZVPN(EZVPN_CLIENT): Event: SOCKET_UP
ezvpn_socket_up
ezvpn_process_socket_up
EZVPN(EZVPN_CLIENT): No Connect ACL checking status change
EzVPN: Local Traffic Feature Deleted
EZVPN(EZVPN_CLIENT): Tunnel UP! Letting user know about it
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): retaining context 14
EZVPN(EZVPN_CLIENT): Current State: IPSEC_ACTIVE
EZVPN(EZVPN_CLIENT): Event: MTU_CHANGED
EZVPN(EZVPN_CLIENT): No state change
EZVPN(EZVPN_CLIENT): retaining context 14
EzVPN(EZVPN_CLIENT): No URL available to configure
%LINK-3-UPDOWN: Interface Loopback10000, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback10000, changed state to up
```

Verify that although an IP address was assigned to the EzVPN Remote, no NAT configuration is applied for protected traffic:

```
R3#ping 150.1.1.1 source fastEthernet0/1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 136.1.100.3 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R3#show ip nat statistics
```

https://t.me/learningnets

```
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Peak translations: 12, occurred 01:53:24 agoOutside interfaces:
 FastEthernet0/0.38
Inside interfaces:
 FastEthernet0/1

Hits: 388  Misses: 0
CEF Translated packets: 378, CEF Punted packets: 0
Expired translations: 50
Dynamic mappings:
-- Inside Source
[Id: 13] access-list EZVPN_CLIENT_internet-list interface FastEthernet0/0.38 refcount 0

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

Verify the EzVPN Remote status and that the IP address has been assigned; notice that HTTP Intercept mode is no longer displayed:

```
R3#show crypto ipsec client ezvpn
Easy VPN Remote Phase: 8

Tunnel name : EZVPN_CLIENT
Inside interface list: FastEthernet0/1
Outside interface: FastEthernet0/0.38
Connect : ACL based with access-list EZVPN_TRAFFICCurrent State: IPSEC_ACTIVE
Last Event: MTU_CHANGEDAddress: 20.0.0.26 (applied on Loopback10000)

Mask: 255.255.255.255
Save Password: Disallowed
Split Tunnel List: 1
       Address    : 150.1.1.1
       Mask       : 255.255.255.255
       Protocol   : 0x0
       Source Port: 0
       Dest Port  : 0
Current EzVPN Peer: 136.1.18.1
```

Check the static routes injected in the EzVPN server's routing table. Notice the /24 prefix corresponding to the NEM subnet:

```
R1#show ip route static | b Gateway
```

https://t.me/learningnets

```
Gateway of last resort is not set

      20.0.0.0/32 is subnetted, 1 subnets
S        20.0.0.26 is directly connected, Virtual-Access1
      136.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        136.1.100.0/24 is directly connected, Virtual-Access1
```

Check the IPSec SAs installed on the EzVPN Remote side. Notice that there are two IPSec SAs, one for the IP address requested from the server and another for the NEM subnet. Because there are no NAT rules, the client implements split-tunneling using IPSec SAs, only encrypting the relevant traffic. IPSec SAs are generated based on the IP subnets pushed in the split-tunnel access-list:

```
R3#show crypto ipsec sa peer 136.1.18.1

interface: FastEthernet0/0.38
    Crypto map tag: FastEthernet0/0.38-head-0, local addr 136.1.38.3

   protected vrf: (none) local  ident (addr/mask/prot/port): (136.1.100.0/255.255.255.0/0/0)
remote ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   current_peer 136.1.18.1 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
#pkts decaps: 5, #pkts decrypt: 5, #pkts verify: 5
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.38.3, remote crypto endpt.: 136.1.18.1
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.38
     current outbound spi: 0xD745123F(3611628095)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0x13137E4E(320044622)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2019, flow_id: NETGX:19, sibling_flags 80000046, crypto map: FastEthernet0/0.38-head-0
        sa timing: remaining key lifetime (k/sec): (4475539/3114)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE

     inbound ah sas:

     inbound pcp sas:
```

https://t.me/learningnets

```
 outbound esp sas:
      spi: 0xD745123F(3611628095)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2020, flow_id: NETGX:20, sibling_flags 80000046, crypto map: FastEthernet0/0.38-head-0
        sa timing: remaining key lifetime (k/sec): (4475539/3114)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE

     outbound ah sas:

     outbound pcp sas:

   protected vrf: (none) local  ident (addr/mask/prot/port): (20.0.0.26/255.255.255.255/0/0)
 remote ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   current_peer 136.1.18.1 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 0, #pkts encrypt: 0, #pkts digest: 0
 #pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.38.3, remote crypto endpt.: 136.1.18.1
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.38
     current outbound spi: 0x6284FBC6(1652882374)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0x432D90D(70441229)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2021, flow_id: NETGX:21, sibling_flags 80000046, crypto map: FastEthernet0/0.38-head-0
        sa timing: remaining key lifetime (k/sec): (4438152/3114)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE

     inbound ah sas:

     inbound pcp sas:
 outbound esp sas:
      spi: 0x6284FBC6(1652882374)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2022, flow_id: NETGX:22, sibling_flags 80000046, crypto map: FastEthernet0/0.38-head-0
        sa timing: remaining key lifetime (k/sec): (4438152/3114)
        IV size: 8 bytes
```

https://t.me/learningnets

```
    replay detection support: Y Status: ACTIVE


  outbound ah sas:

  outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Remote in NEM Plus Mode with DVTI

Task

- Reconfigure the EzVPN server and EzVPN client from the previous task so that R1 learns about remote subnets of EzVPN Remote in NEM mode using OSPF.
- Create a new Loopback100 interface in R3 with the subnet 150.1.100.0/24 and ensure that the traffic from this subnet to R1's Loopback0 is encrypted.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Virtual tunnel interfaces encrypt all traffic routed over a VTI, allowing for dynamic routing with IPSec VPNs. This simplifies network design; you may have separate routing processes for all VPN networks and control traffic encryption by proper route injection. Using VTIs with EzVPN Remote allows the central server and the remote client to exchange complete routing information. For example, if the remote site has multiple routers and runs an IGP, it would be flexible to let the remote router advertise internal networks to the central site over a VTI and vice-versa. This would allow for dynamic discovery of new protected subnets, without the need to push a split ACL down to the client.

In our scenario, we configure R3 to use a VTI. The configuration is straightforward, but the implications are really important. You can dynamically advertise any subnet

https://t.me/learningnets

over the VTI and the relevant traffic will be encrypted. In our case, we advertise the new Loopback100 interface of R3 into OSPF. Notice that R3 uses Network Plus mode. This will force R1 to create two static routes for the NEM subnet and the allocated IP address via the virtual-access interface.

When OSPF is enabled on an IP unnumbered configured interface, it does not require that the remote neighbor be in the same subnet. So we enable OSPF on R1's Loopback0 interface, which is used for the DVTI, and R1 will become OSPF neighbor with R3's dynamically assigned Loopback interface, which will have an IP address from 20.0.0.0/24 subnet. Our example still uses split-access list configuration. This will force R3 to install a static route for remote network pointing to the virtual-access interface instance. In most cases, you no longer need split-access lists with VTI's and dynamic routing, so you may remove the relevant configuration and ensure that the protected network is advertised into the VPN routing protocol.

Configuration

R1:

```
aaa new-model
aaa authentication login EZVPN local
aaa authorization network EZVPN local
!
username CISCO@EZVPN password CISCO
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
ip local pool EZVPN 20.0.0.1 20.0.0.254
crypto isakmp client configuration address-pool local EZVPN
!
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 group-lock
 key CISCO
 pool EZVPN
 acl EZVPN_TUNNEL
!
crypto isakmp profile CRYPTO_MAP
```

https://t.me/learningnets

```
 match identity group EZVPN
 client authentication list EZVPN
 isakmp authorization list EZVPN
 client configuration address respond
 client configuration group EZVPN
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto ipsec profile EZVPN
 set isakmp-profile CRYPTO_MAP
 set transform-set 3DES_MD5
 set reverse-route tag 1
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile EZVPN
!
router ospf 100
 network 150.1.1.0 0.0.0.255 area 0
```

R3:

```
ip access-list extended EZVPN_TRAFFIC
 permit ip 136.1.100.0 0.0.0.255 host 150.1.1.1
!
interface virtual-template1 type tunnel
 tunnel mode ipsec ipv4
!
crypto ipsec client ezvpn EZVPN_CLIENT
 group EZVPN key CISCO
 connect acl EZVPN_TRAFFIC
 mode network-plus
 peer 136.1.18.1
 xauth userid mode interactive
 virtual-interface 1
!
interface FastEthernet0/1
 crypto ipsec client ezvpn EZVPN_CLIENT inside
!
interface FastEthernet0/0.38
 crypto ipsec client ezvpn EZVPN_CLIENT outside
!
interface Loopback100
```

https://t.me/learningnets

```
 ip address 150.1.100.3 255.255.255.0
 ip ospf network point-to-point
!
router ospf 100
 network 20.0.0.0 0.0.0.255 area 0
 network 150.1.100.0 0.0.0.255 area 0
```

So technically, the following commands are not required, but were kept so that we could see that split-tunnel ACL from EzVPN server is ignored, and interesting traffic ACL from EzVPN client as well, because the session was manually initiated. Also, because dynamic routing was used, R1 on the EzVPN server is not required:

```
R1:
ip access-list extended EZVPN_TUNNEL
 permit ip host 150.1.1.1 any
!
crypto isakmp client configuration group EZVPN
 acl EZVPN_TUNNEL
!
crypto ipsec profile EZVPN
 set reverse-route tag 1


R3:

ip access-list extended EZVPN_TRAFFIC
 permit ip 136.1.100.0 0.0.0.255 host 150.1.1.1
!
crypto ipsec client ezvpn EZVPN_CLIENT
 connect acl EZVPN_TRAFFIC
```

Verification

Manually start the EzVPN tunnel using the command **`crypto ipsec client ezvpn connect`** on the EzVPN Remote and **`crypto ipsec client ezvpn xauth`** when prompted for XAUTH:

https://t.me/learningnets

```
R3#crypto ipsec client ezvpn connect
EZVPN(EZVPN_CLIENT): Pending XAuth Request, Please enter the following command:
EZVPN: crypto ipsec client ezvpn xauth

R3#crypto ipsec client ezvpn xauth
Username: CISCO@EZVPN
Password:
```

When the tunnel is UP, verify the static routes injected on both server and client, and also verify that OSPF adjacency is established:

```
R3#show ip route static | b Gateway
Gateway of last resort is not set

      150.1.0.0/16 is variably subnetted, 6 subnets, 2 masks
S        150.1.1.1/32 [1/0] via 0.0.0.0, Virtual-Access1
!
!R1#show ip route static | b Gateway
Gateway of last resort is not set

      20.0.0.0/32 is subnetted, 1 subnets
S        20.0.0.29 [1/0] via 0.0.0.0, Virtual-Access1
      136.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        136.1.100.0/24 [1/0] via 0.0.0.0, Virtual-Access1
!
!R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
150.1.3.3         0   FULL/  -        00:00:33    20.0.0.29       Virtual-Access1
!
!R1#show ip route ospf | b Gateway
Gateway of last resort is not set

      150.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
O        150.1.100.0/24 [110/2] via 20.0.0.29, 00:09:05, Virtual-Access1
```

Verify that regardless of the split-tunneling configuration, the proxy-ACL negotiated is of type `permit ip any any` , because of the VTI configuration used:

```
R1#show crypto session
```

https://t.me/learningnets

```
Crypto session current status

Interface: Virtual-Access1
Username: CISCO@EZVPN
Profile: CRYPTO_MAP
Group: EZVPN
Assigned address: 20.0.0.29 Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
 IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0

        Active SAs: 2, origin: crypto map
```

Verify that traffic for R3's Loopback100 interface is sent over the tunnel:

```
R1#clear crypto sa counters
!                   R1#show crypto ipsec sa peer 136.1.38.3 | i pkts
 #pkts encaps: 1, #pkts encrypt: 1, #pkts digest: 1
 #pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0
     #pkts compressed: 0, #pkts decompressed: 0
     #pkts not compressed: 0, #pkts compr. failed: 0
     #pkts not decompressed: 0, #pkts decompress failed: 0
!R1#ping 150.1.100.3 repeat 50
Type escape sequence to abort.
Sending 50, 100-byte ICMP Echos to 150.1.100.3, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (50/50), round-trip min/avg/max = 1/2/4 ms
!R1#show crypto ipsec sa peer 136.1.38.3 | i pkts
 #pkts encaps: 55, #pkts encrypt: 55, #pkts digest: 55
 #pkts decaps: 51, #pkts decrypt: 51, #pkts verify: 51

     #pkts compressed: 0, #pkts decompressed: 0
     #pkts not compressed: 0, #pkts compr. failed: 0
     #pkts not decompressed: 0, #pkts decompress failed: 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 IOS EzVPN Remote in NEM Mode with PKI

Task

- Configure R1 as an Easy VPN server with the following requirements:
  - Use 3DES/MD5 as cipher/hash for both IPSec phases.
  - Use digital certificates issued by R1 for authentication.
  - Allow remote clients identified by group name **EZVPN** to connect.
  - Enable Xauth against the local user database.
- Configure R3 as EzVPN Remote in Network Extension Mode and manually initiate the tunnel from R3's CLI.
- R1 and R3 should establish an EIGRP adjacency over the IPsec tunnel.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

Using digital signatures for authentication ensures security and scalability for large VPN deployments. Earlier implementation of EzVPN Remote did not support authentication based on digital signatures. Configuring this authentication method is similar to configuring L2L IPsec tunnel with digital signatures. You need to remove the group command under the EzVPN profile in EzVPN Remote router, after you configured CA trustpoint and enrolled with it. However, the question is, how would the EzVPN server determine the proper EzVPN configuration group if the client no longer provides the group name?

There are two solutions to this problem:

- The first one, which is the default method, finds the group name by matching the value of **OU** field from the client's certificate subject name. For example, if certificate **OU** equals **EZVPN**, the EzVPN server will look for EzVPN configuration group named **EZVPN**. As you remember, every ISAKMP profile should have a match statement so that the system can map remote peers to a profile.
- Instead of group names, it is possible to match the hostname identity value (or the domain name) or any other certificate attributes when using digital certificates for authentication.

We must ensure that the EzVPN Remote has a certificate with **OU** information present and that its value equals the configured group name of **EZVPN**. For the EzVPN server configuration, notice the following:

- There is no need to configure a split-tunnel ACL or pre-shared-key.
- There is no need to use a pool because we use Network Extension Mode, so there is no need to allocate an IP address to the EzVPN Remote.
- Because there is no dynamic IP addressing involved, the AAA authorization list is no longer required.

For the EzVPN Remote configuration, notice the following:

- DVTI interface has been configured with `ip unnumbered` , so EIGRP can run over the tunnel interface and establish adjacency.
- There is no need to configure an ACL for interesting traffic because the tunnel is manually initiated.

Configuration

Note that the hostname and domain names required to generate the RSA key pair are preconfigured from the initial configurations, as well as NTP synchronization on all devices with R1. On both routers, we simply configure the PKI enrollment and

https://t.me/learningnets

enable certificate authentication.

R1:

```
crypto pki trustpoint R1-CLIENT
 enrollment url http://150.1.1.1
 subject-name CN=R1.ine.com
 serial-number none
 ip-address none
!
crypto pki authenticate R1-CLIENT
crypto pki enroll R1-CLIENT
!
aaa new-model
aaa authentication login EZVPN local
!
username CISCO password CISCO
!
crypto isakmp policy 100
 authentication rsa-sig
 encryption 3des
 hash md5
 group 2
!
crypto isakmp profile CRYPTO_MAP
 match identity group EZVPN
 client authentication list EZVPN
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto ipsec profile EZVPN
 set isakmp-profile CRYPTO_MAP
 set transform-set 3DES_MD5
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile EZVPN
```

R3:

```
crypto pki trustpoint R1
 enrollment url http://150.1.1.1
```

https://t.me/learningnets

```
 subject-name CN=R1.ine.com, OU=EZVPN
 serial-number none
 ip-address none
!
crypto pki authenticate R1
crypto pki enroll R1
!
crypto ipsec client ezvpn EZVPN_CLIENT
 mode network-extension
 peer 136.1.18.1
 xauth userid mode interactive
 virtual-interface 1
!
interface FastEthernet0/1
 crypto ipsec client ezvpn EZVPN_CLIENT inside
!
interface FastEthernet0/0.38
 crypto ipsec client ezvpn EZVPN_CLIENT outside
!
interface virtual-template1 type tunnel
 ip unnumbered Loopback0
 tunnel mode ipsec ipv4
```

Verification

Enable ISAKMP debugging and manually initiate the EzVPN tunnel:

```
R3#debug crypto isakmp
Crypto ISAKMP debugging is onR3#crypto ipsec client ezvpn connect
R3#crypto ipsec client ezvpn xauth
Username: CISCO
Password: CISCO
```

At some point in the Phase1 negotiation process, before Phase1.5, the following messages appear, indicating that based on the **OU** value present in the certificate, the incoming connection is mapped to the **EZVPN** group:

```
ISAKMP (1045): ID payload
        next-payload : 6
        type         : 9
        Dist. name   : hostname=R3.ine.com,cn=R1.ine.com,ou=EZVPN
        protocol     : 17
        port         : 500
```

https://t.me/learningnets

```
        length       : 84ISAKMP:(0):: UNITY's identity group: OU = EZVPN
ISAKMP:(0):: peer matches CRYPTO_MAP profile

ISAKMP:(1045):Setting client config settings 3004FD58
ISAKMP:(1045):(Re)Setting client xauth list  and state
ISAKMP/xauth: initializing AAA request
ISAKMP:(1045): Profile CRYPTO_MAP assigned peer the group named EZVPN
```

Verify that the IPsec tunnel is using digital certificate authentication:

```
R3#show crypto isakmp sa detail
Codes: C - IKE configuration mode, D - Dead Peer Detection
       K - Keepalives, N - NAT-traversal
       T - cTCP encapsulation, X - IKE Extended Authentication
       psk - Preshared key, rsig - RSA signature
       renc - RSA encryption
IPv4 Crypto ISAKMP SA

C-id  Local           Remote          I-VRF    Status Encr Hash Auth DH Lifetime Cap.
1007  136.1.38.3      136.1.18.1ACTIVE 3des md5  rsig 2
  3w3d     CX
       Engine-id:Conn-id =  SW:7

IPv6 Crypto ISAKMP SA
```

Verify that packets are both encrypted and decrypted over the tunnel (EIGRP packets because there are no prefixes learned through the tunnel EIGRP adjacency):

```
R1#show crypto session detail
Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation
 Interface: Virtual-Access1
Username: CISCO
Profile: CRYPTO_MAP
Uptime: 00:05:03Session status: UP-ACTIVE
Peer: 136.1.38.3 port 500 fvrf: (none) ivrf: (none)
Phase1_id: hostname=R3.ine.com,cn=R1.ine.com,ou=EZVPN
      Desc: (none)
  IKEv1 SA: local 136.1.18.1/500 remote 136.1.38.3/500 Active
          Capabilities:CX connid:1045 lifetime:23:54:42
```

https://t.me/learningnets

```
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0
        Active SAs: 2, origin: crypto map Inbound:  #pkts dec'ed 42 drop 0 life (KB/Sec) 4240505/3296
Outbound: #pkts enc'ed 107 drop 0 life (KB/Sec) 4240498/3296
```

Verify the EIGRP adjacency over the tunnel:

```
R1#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(111)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   136.1.38.3              Vi1                      13 00:00:51    1   100  0  32

1   136.1.18.8              Gi0/0                    11 19:39:31    1   100  0  36
0   136.1.18.2              Gi0/0                    12 19:40:30   16   100  0  20
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 ASA EzVPN Server with PSK

Task

- Configure ASA4 to accept remote VPN connections from Cisco Easy VPN Clients using a group named **EZVPN-ASA**:
  - Use 3DES/MD5 as the cipher/hash for IPSec Phase1/Phase2.
  - Use address pool 21.0.0.0/24 to allocate IP addresses for remote clients, and push the DNS server IP address of 136.1.100.101 to the clients.
  - Allow for split-tunneling to network 136.1.18.0/24.
  - Remote users should be authenticated using the name **CISCO-ASA** and the password **CISCO1234**.
  - Ensure that this user is only allowed to log in under the group **EZVPN-ASA**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Configuration

ASA4:

```
crypto isakmp enable VLAN38
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
```

https://t.me/learningnets

```
!
vpn-addr-assign local
ip local pool EZVPN-ASA 21.0.0.1-21.0.0.254
!
access-list SPLIT_TUNNEL standard permit 136.1.18.0 255.255.255.0
!
username CISCO-ASA password CISCO1234
username CISCO-ASA attributes
 group-lock value EZVPN-ASA
!
group-policy EZVPN-ASA internal
group-policy EZVPN-ASA attributes
 vpn-tunnel-protocol IPSec
 address-pools value EZVPN-ASA
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value SPLIT_TUNNEL
 dns-server value 136.1.100.101
!
tunnel-group EZVPN-ASA type remote-access
tunnel-group EZVPN-ASA general-attributes
  authentication-server-group LOCAL
  default-group-policy EZVPN-ASA
!
tunnel-group EZVPN-ASA ipsec-attributes
  pre-shared-key CISCO
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
crypto dynamic-map DYNAMIC 100 set transform-set 3DES_MD5
crypto dynamic-map DYNAMIC 100 set reverse-route
crypto map VPN 100 ipsec-isakmp dynamic DYNAMIC
crypto map VPN interface VLAN38
!
router eigrp 111
 redistribute static metric 100000 100 255 1 1500
```

The configuration above is for ASA code 8.2.x. Beginning with code 8.3, multiple IPsec-related commands have changed with the addition of IKEv2. You can still use the old commands, although hidden. ASA will automatically convert to the new format:

ASA4:

```
crypto ikev1 enable VLAN38
crypto ikev1 policy 100
 authentication pre-share
```

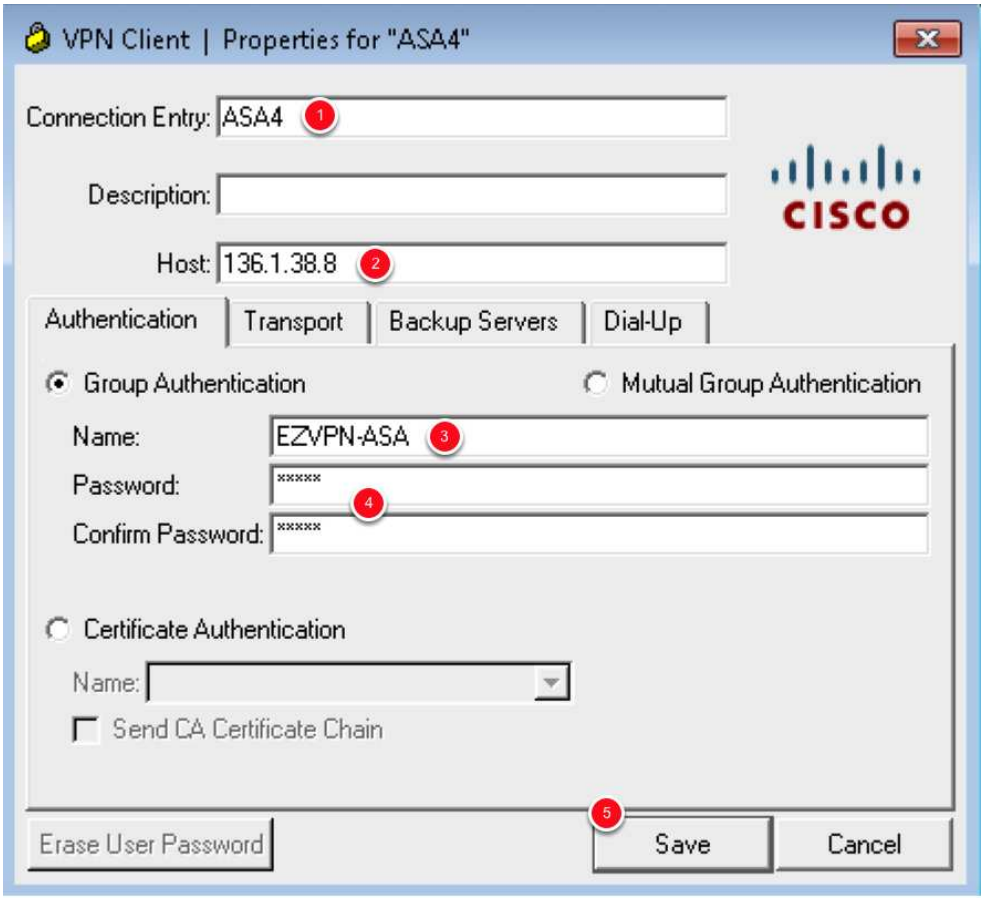
https://t.me/learningnets

```
 encryption 3des
 hash md5
 group 2
!
vpn-addr-assign local
ip local pool EZVPN-ASA 21.0.0.1-21.0.0.254
!
access-list SPLIT_TUNNEL standard permit 136.1.18.0 255.255.255.0
!
username CISCO-ASA password CISCO1234
username CISCO-ASA attributes
 group-lock value EZVPN-ASA
!
group-policy EZVPN-ASA internal
group-policy EZVPN-ASA attributes
 vpn-tunnel-protocol ikev1
 address-pools value EZVPN-ASA
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value SPLIT_TUNNEL
 dns-server value 136.1.100.101
!
tunnel-group EZVPN-ASA type remote-access
tunnel-group EZVPN-ASA general-attributes
  authentication-server-group LOCAL
  default-group-policy EZVPN-ASA
!
tunnel-group EZVPN-ASA ipsec-attributes
  ikev1 pre-shared-key CISCO
!
crypto ipsec ikev1 transform-set 3DES_MD5 esp-3des esp-md5-hmac
crypto dynamic-map DYNAMIC 100 set ikev1 transform-set 3DES_MD5
crypto dynamic-map DYNAMIC 100 set reverse-route
crypto map VPN 100 ipsec-isakmp dynamic DYNAMIC
crypto map VPN interface VLAN38
!
router eigrp 111
 redistribute static metric 100000 100 255 1 1500
```

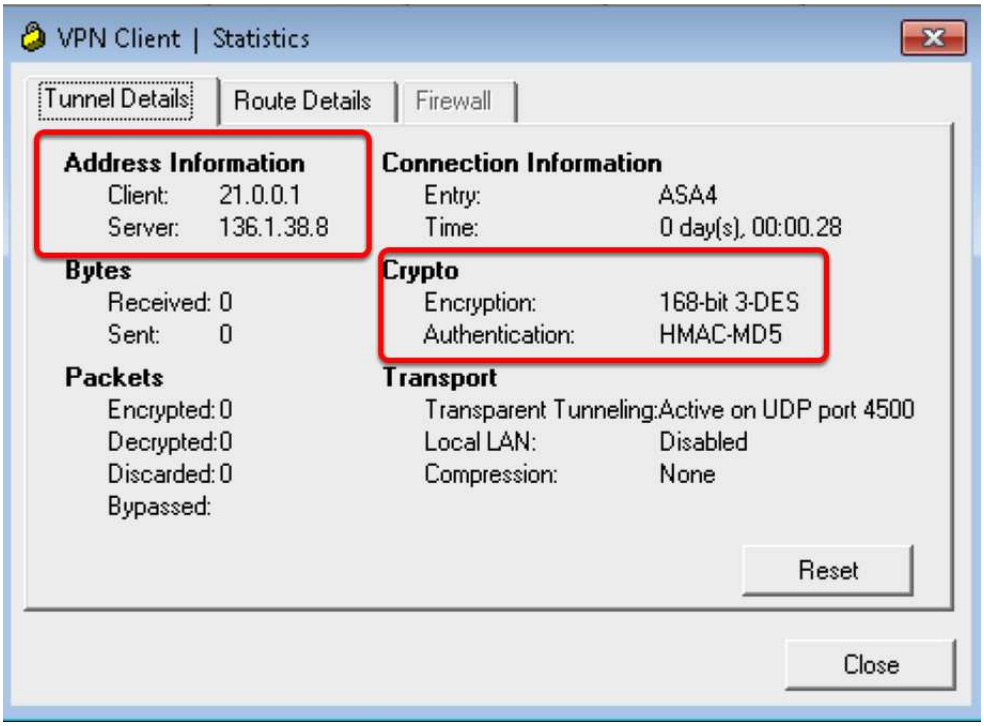
Verification

Connect to **TEST-PC-A** and create a new connection for ASA4:

https://t.me/learningnets

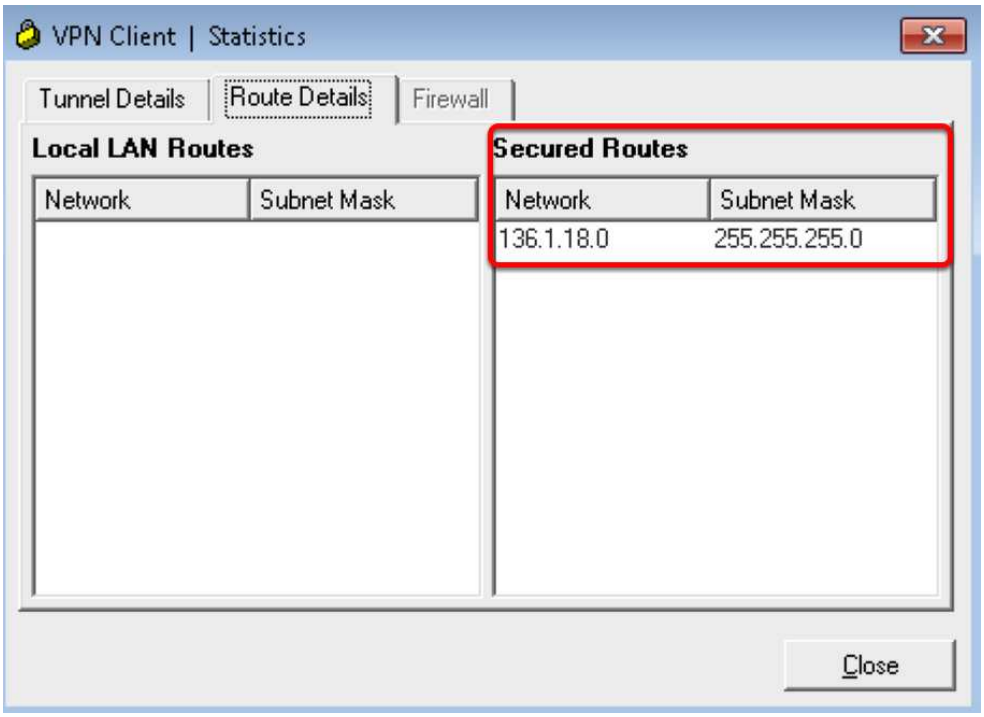


Issue one `debug crypto ikev1 64` on the ASA and initiate the connection from the **TEST-PC-A**. When the connection is established, review the Cisco VPN Client statistics window to inspect the negotiated IP address and algorithms:

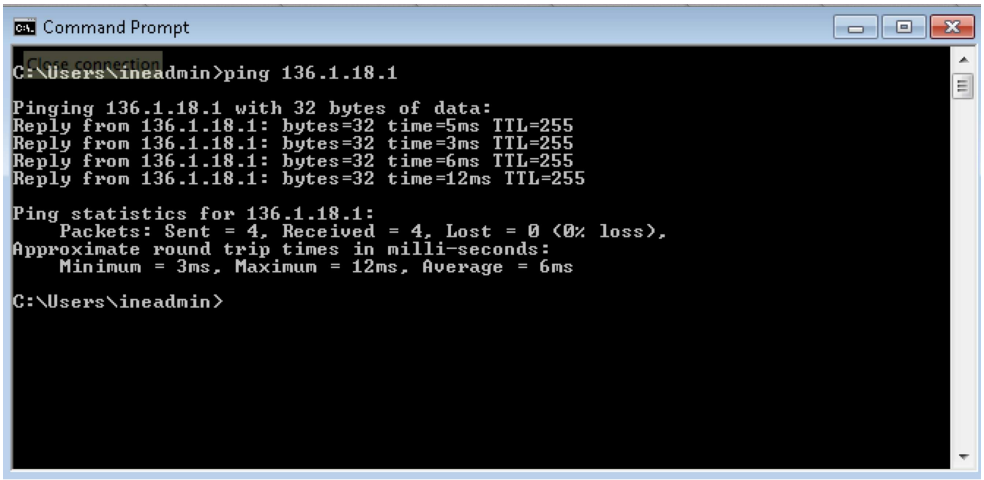


https://t.me/learningnets

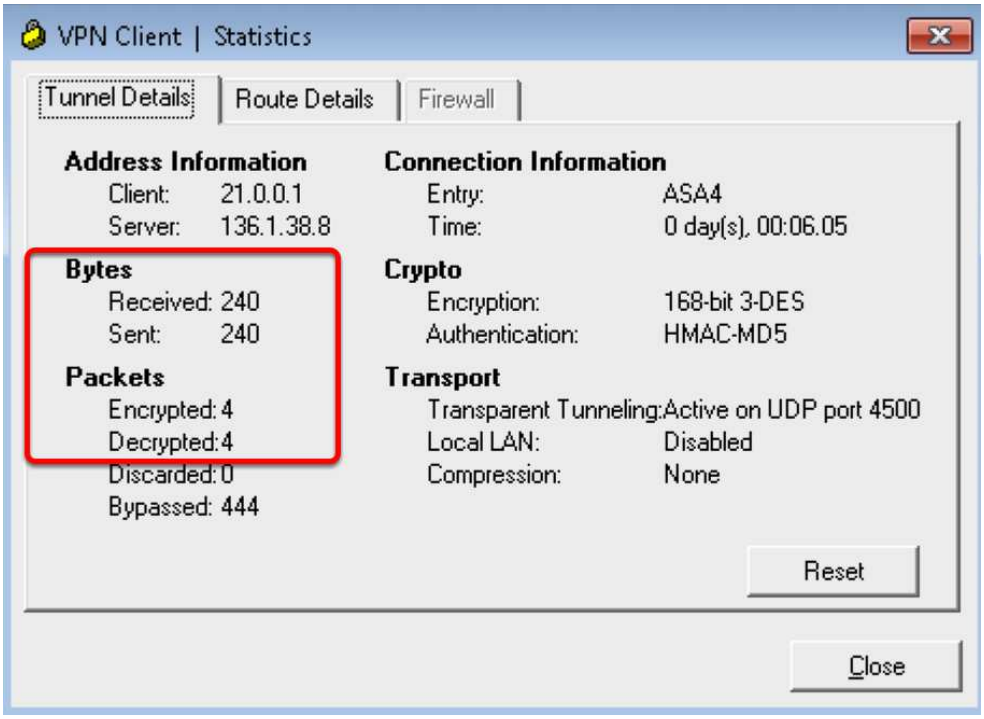
Also review the **Route details** tab to inspect the split-tunneling configuration:



Generate traffic over the tunnel and verify packet statistics again:



https://t.me/learningnets



Verify the VPN connection on the ASA:

https://t.me/learningnets

```
ASA4#show vpn-sessiondb ra-ikev1-ipsec filter name CISCO-ASA


Session Type: IKEv1 IPsec
Username     : CISCO-ASA
              Index        : 4Assigned IP  : 21.0.0.1                Public IP    : 136.1.100.100

Protocol     : IKEv1 IPsec
License      : Other VPNEncryption   : 3DES                    Hashing      : MD5

Bytes Tx     : 0                      Bytes Rx     : 0
Group Policy : EZVPN-ASA              Tunnel Group : EZVPN-ASA

Login Time   : 02:58:12 UTC Mon May 13 2013
Duration     : 0h:00m:25s
Inactivity   : 0h:00m:00s
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none
!
!
ASA4#show crypto ikev1 sa


IKEv1 SAs:

   Active SA: 1
    Rekey SA: 0 (A tunnel will report 1 Active and 1 Rekey SA during rekey)
Total IKE SA: 1
1   IKE Peer: 136.1.100.100

    Type    : user            Role    : responder      Rekey   : noState   : AM_ACTIVE
```

On the ASA side, during the connection process, it receives an IKE AM message with the group ID, based on which the EzVPN server selects the tunnel-group **EZVPN-ASA**:

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=0) with payloads : HDR + SA (1) + KE (4) + NONCE (10)
[IKEv1 DEBUG]IP = 136.1.100.100, processing SA payload
[IKEv1 DEBUG]IP = 136.1.100.100, processing ke payload
[IKEv1 DEBUG]IP = 136.1.100.100, processing ISA_KE payload
[IKEv1 DEBUG]IP = 136.1.100.100, processing nonce payload
[IKEv1 DEBUG]IP = 136.1.100.100, processing ID payload
[IKEv1 DEBUG]IP = 136.1.100.100, processing VID payload
```

https://t.me/learningnets

```
[IKEv1 DEBUG]IP = 136.1.100.100, Received xauth V6 VID
[IKEv1 DEBUG]IP = 136.1.100.100, processing VID payload
[IKEv1 DEBUG]IP = 136.1.100.100, Received DPD VID
[IKEv1 DEBUG]IP = 136.1.100.100, processing VID payload
[IKEv1 DEBUG]IP = 136.1.100.100, Received Fragmentation VID
[IKEv1 DEBUG]IP = 136.1.100.100, IKE Peer included IKE fragmentation capability flags:  Main Mode:        True  Aggr
[IKEv1 DEBUG]IP = 136.1.100.100, processing VID payload
[IKEv1 DEBUG]IP = 136.1.100.100, Received NAT-Traversal ver 02 VID
[IKEv1 DEBUG]IP = 136.1.100.100, processing VID payload
[IKEv1 DEBUG]IP = 136.1.100.100, Received Cisco Unity client VID
[IKEv1]IP = 136.1.38.3, Connection landed on tunnel_group EZVPN-ASA
```

At this point, the EzVPN server attempts to find a matching entry in the local ISAKMP policy database. After this, the server performs authentication and generates a response packet:

```
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing IKE SA payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, IKE SA Proposal # 1, Transform # 10 acceptable  Matches global I

[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing ISAKMP SA payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing ke payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing nonce payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Generating keys for Responder...
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing ID payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Computing hash for ISAKMP
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing Cisco Unity VID payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing xauth V6 VID payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing dpd vid payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing NAT-Traversal VID ver 02 payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing NAT-Discovery payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, computing NAT Discovery hash
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing NAT-Discovery payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, computing NAT Discovery hash
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing Fragmentation VID + extended capabilities payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing VID payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Send Altiga/Cisco VPN3000/Cisco ASA GW VID
```

Notice the rich payload in the response packet. In addition to the selected SA, it contains the key exchange payload and various VENDOR payloads, plus the NAT-T payload:

https://t.me/learningnets

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE SENDING Message (msgid=0) with payloads : HDR + SA (1) + KE (4) + NONCE (10) +
```

The client responds, ending the IKE AM Phase 1. At this point, both devices may determine if any is behind NAT:

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=0) with payloads : HDR + HASH (8) + NOTIFY (11) + NAT-
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Computing hash for ISAKMP
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing notify payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing NAT-Discovery payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, computing NAT Discovery hash
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing NAT-Discovery payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, computing NAT Discovery hash
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing VID payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Processing IOS/PIX Vendor ID payload (version: 1.0.0, capabiliti
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, processing VID payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Received Cisco Unity client VID
[IKEv1]Group = EZVPN-ASA, IP = 136.1.100.100, Automatic NAT Detection Status:     Remote end IS NOT behind a NAT dev

[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing blank hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, constructing qm hash payload
```

The XAUTH phase takes place at this point:

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE SENDING Message (msgid=a62bf863) with payloads : HDR + HASH (8) + ATTR (14) +
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=a62bf863) with payloads : HDR + HASH (8) + ATTR (14) +
```

After XAUTH is successful, the EzVPN server determines what Mode Config attributes will be sent to the client based on the group policy configuration:

```
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, process_attr(): Enter!
[IKEv1 DEBUG]Group = EZVPN-ASA, IP = 136.1.100.100, Processing MODE_CFG Reply attributes.
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: primary DNS = 136.1.
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: secondary DNS = clea
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: primary WINS = clear
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: secondary WINS = cle
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: split tunneling list
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: IP Compression = dis
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: Split Tunneling Poli
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: Browser Proxy Settin
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKEGetUserAttributes: Browser Proxy Bypass
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, User (CISCO-ASA) authenticated.
```

https://t.me/learningnets

```
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing blank hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing qm hash payload
```

The next two IKE packets are for XAUTH successful confirmation from the EzVPN server and XAUTH acknowledgment from the EzVPN client:

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE SENDING Message (msgid=cf218734) with payloads : HDR + HASH (8) + ATTR (14) +
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=cf218734) with payloads : HDR + HASH (8) + ATTR (14) +
```

At this point, the EzVPN client requests Mode Config attributes from the EzVPN server:

```
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, process_attr(): Enter!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Processing cfg ACK attributes
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=acf7a701) with payloads : HDR + HASH (8) + ATTR (14) +
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, process_attr(): Enter!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Processing cfg Request attributes
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for IPV4 addres
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for IPV4 net ma
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for DNS server
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for WINS server
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Received unsupported transaction mode attribute:
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Banner!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Save PW set
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Default Dom
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Split Tunne

[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Split DNS!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for PFS setting
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Client Brow
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for backup ip-s
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Client Smar
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for Application
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Client Type: WinNT  Client Application Version:
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for FWTYPE!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, MODE_CFG: Received request for DHCP hostna
```

The EzVPN server allocates an IP address from the configured pool and sends it to the client, along with all other attributes; the client also negotiates DPD:

```
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Obtained IP addr (21.0.0.1) prior to initi
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Assigned private IP address 21.0.0.1 to remote u
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing blank hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Send Client Browser Proxy Attributes!
```

https://t.me/learningnets

```
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Browser Proxy set to No-Modify. Browser Pr
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Send Cisco Smartcard Removal Disconnect en
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing qm hash payload
[IKEv1]IP = 136.1.100.100, IKE_DECODE SENDING Message (msgid=acf7a701) with payloads : HDR + HASH (8) + ATTR (14) +
[IKEv1 DECODE]IP = 136.1.100.100, IKE Responder starting QM: msg id = d46883e5
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Delay Quick Mode processing, Cert/Trans Ex
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Resume Quick Mode processing, Cert/Trans E
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, PHASE 1 COMPLETED
[IKEv1]IP = 136.1.100.100, Keep-alive type for this connection: DPD

[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Starting P1 rekey timer: 82080 seconds.
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, sending notify message
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing blank hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing qm hash payload
[IKEv1]IP = 136.1.100.100, IKE_DECODE SENDING Message (msgid=50172a0c) with payloads : HDR + HASH (8) + NOTIFY (11)
```

At this point, the EzVPN client initiates IPsec Phase2 Quick Mode negotiations:

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=d46883e5) with payloads : HDR + HASH (8) + SA (1) + NO
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing SA payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing nonce payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing ID payload
[IKEv1 DECODE]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, ID_IPV4_ADDR ID received
21.0.0.1
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Received remote Proxy Host data in ID Payload:
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing ID payload
[IKEv1 DECODE]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, ID_IPV4_ADDR_SUBNET ID received--0.0.0.0-
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Received local IP Proxy Subnet data in ID Payloa
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, QM IsRekeyed old sa not found by addr
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Selecting only UDP-Encapsulated-Tunnel and
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKE Remote Peer configured for crypto map: DYNAM
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing IPSec SA payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IPSec SA Proposal # 11, Transform # 1 acce

[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKE: requesting SPI!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKE got SPI from key engine: SPI = 0xaa3df
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, oakley constucting quick mode
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing blank hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing IPSec SA payload
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Overriding Initiator's IPSec rekeying duration f
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing IPSec nonce payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing proxy ID
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Transmitting Proxy Id:
      Remote host: 21.0.0.1  Protocol 0  Port 0
```

https://t.me/learningnets

```
      Local subnet:  0.0.0.0  mask 0.0.0.0 Protocol 0  Port 0
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Sending RESPONDER LIFETIME notification to
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, constructing qm hash payload
[IKEv1 DECODE]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKE Responder sending 2nd QM pkt: msg id
```

Here we have the final messages in Quick-Mode exchange. After the SA has been established, the EzVPN server injects a reverse-route in the local routing table:

```
[IKEv1]IP = 136.1.100.100, IKE_DECODE SENDING Message (msgid=d46883e5) with payloads : HDR + HASH (8) + SA (1) + NON
[IKEv1]IP = 136.1.100.100, IKE_DECODE RECEIVED Message (msgid=d46883e5) with payloads : HDR + HASH (8) + NONE (0) to
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, processing hash payload
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, loading all IPSEC SAs
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Generating Quick Mode Key!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, NP encrypt rule look up for crypto map DYN
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Generating Quick Mode Key!
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, NP encrypt rule look up for crypto map DYN
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Security negotiation complete for User (CISCO-AS
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, IKE got a KEY_ADD msg for SA: SPI = 0xfc13
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Pitcher: received KEY_UPDATE, spi 0xaa3df7
[IKEv1 DEBUG]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Starting P2 rekey timer: 27360 seconds.
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, Adding static route for client address: 21.0.0.1
[IKEv1]Group = EZVPN-ASA, Username = CISCO-ASA, IP = 136.1.100.100, PHASE 2 COMPLETED (msgid=d46883e5)
```

Check the injected static route:

```
ASA4# show route VLAN38 21.0.0.1

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set
S    21.0.0.1 255.255.255.255 [1/0] via 136.1.38.3, VLAN38
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 ASA EzVPN Server DHCP Address Allocation

Task

- Configure ASA4 to accept remote VPN connections from Cisco Easy VPN Clients using a group named **EZVPN-ASA**:
  - Use 3DES/MD5 as the cipher/hash for IPSec Phase1/Phase2.
  - Use R1 as the DHCP server for IP address allocation of remote clients from pool 21.0.0.0/24, and push the DNS server IP address of 136.1.100.101 to the clients.
  - Allow for split-tunneling to network 136.1.18.0/24.
  - Remote users should be authenticated using the name **CISCO-ASA** and the password **CISCO1234**.
  - Ensure that this user is only allowed to log in under the group **EZVPN-ASA**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **EzVPN IKEv1 IPsec** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Configuring DHCP for VPN address assignment may be attractive if you already use a centralized address allocation server. The ASA firewall may query the external server for DHCP information and use the response for VPN address assignment. We must configure a DHCP server IP address under the tunnel-group general-attributes using the command `dhcp-server <IP_ADDRESS>` . By default, the firewall will

https://t.me/learningnets

set the giaddr field in DHCP packets to the IP address of the interface used to query the DHCP server. This will restrict address allocation to only one subnet, directly connected to the firewall. To overcome this limitation, you may use the group-policy command `dhcp-network-scope <GIADDR>` to specify the giaddr field in DHCP packets. The server will then select the matching pool based on this value. However, there is a caveat here. The DHCP server will reply with a DHCP packet to the IP address specified in giaddr field. Therefore, we need a way to ensure that the firewall advertises this IP address into IGP and may respond to it. The simplest solution to accomplish this is to create a logical interface using Ethernet VLANs and assign it the IP address of giaddr.

Cisco ASA supports IP address assignment for EzVPN clients based on:

- AAA
- DHCP
- POOL

This is the order of processing, where the locally configured pool is a gateway of last resort. Therefore, to assign IP addresses from DHCP, because no AAA settings are in place, there is no need to disable any of the above options. The default configuration can be visualized with the following command:

```
Rack1ASA4#show running-config all vpn-addr-assign
vpn-addr-assign aaa
vpn-addr-assign dhcp
vpn-addr-assign local reuse-delay 0
```

Configuration

ASA4:

```
crypto isakmp enable VLAN38
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
access-list SPLIT_TUNNEL standard permit 136.1.18.0 255.255.255.0
!
username CISCO-ASA password CISCO1234
username CISCO-ASA attributes
 group-lock value EZVPN-ASA
!
```

https://t.me/learningnets

```
group-policy EZVPN-ASA internal
group-policy EZVPN-ASA attributes
 vpn-tunnel-protocol IPSec
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value SPLIT_TUNNEL
 dns-server value 136.1.100.101
 dhcp-network-scope 21.0.0.8
!
tunnel-group EZVPN-ASA type remote-access
tunnel-group EZVPN-ASA general-attributes
 authentication-server-group LOCAL
 default-group-policy EZVPN-ASA
 dhcp-server 136.1.18.1
!
tunnel-group EZVPN-ASA ipsec-attributes
 pre-shared-key CISCO
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
crypto dynamic-map DYNAMIC 100 set transform-set 3DES_MD5
crypto dynamic-map DYNAMIC 100 set reverse-route
crypto map VPN 100 ipsec-isakmp dynamic DYNAMIC
crypto map VPN interface VLAN38
!
interface GigabitEthernet0/1.21
 vlan 21
 nameif EZVPN
 ip address 21.0.0.8 255.255.255.0
!
router eigrp 111
 network 21.0.0.0 255.255.255.0


R1:

ip dhcp pool EZVPN-ASA
 network 21.0.0.0 /24
```

Verification

Before initiating the VPN client connection, issue the following debug on the DHCP server:

https://t.me/learningnets

```
R1#debug ip dhcp server packet
DHCP server packet debugging is on.
```

From the debug enabled on the DHCP server, the DHCP transaction seems complete:

```
DHCPD: DHCPDISCOVER received from client 0063.6973.636f.2d61.3439.332e.3463.6161.2e66.3738.372d.5749.4e37.2d31.362d.
!
DHCPD: Sending DHCPOFFER to client 0063.6973.636f.2d61.3439.332e.3463.6161.2e66.3738.372d.5749.4e37.2d31.362d.564c.4
DHCPD: no option 125
DHCPD: unicasting BOOTREPLY for client a493.4caa.f787 to relay 21.0.0.8.
DHCPD: client's VPN is .
DHCPD: No option 125
!
DHCPD: DHCPREQUEST received from client 0063.6973.636f.2d61.3439.332e.3463.6161.2e66.3738.372d.5749.4e37.2d31.362d.5
DHCPD: Appending system default domain
DHCPD: Using hostname 'WIN7-1' for dynamic update (from hostname option)
!
DHCPD: Sending DHCPACK to client 0063.6973.636f.2d61.3439.332e.3463.6161.2e66.3738.372d.5749.4e37.2d31.362d.564c.414

DHCPD: no option 125
DHCPD: unicasting BOOTREPLY for client a493.4caa.f787 to relay 21.0.0.8.
```

Verify the DHCP binding on the DHCP server.

```
R1#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name21.0.0.5
            0063.6973.636f.2d61.    May 14 2013 05:02 AM    Automatic
                    3439.332e.3463.6161.
                    2e66.3738.372d.5749.
                    4e37.2d31.362d.564c.
                    414e.3138.00
```

Verify that ASA assigned the same IP address to the EzVPN client:

```
ASA4#show vpn-sessiondb ra-ikev1-ipsec filter name CISCO-ASA

Session Type: IKEv1 IPsec
```

https://t.me/learningnets

```
Username     : CISCO-ASA              Index        : 10
Assigned IP  : 21.0.0.5               Public IP    : 136.1.100.100

Protocol     : IKEv1
License      : Other VPN
Encryption   : 3DES                   Hashing      : MD5
Bytes Tx     : 0                      Bytes Rx     : 0
Group Policy : EZVPN-ASA              Tunnel Group : EZVPN-ASA
Login Time   : 05:02:06 UTC Mon May 13 2013
Duration     : 0h:04m:00s
Inactivity   : 0h:00m:00s
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none
```

Verify both IPsec Phases:

```
ASA4#show vpn-sessiondb detail ra-ikev1-ipsec

Session Type: IKEv1 IPsec Detailed

Username     : CISCO-ASA              Index        : 2
Assigned IP  : 21.0.0.5               Public IP    : 136.1.100.100
Protocol     : IKEv1 IPsec
License      : Other VPN
Encryption   : 3DES                   Hashing      : MD5
Bytes Tx     : 240                    Bytes Rx     : 240
Pkts Tx      : 4                      Pkts Rx      : 4
Pkts Tx Drop : 0                      Pkts Rx Drop : 0
Group Policy : EZVPN-ASA              Tunnel Group : EZVPN-ASA
Login Time   : 16:00:52 UTC Tue Oct 8 2013
Duration     : 0h:07m:14s
Inactivity   : 0h:00m:00s
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none

IKEv1 Tunnels: 1
IPsec Tunnels: 1
IKEv1:
  Tunnel ID    : 2.1
  UDP Src Port : 52256                  UDP Dst Port : 500
IKE Neg Mode : Aggressive             Auth Mode    : preSharedKeys
Encryption   : 3DES                   Hashing      : MD5
  Rekey Int (T): 86400 Seconds          Rekey Left(T): 85972 Seconds D/H Group    : 2
  Filter Name  :
  Client OS    : WinNT                  Client OS Ver: 5.0.07.0410
```

https://t.me/learningnets

```
IPsec:

  Tunnel ID    : 2.2 Local Addr   : 0.0.0.0/0.0.0.0/0/0
Remote Addr  : 21.0.0.5/255.255.255.255/0/0
Encryption   : 3DES                  Hashing      : MD5

  Encapsulation: Tunnel
  Rekey Int (T): 28800 Seconds          Rekey Left(T): 28370 Seconds
  Idle Time Out: 30 Minutes             Idle TO Left : 23 Minutes
  Bytes Tx     : 240                    Bytes Rx     : 240
  Pkts Tx      : 4                      Pkts Rx      : 4

NAC:
  Reval Int (T): 0 Seconds              Reval Left(T): 0 Seconds
  SQ Int (T)   : 0 Seconds              EoU Age(T)   : 430 Seconds
  Hold Left (T): 0 Seconds              Posture Token:
  Redirect URL :
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 DMVPN Phase1 with PSK

Task

- Configure DMVPN Phase1 between R1, R2, and R3 as follows:
  - Use R1 as the hub.
  - Protect traffic between VLANs 11, 22, and 33.
  - Use EIGRP 123 as routing protocol.
  - Use 3DES/MD5 as cipher/hash for both IPsec phases.
  - Use DH group2 and a pre-shared-key of **CISCO**.
  - Authenticate NHRP packets using a string of **CISCO**.
  - Use the minimum overhead possible.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.

Overview

Dynamic Multipoint VPN, or DMVPN, is an effective solution for dynamic secure overlay networks. In short, DMVPN is a combination of the following technologies:

- Multipoint GRE (mGRE)
- Next-Hop Resolution Protocol (NHRP)
- Dynamic Routing Protocol (EIGRP, RIP, OSPF, BGP)
- Dynamic IPsec encryption
- Cisco Express Forwarding (CEF)

mGRE

mGRE, the first component of DMVPN, adds functionality to the idea of a tunnel by

https://t.me/learningnets

allowing multiple concurrent destinations for a tunnel. The tunnel endpoints share the same abstract transport medium connecting them, very similar to any NBMA (non-broadcast multiple access) technology such as frame relay or ATM. A special glue technology is needed to map tunnel IP addresses to real or NBMA IP addresses, used by the endpoint routers. This special glue technology is known as Next Hop Resolution Protocol (NHRP), and it serves a function similar to what ARP performs on shared mediums, resolving tunnel IP addresses to their NBMA counterparts. ARP performs IP address to MAC binding.

NHRP

The general idea of NHRP is to have a special next-hop server (NHS) attached to the shared medium. Every other node will register their tunnel-to-NBMA address mapping with the server and query the server when they need to resolve a peer's tunnel address. Note that unlike ARP technology, no broadcast messages are used for resolution, because the medium may not support broadcasting. Therefore, you must provide every node with the NBMA address of the NHS. The use of NHRP allows dynamic discovery of new endpoints added to the DMVPN network and thus frees the administrator from the necessity of extra provisioning. The concept of NHS naturally induces hub-and-spoke topology on a typical DMVPN network. All spokes register their mappings with the hub (the NHS) and use it as the central information repository, querying for the mappings when needed.

Dynamic Routing

The use of dynamic routing inside DMVPN allows new nodes to dynamically learn routing information from the other nodes and advertise their own prefixes. Because of the hub-and-spoke network structure, the hub router is often the one that reflects routing updates between the spokes (if a distance-vector protocol is used) or acts as a designated router for link-state protocols. An important consequence of using NHRP is that every spoke may learn the NBMA address of another spoke and send traffic directly to this node, bypassing the hub. To use this spoke-to-spoke tunnel feature, routing protocols should be configured properly, because there are some limitations imposed by CEF switching. This is discussed in the following scenarios.

Dynamic IPsec Encryption

Because DMVPN has dynamic structure, static crypto maps are not suitable, so we must use IPsec profiles. A multipoint GRE tunnel equipped with an IPsec profile will track any new node registering with the hub and create a new IPsec SA to protect the traffic between the hub and the new spoke. Thus, a single mGRE tunnel may result in multiple IPSec SAs created dynamically upon the request from NHRP. Similarly, every spoke node that has learned another spoke's NBMA address may

https://t.me/learningnets

directly initiate an ISAKMP session with the other spoke and create a new IPsec SA to encrypt the spoke-to-spoke traffic. This promiscuous type of communication requires all nodes to be configured for scalable ISAKMP configuration. In the simplest case, you can use group (or wildcard) pre-shared-keys. In a real-world scenario, however, you may prefer to use a PKI and digital signatures authentication. Note that it is common to use IPsec transport mode for DMVPN to reduce the encapsulation overhead, because packets are already wrapped in GRE. However, in some cases the use of IPsec transport mode is simply mandatory, such as in cases when the spokes are behind NAT; this is required for proper NAT traversal with NHRP and IPsec.

Cisco Express Forwarding

CEF is an important component of DMVPN. Initially, all IP addresses on the shared medium are assigned to a special glean adjacency, meaning that they need an additional NHRP resolution request to be programmed properly. When a packet is sent to a destination across the DMVPN cloud, the destination's next-hop is looked up and found to be on the DMVPN could. At this point, CEF triggers an NHRP request to the NHS and learns the NBMA IP address corresponding to the tunnel IP address. Thus, for a prefix dynamically learned across the DMVPN cloud, CEF will only attempt to resolve the next-hop IP address. No attempt is performed to resolve the prefix itself. To permit direct routing between spokes, all prefixes should be learned by spokes with the next-hop IP addresses pointing to the route origin, which is the other spoke.

The above requirement limits the dynamic IGPs used for DMVPN. For example, if you are using RIP as VPN routing protocol, the next-hop information will be changed to the hub's IP address when reflected between spokes. This will prevent CEF from properly resolving next-hop IP addresses and disallow spoke-to-spoke tunnels. Here are the other protocol limitations:

https://t.me/learningnets

- When using EIGRP as the DMVPN routing protocol, use the command `no ip next-hop-self eigrp <AS_NUMBER>` on the hub router. This will instruct the hub to change the next-hop TLV in relayed routing updates to point to the originating router. Notice that EIGRP also needs the command `no ip split-horizon eigrp <AS_NUMBER>` to be configured on the hub tunnel interface, to permit relaying of EIGRP updates on the same interface it was received on.
- When using OSPF as the DMVPN routing protocol, configure the network type `broadcast` by applying the command `ip ospf network-type broadcast` . This allows for dynamic peer discovery and permits all OSPF nodes to see the routes with unmodified next-hop, because of the shared segment topology. Other network types, such as point-to-multipoint, will prevent the formation of spoke-to-spoke tunnels.

DMVPN Hub Configuration

Except for the tunnel interface where DMVPN specific settings are configured, it is mostly a GRE over IPsec configuration with IPsec profiles. First, IKE Phase1 settings are configured. For example:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
```

Then IKE Phase2 settings are configured. For example:

```
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
```

The next step is to configure the DMVPN tunnel interface on the hub. This tunnel has no destination. We source it off the Loopback0 interface, because this is usually more convenient if there are redundant links from spokes toward the hub.

- Notice the use of the `tunnel key <NUMBER>` command. With older IOS images, the tunnel will not come up until you issue this command. However, this is no longer true. It allows multiplexing multiple tunnels sourced off the same interface, and it is

https://t.me/learningnets

required when multiple mGRE tunnels are sourced of the same interface.

- The next commands, in order of importance, are **`ip nhrp network-id`** and **`ip nhrp authentication .`** The first command is mandatory and identifies an NHRP segment or domain; it basically enables NHRP on the tunnel. It should match on all routers sharing the same NHRP domain. The second command is optional and implements a simple plaintext authentication of NHRP messages to prevent misconfigurations.
- The next important command is **`ip nhrp map multicast dynamic .`** When configured, it instructs the router to replicate multicast traffic to all remote nodes registered via NHRP. It allows for dynamic IGP adjacencies to come up when the IGP uses multicast.
- The last command, **`tunnel protection ipsec profile <NAME>`** , applies the IPsec profile to the tunnel, and we've seen it already used with VTIs.
- Finally, notice the EIGRP routing commands on the tunnel interface. We assume that EIGRP is the DMVPN routing protocol. The two commands **`no ip split-horizon <AS_NUMBER>`** and **`no ip next-hop-self <AS_NUMBER>`** allow EIGRP to exchange routing updates across the hub and preserve the next-hop IP address in the updates to point to the update originators, which are the spokes.

```
interface Tunnel0
 ip address 100.100.100.1 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 tunnel key 123
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map multicast dynamic
 tunnel protection ipsec profile DMVPN
 no ip split-horizon eigrp 123
 no ip next-hop-self eigrp 123
```

DMVPN Spoke Configuration

IKE Phase1 and Phase2 settings must match on the spoke side for IPsec tunnel to be established:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
```

https://t.me/learningnets

```
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
```

Configuring the DMVPN tunnel interface is a bit different. The first two interesting commands are static NHRP maps configured via **`ip nhrp map`**. The first one, **`ip nhrp map <HUB_TUNNEL> <HUB_NBMA>`**, maps the tunnel IP address of the NHS to the NBMA IP address. This tells the local router how to reach the NHS. The second map,** `ip nhrp map multicast <HUB_NBMA>` **, maps all local multicast traffic to the hub. This allows for establishing of IGP adjacencies. Finally, the NHS tunnel IP address is configured using the command **`ip nhrp nhs <HUB_TUNNEL>`**. The optional command **`ip nhrp hold-time <INTERVAL>`** is scoped to let the hub know how long to keep the mapping valid and is sent in NHRP Registration messages. By default, the spoke will attempt to re-register every one-third of the configured hold time, allowing for information refresh.

```
interface Tunnel0
 ip address 100.100.100.2 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 tunnel key 123
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp  nhs 100.100.100.1
 ip nhrp hold-time 60
 tunnel protection ipsec profile DMVPN
```

DMVPN Phase1

DMVPN Phase1 uses a regular Point-to-Point GRE tunnel on all spokes and mGRE tunnel only on the hub, so it has some limitations:

- No direct spoke-to-spoke traffic is possible, because of the P2P tunnel nature on spokes.

https://t.me/learningnets

- All spoke-to-spoke traffic flows through the hub.
- Because a spoke cannot directly communicate with other spokes, the routing protocol used must set the next-hop on all updates to match the hub tunnel IP address.
- Multicast packets can only be exchanged between the hub and spokes, not directly between spokes.

The only advantage that DMVPN Phase1 has over Phase2/Phase3 is a direct result of its design limitations: Spokes do not require specific route information for networks advertised by each remote spoke; a summary or default route received from the hub is good enough.

Configuration

R1:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
interface Tunnel0
 ip address 100.100.100.1 255.255.255.0
 no ip split-horizon eigrp 123
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map multicast dynamic
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.11.0 0.0.0.255
```

https://t.me/learningnets

R2:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
interface Tunnel0
 ip address 100.100.100.2 255.255.255.0
 tunnel source Loopback0
 tunnel destination 150.1.1.1
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.22.0 0.0.0.255
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
```

https://t.me/learningnets

```
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
interface Tunnel0
 ip address 100.100.100.3 255.255.255.0
 tunnel source Loopback0
 tunnel destination 150.1.1.1
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.33.0 0.0.0.255
```

Verification

Verify the dynamic NHRP entries on the hub:

```
R1#show ip nhrp brief
   Target             Via             NBMA            Mode   Intfc   Claimed
100.100.100.2/32     100.100.100.2   150.1.2.2        dynamic  Tu0     <   >
100.100.100.3/32     100.100.100.3   150.1.3.3        dynamic  Tu0     <   >
!
!R1#show ip nhrp detail
100.100.100.2/32 via 100.100.100.2
   Tunnel0 created 00:11:58, expire 01:48:01 Type: dynamic, Flags: unique registered
   NBMA address: 150.1.2.2
100.100.100.3/32 via 100.100.100.3
   Tunnel0 created 00:11:43, expire 01:48:16 Type: dynamic, Flags: unique registered


   NBMA address: 150.1.3.3
```

On spokes, only static NHRP entries are available, for the hub:

```
R2#show ip nhrp brief
```

https://t.me/learningnets

```
   Target             Via             NBMA            Mode   Intfc   Claimed
100.100.100.1/32     100.100.100.1   150.1.1.1 static
   Tu0     <   >
!
!R3#show ip nhrp brief
   Target             Via             NBMA            Mode   Intfc   Claimed
100.100.100.1/32     100.100.100.1   150.1.1.1 static
   Tu0     <   >
```

Verify that EIGRP adjacencies are established and routes are learned on spokes, because EIGRP split-horizon is disabled:

```
R1#show ip eigrp 123 neighbors
EIGRP-IPv4 Neighbors for AS(123)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   100.100.100.3           Tu0                      11 00:03:16    8  1470  0  5
0   100.100.100.2           Tu0                      14 00:03:31   26  1470  0  5
!
!R2#show ip route eigrp 123 | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D        136.1.11.0/24 [90/26882560] via 100.100.100.1, 00:01:57, Tunnel0
D        136.1.33.0/24 [90/28162560] via 100.100.100.1, 00:01:57, Tunnel0
!
!R3#show ip route eigrp 123 | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D        136.1.11.0/24 [90/26882560] via 100.100.100.1, 00:04:14, Tunnel0
D        136.1.22.0/24 [90/28162560] via 100.100.100.1, 00:02:18, Tunnel0
```

Test connectivity between DMVPN protected networks:

```
R2#ping 136.1.33.3 source gigabitEthernet0/0.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.33.3, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.1 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
!
!R1#ping 136.1.33.3 source gigabitEthernet0/0.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.33.3, timeout is 2 seconds:
```

https://t.me/learningnets

```
Packet sent with a source address of 136.1.11.1 !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
!
!R1#ping 136.1.22.2 source gigabitEthernet0/0.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.22.2, timeout is 2 seconds:
Packet sent with a source address of 136.1.11.1 !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Verify that the IPsec tunnel is established between the hub and spokes (spoke-to-spoke is not available in DMVPN Phase1):

```
R1#show crypto session brief
Status: A- Active, U - Up, D - Down, I - Idle, S - Standby, N - Negotiating
        K - No IKE
ivrf = (none)
           Peer     I/F        Username          Group/Phase1_id   Uptime Status
150.1.3.3      Tu0                                150.1.3.3 00:10:02    UA
 150.1.2.2      Tu0                                150.1.2.2 00:10:18    UA
!
!R2#show crypto session brief
Status: A- Active, U - Up, D - Down, I - Idle, S - Standby, N - Negotiating
        K - No IKE
ivrf = (none)
           Peer     I/F        Username          Group/Phase1_id   Uptime Status
150.1.1.1      Tu0                                150.1.1.1 00:10:36    UA
!
!R3#show crypto session brief
Status: A- Active, U - Up, D - Down, I - Idle, S - Standby, N - Negotiating
        K - No IKE
ivrf = (none)
           Peer     I/F        Username          Group/Phase1_id   Uptime Status
150.1.1.1      Tu0                                150.1.1.1 00:10:35    UA
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 DMVPN Phase2 with PSK

Task

- Configure DMVPN Phase2 between R1, R2, and R3 as follows:
  - Use R1 as the hub.
  - Protect traffic between VLANs 11, 22, and 33.
  - Use EIGRP 123 as routing protocol.
  - Use 3DES/MD5 as cipher/hash for both IPsec phases.
  - Use DH group2 and a pre-shared-key of **CISCO**.
  - Authenticate NHRP packets using a string of **CISCO**.
  - Use the minimum overhead possible.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

DMVPN Phase2 allows dynamic spoke-to-spoke tunnels through the use of mGRE tunnels on spokes. This adds the following requirements on the routing protocol:

- Spokes must receive full routing updates for all other spoke-advertised networks; summarization or default route from the hub is not an option.
- When the hub reflects the updates between spokes, it must preserve the original next-hop value so that each spoke has in its routing table entries for remote spoke

https://t.me/learningnets

subnets with a next-hop value equal to the tunnel IP address of remote spokes.

To better understand the process, here is a brief description of the NHRP packet types (NHC is Next Hop Client, and NHS is Next Hop Server):

- **NRHP Registration Request** is used by NHC (spoke) routers to register their NBMA address with the configured NHS (hub).
- **NHRP Registration Reply** is used by NHS (hub) routers to acknowledge NHC (spoke) registration.
- **NHRP Resolution Request** is used by NHC (spoke) routers to resolve the NBMA address of a remote NHC (spoke) router.
- **NHRP Resolution Reply** is used by NHC (spoke) routers to respond to Resolution Requests received from other NHC (spoke) routers.

For example, an NHRP Registration Request/Reply exchange looks like the following:

```
NHRP: Send Registration Request via Tunnel0 vrf 0, packet size: 105
 src: 100.100.100.2, dst: 100.100.100.1
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 105 extoff: 52
 (M) flags: "unique nat ", reqid: 65566 src NBMA: 150.1.2.2
src protocol: 100.100.100.2, dst protocol: 100.100.100.1
 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
!
!NHRP: Receive Registration Reply via Tunnel0 vrf 0, packet size: 125
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 125 extoff: 52
 (M) flags: "unique nat ", reqid: 65566 src NBMA: 150.1.2.2
  src protocol: 100.100.100.2, dst protocol: 100.100.100.1

 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
```

https://t.me/learningnets

Each NHRP mapping contains one or multiple NHRP flags that actually define the router's behavior for that entry; not all flags are visible through show outputs (on debug, all flags are visible), but those available can be seen with `show ip nhrp` . The most important flags are:

- **authoritative**: When an NHC registers with an NHS, if it sets the authoritative flag, it delegates the NHS to act as a proxy so that the NHS can directly respond to NHRP resolution requests destined to this NHC (so resolution request is not forwarded to the NHC to answer).
- **registered**: When the NHS receives a Registration Request from an NHC, it creates the entry and adds this flag to it.
- **unique**: This flag is marked by NHC in the registration process and scoped to prevent misconfigurations or DOS attacks; when an NHC is registered and the NHS creates an NHRP mapping for it, for the lifetime of the entry the NHS does not accept new registration requests with different NBMA addresses for this entry.
- **used**: This flag basically says that the NHRP entry is used/active by the router; initially it was scoped to define that for this entry packets will be processed switched; however, with current IOS versions it is present even though packets are CEF switched.
- **router**: This flag describes an NHRP entry for a remote NHC.
- **nat**: This flag is used by NHC to signal in NHRP packets that NAT-aware DMVPN/NHRP is supported by their running IOS code.

In older IOS versions, the magic behind Phase2 was identified by these steps, assuming that SpokeA wants to communicate directly with SpokeB through a dynamically created tunnel:

1. SpokeA initiates traffic toward SpokeB through the hub and at the same time initiates an NHRP Resolution Request to the hub to resolve the NBMA address of SpokeB.
2. The hub replies to SpokeA with an NHRP Resolution Reply message, specifying the NBMA address of SpokeB.
3. SpokeB receives traffic from SpokeA and replies through the hub as well, initiating an NHRP Resolution Request to the hub to resolve the NBMA address of SpokeA.
4. The hub replies to SpokeB with an NHRP Resolution Reply message, specifying the NBMA address of SpokeA.
5. Both spokes have resolved their NBMA addresses (not exactly at the same time) and data-plane traffic is switched over via a direct spoke-to-spoke IPsec tunnel.

Note that spoke-to-spoke NBMA resolution via NHRP was always done by the hub because of the **authoritative** flag that entitled the hub to respond in the name of the

https://t.me/learningnets

clients. With current IOS versions, the behavior has changed and the **authoritative** flag is no longer present (the process will be debugged in the Verification section). Altghough by design it was supposed to work differently, the process was coded as follows:

1. SpokeA initiates traffic toward SpokeB through the hub and at the same time initiates an NHRP Resolution Request to the hub to resolve the NBMA address of SpokeB.
2. The hub forwards the received Resolution Request from SpokeA to SpokeB.
3. SpokeB receives traffic from SpokeA and replies through the hub as well, at the same time initiating an NHRP Resolution Request to the hub to resolve the NBMA address of SpokeA.
4. SpokeB receives the NHRP Resolution Request from SpokeA but relayed through the hub and responds also through the hub.
5. The hub forwards the received Resolution Request from SpokeB to SpokeA.
6. The hub forwards the received Resolution Reply from SpokeB to SpokeA.
7. SpokeA receives the Resolution Request from SpokeB and replies through the hub as well.
8. The hub forwards the received Resolution Reply from SpokeA to SpokeB.
9. Both spokes have resolved their NBMA addresses (not exactly at the same time) and data-plane traffic is switched over via a direct spoke-to-spoke IPsec tunnel.

Until the spoke-to-spoke NHRP mappings are created and the spoke-to-spoke IPsec tunnel is established, all spoke-to-spoke traffic passes through the hub. The spoke-to-spoke NHRP mapping traffic always passes through the hub; there is no direct spoke-to-spoke NHRP traffic. The spoke-to-spoke data traffic flows initially through the hub; after the control-plane converges, data-plane traffic flows directly between spokes, bypassing the hub.

As we will see in the debug outputs from the Verification section, when spoke-to-spoke traffic is detected, both spokes send NHRP Resolution Request messages for the remote spoke tunnel IP address, not for the remote spoke-advertised subnet. This is important, because in DMVPN Phase3, behavior is different. From the configuration perspective, as compared with DMVPN Phase1, we must convert spoke tunnels to mGRE, and if EIGRP is used, configure it on the hub to preserve the next-hop address when relaying updates from spoke to spoke.

Configuration

```
R1:
crypto isakmp policy 100
 authentication pre-share
```

https://t.me/learningnets

```
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
interface Tunnel0
 ip address 100.100.100.1 255.255.255.0
 no ip split-horizon eigrp 123
 no ip next-hop-self eigrp 123
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map multicast dynamic
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.11.0 0.0.0.255
```

R2:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
interface Tunnel0
```

https://t.me/learningnets

```
 ip address 100.100.100.2 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.22.0 0.0.0.255
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
interface Tunnel0
 ip address 100.100.100.3 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
```

https://t.me/learningnets

```
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.33.0 0.0.0.255
```

Verification

At this point, there should be no spoke-to-spoke tunnel created, because no spoke-to-spoke traffic was initiated. Review the NHRP table on both the hub and spokes, and identify the previously discussed NHRP flags:

```
R1#show ip nhrp
100.100.100.2/32 via 100.100.100.2
   Tunnel0 created 00:03:26, expire 01:56:33 Type: dynamic, Flags: unique registered
   NBMA address: 150.1.2.2
100.100.100.3/32 via 100.100.100.3
   Tunnel0 created 00:03:31, expire 01:56:29 Type: dynamic, Flags: unique registered
   NBMA address: 150.1.3.3
!
!R2#show ip nhrp
100.100.100.1/32 via 100.100.100.1
   Tunnel0 created 00:04:35, never expire  Type: static, Flags: used

   NBMA address: 150.1.1.1
```

Less detailed but useful information can be verified using the **`show dmvpn`** command; notice the dynamic or static type of mapping:

```
R1#show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel0, IPv4 NHRP Details
Type:Hub, NHRP Peers:2,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
1        150.1.2.2    100.100.100.2    UP 00:04:43     D
1        150.1.3.3    100.100.100.3    UP 00:04:48     D
!
!R2#show dmvpn
```

https://t.me/learningnets

```
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel0, IPv4 NHRP Details
Type:Spoke, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
1       150.1.1.1   100.100.100.1    UP 00:04:58     S
```

Although classic **`show crypto isakmp sa`** and **`show crypto ispec sa`** commands can be used, DMVPN IPsec sessions can be verified using the **`show dmvpn detail`** command; note that there are no IPsec sessions directly between spokes:

```
R1#show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface Tunnel0 is up/up, Addr. is 100.100.100.1, VRF ""
   Tunnel Src./Dest. addr: 150.1.1.1/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled
Type:Hub, Total NBMA Peers (v4/v6): 2

# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1      150.1.2.2   100.100.100.2    UP 00:05:43    D   100.100.100.2/32

    1      150.1.3.3   100.100.100.3    UP 00:05:48    D   100.100.100.3/32



Crypto Session Details:
--------------------------------------------------------------------------------
Interface: Tunnel0
```

https://t.me/learningnets

```
Session: [0x2CC8A138]
  IKEv1 SA: local 150.1.1.1/500 remote 150.1.2.2/500 Active
          Capabilities:(none) connid:1042 lifetime:23:54:16 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.2.2 IPSEC FLOW: permit 47 host 150.1.1.1 host 150.1.2.2
        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 81 drop 0 life (KB/Sec) 4273745/3256
        Outbound: #pkts enc'ed 80 drop 0 life (KB/Sec) 4273744/3256
   Outbound SPI : 0xE16659C4, transform : esp-3des esp-md5-hmac
    Socket State: Open
 Interface: Tunnel0
Session: [0x2CC8A040]
  IKEv1 SA: local 150.1.1.1/500 remote 150.1.3.3/500 Active
          Capabilities:(none) connid:1041 lifetime:23:54:11 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.3.3 IPSEC FLOW: permit 47 host 150.1.1.1 host 150.1.3.3
        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 85 drop 0 life (KB/Sec) 4359763/3251
        Outbound: #pkts enc'ed 84 drop 0 life (KB/Sec) 4359763/3251
   Outbound SPI : 0xE6DC5505, transform : esp-3des esp-md5-hmac
    Socket State: Open
!
!R2#show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==============================================================================

Interface Tunnel0 is up/up, Addr. is 100.100.100.2, VRF ""
   Tunnel Src./Dest. addr: 150.1.2.2/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled

IPv4 NHS:
100.100.100.1  RE priority = 0 cluster = 0
Type:Spoke, Total NBMA Peers (v4/v6): 1

# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1      150.1.1.1   100.100.100.1    UP 00:07:29    S   100.100.100.1/32



Crypto Session Details:
--------------------------------------------------------------------------------
```

https://t.me/learningnets

```
 Interface: Tunnel0

Session: [0x2CBDA83C]
  IKEv1 SA: local 150.1.2.2/500 remote 150.1.1.1/500 Active
          Capabilities:(none) connid:1042 lifetime:23:52:30 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.1.1 IPSEC FLOW: permit 47 host 150.1.2.2 host 150.1.1.1

        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 103 drop 0 life (KB/Sec) 4283981/3150
        Outbound: #pkts enc'ed 104 drop 0 life (KB/Sec) 4283981/3150
   Outbound SPI : 0xFA2FF074, transform : esp-3des esp-md5-hmac
    Socket State: Open
```

Notice that on each spoke, remote spoke subnets are learned with a next-hop of remote spoke tunnel IP address:

```
R2#show ip route eigrp 123 | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D        136.1.11.0/24 [90/26882560] via 100.100.100.1, 00:10:17, Tunnel0D 136.1.33.0/24
 [90/28162560] via 100.100.100.3
, 00:10:17, Tunnel0
```

If we look on R2 at the CEF entry for R3's remote subnet of 136.1.33.0/24, it shows an **incomplete** adjacency for the next-hop tunnel IP address of 100.100.100.3. If we look at the CEF entry for R3's tunnel IP address of 100.100.100.3, it shows as **adjacency punt**, which means that packets will be processed-switched to trigger the NHRP Resolution Request Process:

```
R2#show ip cef 136.1.33.0 internal
136.1.33.0/24, epoch 0, RIB[I], refcount 5, per-destination sharing
  sources: RIB
  ifnums:
   Tunnel0(13): 100.100.100.3
  path 2CBD929C, path list 2CF73068, share 1/1, type attached nexthop, for IPv4
  nexthop 100.100.100.3 Tunnel0, adjacency IP adj out of Tunnel0, addr 100.100.100.3 (incomplete)
 output chain: IP adj out of Tunnel0, addr 100.100.100.3 (incomplete)
!
!R2#show ip cef 100.100.100.3 internal
100.100.100.0/24, epoch 0, flags attached, connected, cover dependents, need deagg, RIB[C], refcount 5, per-destinat
  sources: RIB
  subblocks:
   gsb Connected chain head(1): 0x2D2315C8
```

https://t.me/learningnets

```
   Covered dependent prefixes: 3
     need deagg: 2
     notify cover updated: 1
  ifnums:
   Tunnel0(13)
  path 2CBD945C, path list 2CF731F8, share 1/1, type connected prefix, for IPv4
connected to Tunnel0, adjacency punt

  output chain: punt
```

Perform the following debugs on one spoke and the hub; generate spoke-to-poke traffic to identify the NHRP steps involved:

```
R1#debug nhrp
NHRP protocol debugging is onR1#debug nhrp packet
NHRP activity debugging is on
!R2#debug nhrp
NHRP protocol debugging is onR2#debug nhrp packet
NHRP activity debugging is on
!R2#ping 136.1.33.3 repeat 2
Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 136.1.33.3, timeout is 2 seconds:!!

Success rate is 100 percent (2/2), round-trip min/avg/max = 4/4/4 ms
```

The spoke identifies that 136.1.33.0/24 is available with a next-hop of 100.100.100.3, and it detects that it has no NHRP mapping and queues the NHRP Resolution Request for it:

```
R2:
NHRP: NHRP could not map 100.100.100.3 to NBMA, cache entry not found
NHRP: MACADDR: if_in null netid-in 0 if_out Tunnel0 netid-out 123
NHRP: Sending packet to NHS 100.100.100.1 on Tunnel0
NHRP: Checking for delayed event NULL/100.100.100.3 on list (Tunnel0).
NHRP: No node found.
NHRP: Adding Tunnel Endpoints (VPN: 100.100.100.3, NBMA: 150.1.1.1)
NHRP: Successfully attached NHRP subblock for Tunnel Endpoints (VPN: 100.100.100.3, NBMA: 150.1.1.1)
NHRP: Enqueued NHRP Resolution Request for destination: 100.100.100.3

NHRP: Checking for delayed event NULL/100.100.100.3 on list (Tunnel0).
NHRP: No node found.
NHRP-ATTR:  Requester Ext Len: Total ext_len  with NHRP attribute VPE 33
```

https://t.me/learningnets

At this point, it sends the NHRP Resolution Request for 100.100.100.3, which is R3's tunnel IP address, and it sends the request to the hub (source/destination of the packet is R2/R1):

```
R2:
NHRP: Sending NHRP Resolution Request for dest: 100.100.100.3 to nexthop: 100.100.100.1 using our src: 100.100.100.2
NHRP: Attempting to send packet via DEST 100.100.100.1
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.1.1

NHRP: Send Resolution Request via Tunnel0 vrf 0, packet size: 85
 src: 100.100.100.2, dst: 100.100.100.1
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 85 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 6
     src NBMA: 150.1.2.2
     src protocol: 100.100.100.2, dst protocol: 100.100.100.3
 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: 109 bytes out Tunnel0
```

At the same time, R2 receives an NHRP Resolution Request from R3, but not directly; it was sent through the hub, and the hub forwarded it to R2. It is clear that the packet was relayed through a node because it has a TTL value of 254, and it says **This is a forwarded packet**. We'll also see in the hub debugs that it actually forwarded it:

```
R2:
NHRP: Receive Resolution Request via Tunnel0 vrf 0, packet size: 105
 (F) afn: AF_IP(1), type: IP(800), hop: 254
, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 105 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 5
     src NBMA: 150.1.3.3
     src protocol: 100.100.100.3, dst protocol: 100.100.100.2
 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: netid_in = 123, to_us = 1
NHRP: nhrp_rtlookup for destination on 100.100.100.2 yielded interface Tunnel0, prefixlen 24
NHRP-ATTR: smart spoke feature and attributes are not configured,

NHRP-ATTR: In nhrp_process_recv_resolution_request eem_decision : TRUE, time : 0, LINE: 6955
```

https://t.me/learningnets

```
NHRP: This is a forwarded packet

NHRP: nhrp_rtlookup on 100.100.100.2 yielded interface Tunnel0, prefixlen 24
NHRP: Request was to us, responding with ouraddress
```

R2 received a Resolution Request from R3, and it still waits for a Resolution Response from R3. Until R2 receives a Resolution Reply or Resolution Request from R3, you will see several messages on R2 saying that it delays to respond to R3. This is because it waits to create the NHRP binding for R3 so that it can respond directly to it. Although the NHRP reply from R2 has source/destination IP addresses of R2/R3, it is still relayed through the hub. Delay messages on R2 are similar to: **`NHRP: Process delayed resolution request src:100.100.100.3 dst:100.100.100.2 .`**

```
R2:
NHRP: Send Resolution Reply via Tunnel0 vrf 0, packet size: 133

 src: 100.100.100.2, dst: 100.100.100.3
 (F) afn: AF_IP(1), type
Rack1R2#: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 133 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 5
     src NBMA: 150.1.3.3
     src protocol: 100.100.100.3, dst protocol: 100.100.100.2
 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.2.2
       client protocol: 100.100.100.2
NHRP: 157 bytes out Tunnel0
```

R2 receives the Resolution Reply from R3, which was also relayed through the hub:

```
R2:
NHRP: Receive Resolution Reply via Tunnel0 vrf 0, packet size: 153
 (F) afn: AF_IP(1), type: IP(800), hop: 254
, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 153 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 6
     src NBMA: 150.1.2.2
     src protocol: 100.100.100.2, dst protocol: 100.100.100.3
 (C-1) code: no error(0)
```

https://t.me/learningnets

```
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.3.3
       client protocol: 100.100.100.3
NHRP: netid_in = 0, to_us = 1NHRP: This is a forwarded packet
NHRP: Checking for delayed event NULL/100.100.100.3 on list (Tunnel0).
NHRP: No node found.
NHRP: No need to delay processing of resolution event nbma src:150.1.2.2 nbma dst:150.1.3.3
NHRP-ATTR: In nhrp_cache_pak LINE: 1391
 NHRP: Adding Tunnel Endpoints (VPN: 100.100.100.3, NBMA: 150.1.3.3)
NHRP: Cleanup up 0 stale cache entries
NHRP: Successfully attached NHRP subblock for Tunnel Endpoints (VPN: 100.100.100.3, NBMA: 150.1.3.3)
```

Review the result of the debug output on the hub site. First the hub receives the Resolution Request from R2, asking about R3's NBMA:

```
R1:
NHRP: Receive Resolution Request via Tunnel0 vrf 0, packet size: 85
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 85 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 6
     src NBMA: 150.1.2.2 src protocol: 100.100.100.2, dst protocol: 100.100.100.3

 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: netid_in = 123, to_us = 0
```

R1 forwards the Resolution Request to R3:

```
R1:
NHRP: nhrp_rtlookup for destination on 100.100.100.3 yielded interface Tunnel0, prefixlen 24
NHRP: netid_out 123, netid_in 123
NHRP-ATTR: In nhrp_recv_resolution_request NHRP Resolution Request packet is forwarded to 100.100.100.3
 NHRP: Attempting to forward to destination: 100.100.100.3
NHRP: Forwarding: NHRP SAS picked source: 100.100.100.1 for destination: 100.100.100.3
NHRP: Attempting to send packet via DEST 100.100.100.3
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.3.3
NHRP: Forwarding Resolution Request via Tunnel0 vrf 0, packet size: 105

 src: 100.100.100.1, dst: 100.100.100.3
 (F) afn: AF_IP(1), type: IP(800), hop: 254, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
```

https://t.me/learningnets

```
     pktsz: 105 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 6
     src NBMA: 150.1.2.2
     src protocol: 100.100.100.2, dst protocol: 100.100.100.3
 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: 129 bytes out Tunnel0
```

R1 receives a Resolution Request from R3, asking about R2's NBMA:

```
R1:
NHRP: Receive Resolution Request via Tunnel0 vrf 0, packet size: 85
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 85 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 5
     src NBMA: 150.1.3.3 src protocol: 100.100.100.3, dst protocol: 100.100.100.2

 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: netid_in = 123, to_us = 0
```

R1 forwards the Resolution Request to R2:

```
R1:
NHRP: nhrp_rtlookup for destination on 100.100.100.2 yielded interface Tunnel0, prefixlen 24
NHRP: netid_out 123, netid_in 123
NHRP-ATTR: In nhrp_recv_resolution_request NHRP Resolution Request packet is forwarded to 100.100.100.2
 NHRP: Attempting to forward to destination: 100.100.100.2
NHRP: Forwarding: NHRP SAS picked source: 100.100.100.1 for destination: 100.100.100.2
NHRP: Attempting to send packet via DEST 100.100.100.2
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.2.2
NHRP: Forwarding Resolution Request via Tunnel0 vrf 0, packet size: 105

 src: 100.100.100.1, dst: 100.100.100.2
 (F) afn: AF_IP(1), type: IP(800), hop: 254, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 105 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 5
     src NBMA: 150.1.3.3
     src protocol: 100.100.100.3, dst protocol: 100.100.100.2
 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
```

https://t.me/learningnets

```
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: 129 bytes out Tunnel0
```

R3 replies with a Resolution Reply for R2's Resolution Request, and this is received by R1, because both the request and reply are relayed through the hub:

```
R1:
NHRP: Receive Resolution Reply via Tunnel0 vrf 0, packet size: 133
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 133 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 5
     src NBMA: 150.1.3.3 src protocol: 100.100.100.3, dst protocol: 100.100.100.2

 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.2.2
       client protocol: 100.100.100.2
NHRP: netid_in = 0, to_us = 0
```

R1 forwards R3's Resolution Reply to R2:

https://t.me/learningnets

```
R1:
NHRP: Forwarding Resolution Reply to 100.100.100.3

NHRP: Attempting to forward to destination: 100.100.100.3
NHRP: Finding next idb with in_pak id: 0
NHRP: Forwarding: NHRP SAS picked source: 100.100.100.1 for destination: 100.100.100.3

NHRP: Attempting to send packet via DEST 100.100.100.3
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.3.3
NHRP: Forwarding Resolution Reply via Tunnel0 vrf 0, packet size: 153

 src: 100.100.100.1, dst: 100.100.100.3
 (F) afn: AF_IP(1), type: IP(800), hop: 254, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 153 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 5
     src NBMA: 150.1.3.3 src protocol: 100.100.100.3, dst protocol: 100.100.100.2

 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.2.2
       client protocol: 100.100.100.2
NHRP: 177 bytes out Tunnel0
```

R2 replies with a Resolution Reply for R3's Resolution Request, and this is received by R1, because both the request and reply are relayed through the hub:

https://t.me/learningnets

```
R1:
NHRP: Receive Resolution Reply via Tunnel0 vrf 0, packet size: 133

 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 133 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 6
     src NBMA: 150.1.2.2 src protocol: 100.100.100.2, dst protocol: 100.100.100.3

 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.3.3
       client protocol: 100.100.100.3
NHRP: netid_in = 0, to_us = 0
```

R1 forwards R2's Resolution Reply to R3:

```
R1:
NHRP: Forwarding Resolution Reply to 100.100.100.2
NHRP: Attempting to forward to destination: 100.100.100.2
NHRP: Finding next idb with in_pak id: 0
NHRP: Forwarding: NHRP SAS picked source: 100.100.100.1 for destination: 100.100.100.2
NHRP: Attempting to send packet via DEST 100.100.100.2
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.2.2
NHRP: Forwarding Resolution Reply via Tunnel0 vrf 0, packet size: 153
 src: 100.100.100.1, dst: 100.100.100.2
 (F) afn: AF_IP(1), type: IP(800), hop: 254, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 153 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 6
     src NBMA: 150.1.2.2 src protocol: 100.100.100.2, dst protocol: 100.100.100.3

 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.3.3
       client protocol: 100.100.100.3
NHRP: 177 bytes out Tunnel0
```

https://t.me/learningnets

The CEF adjacency of R2 for R3's remote subnet should no longer be **`incomplete`**, because the NHRP exchange between R2 and R3 was successfully completed:

```
R2#show ip cef 136.1.33.0 internal
136.1.33.0/24, epoch 0, RIB[I], refcount 5, per-destination sharing
  sources: RIB
  ifnums:
   Tunnel0(13): 100.100.100.3
  path 2CBD929C, path list 2CF73068, share 1/1, type attached nexthop, for IPv4
  nexthop 100.100.100.3 Tunnel0, adjacency IP midchain out of Tunnel0, addr 100.100.100.3 2C749380
output chain: IP midchain out of Tunnel0, addr 100.100.100.3 2C749380 IP adj out of GigabitEthernet0/0.28, addr 136.
!
!R2#show ip cef 100.100.100.3 internal

100.100.100.3/32, epoch 0, flags attached, refcount 5, per-destination sharing
  sources: Adj
  subblocks:
   Adj source: IP midchain out of Tunnel0, addr 100.100.100.3 2C749380
    Dependent covered prefix type adjfib cover 100.100.100.0/24
  ifnums:
   Tunnel0(13): 100.100.100.3
  path 2CBD953C, path list 2CF731A8, share 1/1, type adjacency prefix, for IPv4
  attached to Tunnel0, adjacency IP midchain out of Tunnel0, addr 100.100.100.3 2C749380
output chain: IP midchain out of Tunnel0, addr 100.100.100.3 2C749380 IP adj out of GigabitEthernet0/0.28, addr 136.
```

Verify the spoke-to-spoke dynamic NHRP mappings:

```
R2#show ip nhrp 100.100.100.3
100.100.100.3/32 via 100.100.100.3
   Tunnel0 created 00:03:26, expire 01:56:35
   Type: dynamic, Flags: router
   NBMA address: 150.1.3.3
!
!R2#show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel0, IPv4 NHRP Details
Type:Spoke, NHRP Peers:2,
```

https://t.me/learningnets

```
 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
     1       150.1.1.1   100.100.100.1    UP 00:23:57     S
1        150.1.3.3   100.100.100.3    UP 00:03:35     D
```

Verify that spoke-to-spoke IPsec tunnel is established:

```
R2#show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==============================================================================

Interface Tunnel0 is up/up, Addr. is 100.100.100.2, VRF ""
   Tunnel Src./Dest. addr: 150.1.2.2/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled

IPv4 NHS:
100.100.100.1  RE priority = 0 cluster = 0
Type:Spoke, Total NBMA Peers (v4/v6): 3

# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1      150.1.1.1   100.100.100.1    UP 00:26:06    S   100.100.100.1/32

    1      150.1.2.2   100.100.100.2    UP 00:05:46  DLX   100.100.100.2/32

    1      150.1.3.3   100.100.100.3    UP 00:05:44    D   100.100.100.3/32



Crypto Session Details:
--------------------------------------------------------------------------------
Interface: Tunnel0
Session: [0x2CBDA83C]
  IKEv1 SA: local 150.1.2.2/500 remote 150.1.1.1/500 Active
          Capabilities:(none) connid:1042 lifetime:23:33:53 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.1.1 IPSEC FLOW: permit 47 host 150.1.2.2 host 150.1.1.1
        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 352 drop 0 life (KB/Sec) 4283952/2033
```

https://t.me/learningnets

```
        Outbound: #pkts enc'ed 353 drop 0 life (KB/Sec) 4283952/2033
    Outbound SPI : 0xFA2FF074, transform : esp-3des esp-md5-hmac
     Socket State: Open
Interface: Tunnel0
Session: [0x2CBDA744]
  IKEv1 SA: local 150.1.2.2/500 remote 150.1.3.3/500 Active
          Capabilities:(none) connid:1043 lifetime:23:54:13
  IKEv1 SA: local 150.1.2.2/500 remote 150.1.3.3/500 Active
          Capabilities:(none) connid:1044 lifetime:23:54:13 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.3.3 IPSEC FLOW: permit 47 host 150.1.2.2 host 150.1.3.3

        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 1 drop 0 life (KB/Sec) 4244677/3253
        Outbound: #pkts enc'ed 0 drop 0 life (KB/Sec) 4244677/3253
    Outbound SPI : 0x8FBC7736, transform : esp-3des esp-md5-hmac
     Socket State: Open
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 DMVPN Phase3 with PSK

Task

- Configure DMVPN Phase3 between R1, R2, and R3 as follows:
  - Use R1 as the hub.
  - Protect traffic between VLANs 11, 22, and 33.
  - Use EIGRP 123 as routing protocol.
  - Use 3DES/MD5 as cipher/hash for both IPsec phases.
  - Use DH group2 and a pre-shared-key of **CISCO**.
  - Authenticate NHRP packets using a string of **CISCO**.
  - Use the minimum overhead possible.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

DMVPN Phase3 is scoped to remediate some design limitations of Phase2 and also enhance performance and scalability:

- The hub no longer modifies the next-hop value of routing updates when relayed between spokes, which gives you the Phase1 benefit of having the hub send only a default route or a summary to all spokes.
- It allows for hierarchical DMVPN deployments.

https://t.me/learningnets

- It resolves possible design issues with redundant hubs within the same DMVPN cloud.
- All spoke-to-spoke packets are now CEF switched, even before the spoke-to-spoke tunnel is established, because there are no more incomplete CEF adjacencies.

To achieve all the benefits above, Phase3 adds two new NHRP features:

- **NHRP Redirect** is configured on the hub and functions like IP redirect. If the hub receives traffic inbound on its mGRE interface from a spoke and forwards it outbound on the same mGRE interface to another spoke, it sends an NHRP Redirect message to the traffic initiator spoke, forcing it to find a better direct path for spoke-to-spoke traffic via NHRP Resolution Request.
- **NHRP Shortcut** is configured on the spoke side and functions as a CEF shortcut: When the NHRP Resolution Request/Reply process between spokes is finished, it forces an update to the CEF table. It basically overrides the next-hop value for a remote spoke network from the default initial hub tunnel IP address to the NHRP resolved remote spoke tunnel IP address, and this is visible both in CEF and RIB.

DMVPN Phase3 includes the following additional NHRP flags:

- **implicit** says that the NHRP mapping was learned from an NHRP packet exchange that was not scoped for this entry; on spokes, all NHRP entries for remote spoke tunnel IP addresses are implicit, because in Phase3 the NHRP Resolution Request is not issued for remote spoke tunnel IP addresses.
- **nho** means Next-Hop-Override and says that this entry modified the CEF table because shortcut switching is enabled.

When spoke-to-spoke traffic is generated, the following happens:

- SpokeA initiates traffic towards SpokeB LAN through the hub.
- The hub forwards the packet to SpokeB and at the same time sends an NHRP redirect message to SpokeA.
- SpokeA receives the redirect message and sends an NHRP Resolution Request through the hub for LAN network of SpokeB.
- The hub receives the NHRP Resolution Request message from SpokeA and relays it to SpokeB.
- SpokeB receives traffic from SpokeA and replies through the hub.
- The hub forwards the packet to SpokeA and at the same time sends an NHRP redirect message to SpokeB.
- SpokeB receives the redirect message and sends an NHRP Resolution Request

https://t.me/learningnets

through the hub for LAN network of SpokeA.
- The hub receives the NHRP Resolution Request message from SpokeB and forwards it to SpokeA.
- SpokeA and SpokeB negotiate a spoke-to-spoke IPsec tunnel.
- SpokeB receives the NHRP Resolution Request message from SpokeA and responds directly to SpokeA.
- SpokeA receives the NHRP Resolution Request message from SpokeB and responds directly to SpokeB.
- Traffic flows directly between spokes.

As shown in the Verification section, spoke-to-spoke data traffic flows initially through the hub; after control-plane (NHRP and IPsec) converges, data-plane traffic flows directly between spokes, bypassing the hub. From the control-plane perspective as related to NHRP, the behavior is different from DMVPN Phase2. Spoke-to-spoke NHRP Resolution Request flows through the hub, but spoke-to-spoke NHRP Resolution Reply flows directly between spokes, bypassing the hub.

As shown in the debug outputs in the Verification section, when spoke-to-spoke traffic is detected, both spokes send NHRP Resolution Request messages for the remote spoke advertised subnet, not for the remote spoke tunnel IP address. This is exactly the opposite if compared with DMVPN Phase2. This is because in DMVPN Phase3, the hub does not modify the next-hop value when relaying routing updates between spokes. So on spokes, all DMVPN learned subnets will have a next-hop equal to the hub IP address. For this reason, a spoke initiating traffic to a remote spoke advertised network is not aware of the remote spoke tunnel IP address, so it cannot initiate an NHRP Resolution Request for it; instead it generates an NHRP Resolution Request for the remote spoke network. To make a mapping of this difference between Phase2 and Phase3 with other technologies known to us, we can say that:

- NHRP Resolution Requests in DMVPN Phase2 is done for the next-hop IP address, which is identified from the routing table; this is similar to how Ethernet ARP works for a static route configured with a next-hop IP address.
- NHRP Resolution Requests in DMVPN Phase3 is done for the destination network being accessed; this is similar to how Ethernet ARP works for a static route configured with a multipoint exit interface.

https://t.me/learningnets

From the configuration perspective, compared with DMVPN Phase2, we must ensure that the hub changes the next-hop values in routing updates relayed from spoke to spoke. On the hub side we must enable NHRP redirect, and on the spokes side we must enable NHRP shortcut.

Configuration

R1:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
interface Tunnel0
 ip address 100.100.100.1 255.255.255.0
 no ip split-horizon eigrp 123
 ip next-hop-self eigrp 123
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map multicast dynamic
 ip nhrp redirect
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.11.0 0.0.0.255
```

R2:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
```

https://t.me/learningnets

```
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
interface Tunnel0
 ip address 100.100.100.2 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp shortcut
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.22.0 0.0.0.255
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
```

https://t.me/learningnets

```
!
interface Tunnel0
 ip address 100.100.100.3 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp shortcut
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.33.0 0.0.0.255
```

Verification

Initially, there is no spoke-to-spoke traffic. Verify the NHRP mappings and that the routing table changes if compared with Phase2:

```
R1#show ip nhrp
100.100.100.2/32 via 100.100.100.2
   Tunnel0 created 00:07:35, expire 01:52:24
   Type: dynamic, Flags: unique registered
   NBMA address: 150.1.2.2
100.100.100.3/32 via 100.100.100.3
   Tunnel0 created 00:07:32, expire 01:52:28
   Type: dynamic, Flags: unique registered
   NBMA address: 150.1.3.3
!
!R1#show ip eigrp 123 neighbors
EIGRP-IPv4 Neighbors for AS(123)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
1   100.100.100.3           Tu0                      13 00:07:46    2  1512  0  29
0   100.100.100.2           Tu0                      12 00:07:51   12  1512  0  46
!
!R2#show ip nhrp
100.100.100.1/32 via 100.100.100.1
   Tunnel0 created 00:05:51, never expire
   Type: static, Flags: used
```

https://t.me/learningnets

```
   NBMA address: 150.1.1.1
!
!R2#show dmvpn
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel0, IPv4 NHRP Details
Type:Spoke, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
1        150.1.1.1   100.100.100.1    UP 00:05:58     S
!
!R2#show ip route eigrp 123 | b Gateway
Gateway of last resort is not set

     136.1.0.0/16 is variably subnetted, 9 subnets, 2 masks
D       136.1.11.0/24 [90/26882560] via 100.100.100.1, 01:20:15, Tunnel0D 136.1.33.0/24
 [90/28162560] via 100.100.100.1
, 01:20:15, Tunnel0
```

Verify that R2 has an IPsec tunnel only with the hub:

```
R2#show dmvpn detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface Tunnel0 is up/up, Addr. is 100.100.100.2, VRF ""
   Tunnel Src./Dest. addr: 150.1.2.2/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled

IPv4 NHS:
100.100.100.1  RE priority = 0 cluster = 0
Type:Spoke, Total NBMA Peers (v4/v6): 1
```

https://t.me/learningnets

```
# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1      150.1.1.1   100.100.100.1    UP 00:06:10    S   100.100.100.1/32



Crypto Session Details:
--------------------------------------------------------------------------------
Interface: Tunnel0
Session: [0x2CBDA744]
  IKEv1 SA: local 150.1.2.2/500 remote 150.1.1.1/500 Active
          Capabilities:(none) connid:1049 lifetime:23:53:49 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.1.1 IPSEC FLOW: permit 47 host 150.1.2.2 host 150.1.1.1

        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 92 drop 0 life (KB/Sec) 4341709/3229
        Outbound: #pkts enc'ed 91 drop 0 life (KB/Sec) 4341709/3229
   Outbound SPI : 0xB30473C4, transform : esp-3des esp-md5-hmac
    Socket State: Open
```

Verify that R2 no longer has an incomplete CEF adjacency for R3's subnet, because the next-hop value from the routing table points to the hub now; it no longer points to the remote spoke. However, the CEF adjacency for R3's tunnel IP address is still **`adjacency punt`** because it is a directly connected host in the DMVPN cloud/subnet, which is not resolved through NHRP; therefore, traffic destined to it will be processed switched:

```
R2#show ip cef 136.1.33.0 internal
136.1.33.0/24, epoch 0, RIB[I], refcount 5, per-destination sharing
  sources: RIB
  ifnums:
   Tunnel0(13): 100.100.100.1
  path 2CBD95AC, path list 2CF73298, share 1/1, type attached nexthop, for IPv4
  nexthop 100.100.100.1 Tunnel0, adjacency IP midchain out of Tunnel0, addr 100.100.100.1 2C555EC0
 output chain: IP midchain out of Tunnel0, addr 100.100.100.1 2C555EC0 IP adj out of GigabitEthernet0/0.28, addr 136.
!
!R2#show ip cef 100.100.100.3 internal
100.100.100.0/24, epoch 0, flags attached, connected, cover dependents, need deagg, RIB[C], refcount 5, per-destinat
  sources: RIB
  subblocks:
   gsb Connected chain head(1): 0x2D2315C8
   Covered dependent prefixes: 3
     need deagg: 2
     notify cover updated: 1
  ifnums:
```

https://t.me/learningnets

```
   Tunnel0(13)
  path 2CBD94CC, path list 2CF73338, share 1/1, type connected prefix, for IPv4
connected to Tunnel0, adjacency punt

  output chain: punt
```

Generate spoke-to-spoke traffic and enable debugs on R2 and R1 to identify the NHRP message exchange process; spoke-to-spoke ICMP traffic will not work as we ping a host on the network, which does not exist:

```
R2#debug nhrp
NHRP protocol debugging is on
!R2#debug nhrp packet
NHRP activity debugging is on
!
!R1#debug nhrp
NHRP protocol debugging is on
!R1#debug nhrp packet
NHRP activity debugging is on
!R2#ping 136.1.33.3 source gigabitEthernet 0/0.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.33.3, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!R2#ping 136.1.33.30 source gigabitEthernet0/0.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.33.30, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 .....

Success rate is 0 percent (0/5)
```

Investigate NHRP packets on the hub side; only NHRP Resolution Request packets are relayed via the hub, and NHRP Resolution Reply packets are sent directly between spokes. The hub forwards packets out the same mGRE interface it was received on, so it generates an NHRP redirect message to R2:

```
R1:
NHRP: Tunnels gave us remote_nbma: 150.1.2.2 for Redirect
NHRP: Attempting to Redirect, remote_nbma: 150.1.2.2
NHRP: inserting (150.1.2.2/136.1.33.30) in redirect table
NHRP-ATTR:  Requester Ext Len: Total ext_len  with NHRP attribute VPE 29

NHRP: Attempting to send packet via DEST 136.1.22.2
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.2.2
```

https://t.me/learningnets

```
NHRP: Send Traffic Indication via Tunnel0 vrf 0, packet size: 97
 src: 100.100.100.1, dst: 136.1.22.2
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 97 extoff: 68 (M) traffic code: redirect(0)

     src NBMA: 150.1.1.1
     src protocol: 100.100.100.1, dst protocol: 136.1.22.2
     Contents of nhrp traffic indication packet:
        45 00 00 64 00 43 00 00 FE 01 75 33 88 01 16 02
        88 01 21 1E 08 00 5B 5D 00 11 00
NHRP: 121 bytes out Tunnel0
```

R3 responds to the traffic received from R2 through the hub also. The hub issues an NHRP redirect message for R3, because traffic was received and forwarded on the same interface, which is NHRP redirect enabled:

```
R1:
NHRP: Tunnels gave us remote_nbma: 150.1.3.3 for Redirect
NHRP: Attempting to Redirect, remote_nbma: 150.1.3.3
NHRP: inserting (150.1.3.3/136.1.22.2) in redirect table
NHRP-ATTR:  Requester Ext Len: Total ext_len  with NHRP attribute VPE 29

NHRP: Attempting to send packet via DEST 136.1.33.3
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.3.3
NHRP: Send Traffic Indication via Tunnel0 vrf 0, packet size: 97
 src: 100.100.100.1, dst: 136.1.33.3
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 97 extoff: 68 (M) traffic code: redirect(0)

     src NBMA: 150.1.1.1
     src protocol: 100.100.100.1, dst protocol: 136.1.33.3
     Contents of nhrp traffic indication packet:
        45 00 00 64 00 46 00 00 FE 01 75 4B 88 01 21 03
        88 01 16 02 00 00 EB 7F 00 12 00
NHRP: 121 bytes out Tunnel0
```

R2 receives the NHRP redirect message from the hub and sends an NHRP Resolution Request message for R3's remote subnet through the hub, which forwards it to R3:

```
R2:
NHRP: Receive Resolution Request via Tunnel0 vrf 0, packet size: 85
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
```

https://t.me/learningnets

```
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 85 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 14
     src NBMA: 150.1.2.2
     src protocol: 100.100.100.2, dst protocol: 136.1.33.3
 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: netid_in = 123, to_us = 0
NHRP: nhrp_rtlookup for destination on 136.1.33.3 yielded interface Tunnel0, prefixlen 24
NHRP: netid_out 123, netid_in 123
NHRP-ATTR: In nhrp_recv_resolution_request NHRP Resolution Request packet is forwarded to 136.1.33.3


NHRP: Attempting to forward to destination: 136.1.33.3
NHRP: Forwarding: NHRP SAS picked source: 100.100.100.1 for destination: 136.1.33.3
NHRP: Attempting to send packet via DEST 136.1.33.3
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.3.3
NHRP: Forwarding Resolution Request via Tunnel0 vrf 0, packet size: 105
 src: 100.100.100.1, dst: 136.1.33.3
 (F) afn: AF_IP(1), type: IP(800), hop: 254, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 105 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 14
     src NBMA: 150.1.2.2
     src protocol: 100.100.100.2, dst protocol: 136.1.33.3
 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
  addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: 129 bytes out Tunnel0
```

The NHRP Resolution Reply message from R3 to R2 is sent directly between spokes; it is not visible on the hub side. R3 also receives the NHRP redirect message from the hub and sends an NHRP Resolution Request message for R2's remote subnet through the hub, which forwards it to R2:

```
R3:
NHRP: Receive Resolution Request via Tunnel0 vrf 0, packet size: 85
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 85 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 12
     src NBMA: 150.1.3.3 protocol: 100.100.100.3, dst protocol: 136.1.22.2
 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
```

https://t.me/learningnets

```
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: netid_in = 123, to_us = 0
NHRP: nhrp_rtlookup for destination on 136.1.22.2 yielded interface Tunnel0, prefixlen 24
NHRP: netid_out 123, netid_in 123
NHRP-ATTR: In nhrp_recv_resolution_request NHRP Resolution Request packet is forwarded to 136.1.22.2


NHRP: Attempting to forward to destination: 136.1.22.2
NHRP: Forwarding: NHRP SAS picked source: 100.100.100.1 for destination: 136.1.22.2
NHRP: Attempting to send packet via DEST 136.1.22.2
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.2.2
NHRP: Forwarding Resolution Request via Tunnel0 vrf 0, packet size: 105
 src: 100.100.100.1, dst: 136.1.22.2
 (F) afn: AF_IP(1), type: IP(800), hop: 254, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 105 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 12
     src NBMA: 150.1.3.3
     src protocol: 100.100.100.3, dst protocol: 136.1.22.2
 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: 129 bytes out Tunnel0
```

On the spoke side, it first receives the NHRP Redirect message from the hub, so the spoke does not initially send an NHRP Resolution Request message as in Phase2:

```
R2:
NHRP: Receive Traffic Indication via Tunnel0 vrf 0, packet size: 97
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 97 extoff: 68 (M) traffic code: redirect(0)

     src NBMA: 150.1.1.1
     src protocol: 100.100.100.1, dst protocol: 136.1.22.2
     Contents of nhrp traffic indication packet:
        45 00 00 64 00 46 00 00 FE 01 75 4B 88 01 16 02
        88 01 21 03 08 00 E3 7F 00 12 00
NHRP: netid_in = 123, to_us = 0
NHRP: nhrp_rtlookup yielded GigabitEthernet0/0.22
NHRP: netid_out 0, netid_in 123
NHRP: Parsing NHRP Traffic Indication
```

At this point, R2 is instructed by the hub to send an NHRP Resolution Request for

https://t.me/learningnets

the remote spoke subnet, and it does so via the hub:

```
R2:
NHRP: Sending NHRP Resolution Request for dest: 136.1.33.3 to nexthop: 100.100.100.1 using our src: 100.100.100.2
NHRP: Attempting to send packet via DEST 100.100.100.1
NHRP: Encapsulation succeeded.  Sending NHRP Control Packet  NBMA Address: 150.1.1.1

NHRP: Send Resolution Request via Tunnel0 vrf 0, packet size: 85
 src: 100.100.100.2, dst: 100.100.100.1
 (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 85 extoff: 52
 (M) flags: "router auth src-stable nat ", reqid: 14
     src NBMA: 150.1.2.2
     src protocol: 100.100.100.2, dst protocol: 136.1.33.3
 (C-1) code: no error(0)
       prefix: 32, mtu: 17882, hd_time: 7200
       addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: 109 bytes out Tunnel0
```

At the same time, R2 receives an NHRP Resolution Request from R3, also sent via the hub. You can see this in the TTL, which has a value of 254 and a message saying **This is a forwarded packet**:

https://t.me/learningnets

R2:

```
NHRP: Receive Resolution Request via Tunnel0 vrf 0, packet size: 105
  (F) afn: AF_IP(1), type: IP(800), hop: 254
, ver: 1
      shtl: 4(NSAP), sstl: 0(NSAP)
      pktsz: 105 extoff: 52
  (M) flags: "router auth src-stable nat ", reqid: 12
      src NBMA: 150.1.3.3
      src protocol: 100.100.100.3, dst protocol: 136.1.22.2
  (C-1) code: no error(0)
        prefix: 32, mtu: 17916, hd_time: 7200
        addr_len: 0(NSAP), subaddr_len: 0(NSAP), proto_len: 0, pref: 0
NHRP: netid_in = 123, to_us = 0
NHRP: nhrp_rtlookup for destination on 136.1.22.2 yielded interface GigabitEthernet0/0.22, prefixlen 24

NHRP: netid_out 0, netid_in 123
NHRP-ATTR: smart spoke and attributes are not configured
NHRP-ATTR: In nhrp_process_recv_resolution_request eem_decision : TRUE, time : 0, LINE: 6955
NHRP: This is a forwarded packet
```

R3 directly responds to R2 with an NHRP Resolution Reply message, which can be seen in the source/destination IP of the packet:

R3:

```
NHRP: Send Resolution Reply via Tunnel0 vrf 0, packet size: 133
  src: 100.100.100.2, dst: 100.100.100.3
  (F) afn: AF_IP(1), type: IP(800), hop: 255, ver: 1
      shtl: 4(NSAP), sstl: 0(NSAP)
      pktsz: 133 extoff: 60
  (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 12 src NBMA: 150.1.3.3
 src protocol: 100.100.100.3, dst protocol: 136.1.22.2

  (C-1) code: no error(0)
        prefix: 24, mtu: 17882, hd_time: 7200
        addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
        client NBMA: 150.1.2.2
        client protocol: 100.100.100.2
NHRP: 157 bytes out Tunnel0
```

https://t.me/learningnets

Finally, R2 receives the NHRP Resolution Reply directly from R3. This can be seen in the TTL, which has a value of 255 and therefore has not been relayed:

```
NHRP: Receive Resolution Reply via Tunnel0 vrf 0, packet size: 133
 (F) afn: AF_IP(1), type: IP(800), hop: 255
, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 133 extoff: 60
 (M) flags: "router auth dst-stable unique src-stable nat ", reqid: 14 src NBMA: 150.1.2.2
src protocol: 100.100.100.2, dst protocol: 136.1.33.3

 (C-1) code: no error(0)
       prefix: 32, mtu: 17916, hd_time: 7200
       addr_len: 4(NSAP), subaddr_len: 0(NSAP), proto_len: 4, pref: 0
       client NBMA: 150.1.3.3
       client protocol: 100.100.100.3
NHRP: netid_in = 0, to_us = 1
```

Notice the flags present in the NHRP mapping of the remote spoke:

```
R2#show ip nhrp 136.1.33.0
136.1.33.0/24 via 100.100.100.3
   Tunnel0 created 00:00:21, expire 01:59:38 Type: dynamic, Flags: router used rib nho

   NBMA address: 150.1.3.3
```

Verify that the NHO bit from the NHRP mapping resulted in an NHO in the RIB/CEF table and traffic between spokes is direct, bypassing the hub:

```
R2#show ip route next-hop-override | s 136.1.33.0
D   %    136.1.33.0/24 [90/28162560] via 100.100.100.1, 01:09:31, Tunnel0
[NHO][90/1] via 100.100.100.3, 00:01:59, Tunnel0
!
!R2#show ip cef 136.1.33.0
136.1.33.0/24 nexthop 100.100.100.3 Tunnel0
```

Verify that a spoke-to-spoke direct tunnel has been established:

```
R2#show dmvpn peer network 136.1.33.0/24 detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
```

https://t.me/learningnets

```
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface Tunnel0 is up/up, Addr. is 100.100.100.2, VRF ""
   Tunnel Src./Dest. addr: 150.1.2.2/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled

IPv4 NHS:
100.100.100.1  RE priority = 0 cluster = 0
Type:Unknown, Total NBMA Peers (v4/v6): 1

# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1      150.1.3.3   100.100.100.3    UP 00:04:16  DT2      136.1.33.0/24



Crypto Session Details:
--------------------------------------------------------------------------------
Interface: Tunnel0
Session: [0x2CBDA744]
  IKEv1 SA: local 150.1.2.2/500 remote 150.1.3.3/500 Active
          Capabilities:(none) connid:1074 lifetime:23:55:43 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.3.3 IPSEC FLOW: permit 47 host 150.1.2.2 host 150.1.3.3

        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 1 drop 0 life (KB/Sec) 4179668/3343
        Outbound: #pkts enc'ed 10 drop 0 life (KB/Sec) 4179666/3343
   Outbound SPI : 0x904017FB, transform : esp-3des esp-md5-hmac
    Socket State: Open
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 DMVPN NAT Aware with PSK

Task

- Configure DMVPN Phase3 between R1, R2, and R3 as follows:
  - Use R1 as the hub.
  - Protect traffic between VLANs 11, 22, and 33.
  - Use EIGRP 123 as routing protocol.
  - Use 3DES/MD5 as cipher/hash for both IPsec phases.
  - Use DH group2 and a pre-shared-key of **CISCO**.
  - Authenticate NHRP packets using a string of **CISCO**.
  - Use the minimum overhead possible.
- Configure ASA4 so that R3's loopback address is statically NAT'ed to 136.1.18.3.
- Ensure that dynamic spoke-to-spoke tunnels are still established.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Starting with IOS 12.4(15)T, new NHRP protocol enhancements allow hubs and spokes in DMVPN cloud behind NAT devices. A new bit has been added for NHRP to signal that it supports routers behind NAT devices. The NAT identification is as follows:

https://t.me/learningnets

- The hub determines whether a registered spoke is behind a NAT device by comparing the source IP address in the GRE/IP header with the source NBMA address in the NHRP payload from the NHRP Registration Request message; if these two don't match, the spoke is behind a NAT device.
- The hub determines whether it is behind a NAT device by comparing its configured NBMA address with the NBMA address used by spokes in the NHRP payload of an NHRP Registration Request message.

When deploying NAT with DMVPN, the following restrictions should be considered:

- IPsec transport mode must be used.
- Any spoke/hub should use the same NAT IP address when communicating with any other DMVPN peer.
- Any type of NAT deployment is supported for spokes, whereas hubs must be statically NAT'ed:
  - When any two spokes are behind PAT devices, PAT functionality prevents the spoke-to-spoke tunnel from being established.
  - If a spoke is behind a dynamic NAT, making its NAT NBMA address subject to changes, you should disable the `unique` flag on the spoke with the command `ip nhrp registration no-unique` so that its registered NBMA on the hub can be overridden at any time.

Configuration

R3's NBMA address (Loopback0) must be NAT'ed for both spoke-to-hub and spoke-to-spoke traffic using the same IP address; otherwise, the difference in NBMA address values prevents the spoke-to-spoke tunnel from being established. Note that as with any DMVPN implementation, if the spoke-to-spoke tunnel fails to establish, traffic will flow through the hub.

ASA4:

```
object network R3_R2_STATIC
 host 150.1.3.3
 nat (VLAN38,VLAN28) static 136.1.18.3
object network R3_R1_STATIC
 host 150.1.3.3
 nat (VLAN38,VLAN18) static 136.1.18.3
```

R1:

```
crypto isakmp policy 100
 authentication pre-share
```

https://t.me/learningnets

```
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
interface Tunnel0
 ip address 100.100.100.1 255.255.255.0
 no ip split-horizon eigrp 123
 ip next-hop-self eigrp 123
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map multicast dynamic
 ip nhrp redirect
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.11.0 0.0.0.255
```

R2:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
```

https://t.me/learningnets

```
interface Tunnel0
 ip address 100.100.100.2 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp shortcut
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.22.0 0.0.0.255
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
 mode transport
!
crypto ipsec profile DMVPN
 set transform-set 3DES_MD5
!
!
interface Tunnel0
 ip address 100.100.100.3 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.1 150.1.1.1
 ip nhrp map multicast 150.1.1.1
 ip nhrp shortcut
 ip nhrp nhs 100.100.100.1
 tunnel protection ipsec profile DMVPN
!
```

https://t.me/learningnets

```
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.33.0 0.0.0.255
```

Verification

Verify that the hub detected R3 as behind a NAT device, the **NBMA address** is from the GRE/IP header, and the **Claimed NBMA address** is from the NHRP payload:

```
R1#show ip nhrp 100.100.100.3
100.100.100.3/32 via 100.100.100.3
   Tunnel0 created 00:40:08, expire 01:59:27
   Type: dynamic, Flags: unique registered used  NBMA address: 136.1.18.3
 (Claimed NBMA address: 150.1.3.3)
```

Verify that the IPsec tunnel between hub and spoke is established based on the NAT'ed NBMA address of R3:

```
R1#show dmvpn peer nbma 136.1.18.3
Legend: Attrb --> S - Static, D - Dynamic
, I - Incomplete N - NATed
, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel0, IPv4 NHRP Details
Type:Hub, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
1       136.1.18.3    100.100.100.3    UP 00:40:44     DN

!
!R1#show dmvpn peer nbma 136.1.18.3 detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================
```

https://t.me/learningnets

```
Interface Tunnel0 is up/up, Addr. is 100.100.100.1, VRF ""
   Tunnel Src./Dest. addr: 150.1.1.1/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled
Type:Hub, Total NBMA Peers (v4/v6): 1

# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
1     136.1.18.3   100.100.100.3    UP 00:40:48   DN   100.100.100.3/32
 Claimed Addr. 150.1.3.3



Crypto Session Details:
--------------------------------------------------------------------------------
 Interface: Tunnel0
Session: [0x2C8FB2A0]
  IKEv1 SA: local 150.1.1.1/4500 remote 136.1.18.3/4500 Active
          Capabilities:N connid:1002 lifetime:23:19:01 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.3.3 IPSEC FLOW: permit 47 host 150.1.1.1 host 136.1.18.3

        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 538 drop 0 life (KB/Sec) 4337920/1151
        Outbound: #pkts enc'ed 536 drop 0 life (KB/Sec) 4337988/1151
   Outbound SPI : 0x7F7368B8, transform : esp-3des esp-md5-hmac
    Socket State: Open
```

Generate spoke-to-spoke traffic to trigger the dynamic spoke-to-spoke tunnel:

https://t.me/learningnets

```
R2#ping 136.1.33.3 source gigabitEthernet0/0.22

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.33.3, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
!
!R2#ping 136.1.33.30 source gigabitEthernet0/0.22

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.33.30, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 .....

Success rate is 0 percent (0/5)
```

Verify that R2 identifies from the NHRP mapping that R3 is behind a NAT device:

```
R2#show ip nhrp 100.100.100.3
100.100.100.3/32 via 100.100.100.3
   Tunnel0 created 00:26:08, expire 01:59:07
   Type: dynamic, Flags: router implicit NBMA address: 136.1.18.3
 (Claimed NBMA address: 150.1.3.3)
```

Verify that the NHO CEF/RIB entry is created:

```
R2#show ip cef 136.1.33.0
136.1.33.0/24 nexthop 100.100.100.3 Tunnel0
!
!R2#show ip route next-hop-override | s 136.1.33.0
D   %    136.1.33.0/24 [90/28162560] via 100.100.100.1, 01:25:16, Tunnel0
 [NHO][90/1] via 100.100.100.3, 00:10:54, Tunnel0
!
!R2#show ip nhrp 136.1.33.0
136.1.33.0/24 via 100.100.100.3
   Tunnel0 created 00:11:43, expire 01:48:16
   Type: dynamic, Flags: router rib nho NBMA address: 136.1.18.3
 (Claimed NBMA address: 150.1.3.3)
```

Verify that the spoke-to-spoke IPsec tunnel is established based on R3's NAT'ed

https://t.me/learningnets

address:

```
R2#show dmvpn peer network 136.1.33.0/24
Legend: Attrb --> S - Static, D - Dynamic
, I - Incomplete N - NATed
, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface: Tunnel0, IPv4 NHRP Details
Type:Unknown, NHRP Peers:1,

 # Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb
 ----- --------------- --------------- ----- -------- -----
1       136.1.18.3    100.100.100.3    UP 00:01:49  DNT2
!
!R2#show dmvpn peer network 136.1.33.0/24 detail
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
        N - NATed, L - Local, X - No Socket
        # Ent --> Number of NHRP entries with same NBMA peer
        NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
        UpDn Time --> Up or Down Time for a Tunnel
==========================================================================

Interface Tunnel0 is up/up, Addr. is 100.100.100.2, VRF ""
   Tunnel Src./Dest. addr: 150.1.2.2/MGRE, Tunnel VRF ""
   Protocol/Transport: "multi-GRE/IP", Protect "DMVPN"
   Interface State Control: Disabled
   nhrp event-publisher : Disabled

IPv4 NHS:
100.100.100.1  RE priority = 0 cluster = 0
Type:Unknown, Total NBMA Peers (v4/v6): 1

# Ent  Peer NBMA Addr Peer Tunnel Add State  UpDn Tm Attrb    Target Network
----- --------------- --------------- ----- -------- ----- -----------------
    1     136.1.18.3   100.100.100.3    UP 00:01:52 DNT2      136.1.33.0/24
              Claimed Addr. 150.1.3.3



Crypto Session Details:
--------------------------------------------------------------------------------
Interface: Tunnel0
```

https://t.me/learningnets

```
Session: [0x300022C0]
  IKEv1 SA: local 150.1.2.2/4500 remote 136.1.18.3/4500 Active
          Capabilities:N connid:1003 lifetime:23:58:06 Crypto Session Status: UP-ACTIVE
  fvrf: (none), Phase1_id: 150.1.3.3 IPSEC FLOW: permit 47 host 150.1.2.2 host 136.1.18.3

        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 12 drop 0 life (KB/Sec) 4197089/3487
        Outbound: #pkts enc'ed 12 drop 0 life (KB/Sec) 4197091/3487
   Outbound SPI : 0x7C10BC27, transform : esp-3des esp-md5-hmac
    Socket State: Open
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 GETVPN Unicast Rekey with PSK

Task

- Configure GET VPN between R1, R2, and R3 as follows:
  - R1 should be the KS and R2/R3 the GMs.
  - Protect traffic between /16 subnet of Loopback0 interfaces.
  - Use 3DES/MD5 encryption for both IPSec phases.
  - Use a pre-shared-key of **CISCO** for authentication of GMs with the KS.
  - Use unicast for rekeying.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

GET VPN is a VPN technology that was originally designed to provide scalable encryption of multicast traffic. The primary purpose of GET VPN is to allow more than two devices to exchange traffic securely without requiring a high number of IPsec SAs. Traditional IPsec VPNs allow only for point-to-point security association (SA), so any communication involving more than two devices requires either a mesh of point-to-point tunnels or an overlay infrastructure with double encapsulation.

The key feature of GET VPN is preservation of the original IP packet's header, eliminating the need for overlay encapsulation and routing. Instead of adding a second layer of encapsulation built on the tunnel endpoints, GET VPN simply

https://t.me/learningnets

replicates the original packet's header using IPSec Tunnel mode. Preserving the original header makes it possible to use the core network's routing for optimal multicast packet replication (not just edge replication) and for better traffic engineering in the core network. This model perfectly fits private WANs built over MPLS cores (VPN in VPN), but it might not work over Internet if the customer is using private addressing (no overlay encapsulation). This is why GET VPN is usually positioned as an enterprise solution to be deployed over modern MPLS networks. For over-the-Internet VPNs where public IP addresses are required, DMVPN or GET VPN coupled with DMVPN is a viable solution.

To allow multiple devices to communicate securely using their original IP header and the same encryption policy, a centralized device is introduced in GET VPN topology. The central device is called Key Server (KS), and all other endpoints (routers) are called Group Members (GM). The protocol used between the KS and GMs is known as ISAKMP GDOI (Group Domain of Interpretation), which allows for establishing of shared group encryption keys.

GM and KS interaction

All GMs register with the KS, establishing Phase1 IKEv1 SA, and authenticate themselves using either pre-shared-keys or RSA signatures. After this, GDOI negotiations begin (protected by IKE SAs), and every member receives two group keys: KEK (Key Encryption Key) and TEK (Traffic Encryption Key). The KEK is used for protection of control plane messages (GDOI negotiations), whereas all nodes use TEK to encrypt data plane traffic (Phase2). Both keys are shared by all nodes and periodically refreshed by the KS. The original IKE SA established by every GM with the KS will expire in time, leaving only GDOI SA to negotiate group policies. The purpose of the initial SA is to provide a GM with the KEK and TEK. The KS may use either multicast or unicast messages to refresh KEK/TEK, with the multicast option allowing for higher scalability and bandwidth conservation. However, unicast re-keying allows for acknowledged key distribution, because the KS confirms that every GM receives the key update.

Central Policy

In addition to obtaining the encryption keys, GMs also download group policy from the KS. The policy is defined as an extended ACL, specifying the traffic to be encrypted during IPSec Phase2. This allows for concentrating all configurations in the KS and for configuring GMs with just the settings necessary to connect to the KS.

Replay Protection

https://t.me/learningnets

Replay attack is a type of injection attack in which the same IPsec packet is being inserted in the legitimate flow to affect the endpoint's behavior. IPsec prevents replay by running sequence numbers for each end of the tunnel connection and incrementing for every next packet sent. If the packet received does not have the expected sequence number, it is dropped. This technique, however, does not work for group encryption, because there could be many endpoints involved. Instead, the KS runs a special pseudo-time virtual clock and distributes pseudo-time to the group members. This is an independent protocol that does not require NTP to work properly. Every GM syncs the pseudo-time with the KS and timestamps all outgoing packets. A receiving side checks that the timestamp is within the acceptable window, such as plus/minus 10 seconds. If it is not, the packet is dropped.

Configuration: KS

Here is a short overview of configuring a KS for GET VPN. As usual, you start by defining an ISAKMP policy, which is used for GM's registration. Here you specify the authentication type and ISAKMP SA parameters:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
```

The sample above enables the use of pre-shared keys for authentication and configures a wildcard key. After the ISAKMP policy, you have to define an IPsec transform-set and IPsec profile:

```
crypto ipsec transform-set GETVPN esp-3des esp-md5-hmac
!
crypto ipsec profile GETVPN_PROFILE
  set transform-set GETVPN
```

The next step is to generate RSA keys to sign the re-keying messages and define an access-list for interesting traffic. This access-list will be downloaded by every member to know what traffic is protected by GET VPN:

https://t.me/learningnets

```
crypto key generate rsa general-keys label GETVPN_KEYS modulus 1024
!
access-list 100 permit ip 150.1.0.0 0.0.0.255 any
```

Finally, you may configure GDOI policy in the KS. This is what defines the GDOI SA settings, the re-keying process, and other settings. For every GDOI group, you must define an identity and the RSA key-pair to sign re-keying messages. In addition to that, at least one IPsec profile must be defined, to specify the data-plane protection options:

```
crypto gdoi group GETVPN_GROUP
  identity number 1234
  server local
   rekey authentication mypubkey rsa GETVPN_KEYS
   rekey transport unicast
   address ipv4 150.1.1.1
   sa ipsec 1
    profile GETVPN_PROFILE
    match address ipv4 100
    replay time window-size 5
```

Here you can see the IPsec profile associated with the GDOI group, the interesting traffic being defined, and the replay window size in seconds. Notice the address command that specifies the local address to be used for the server. This IP address is needed when KS communicates to another KS, such as for redundancy purposes.

Configuration: GM

Group member configuration is much simpler than KS configuration, because the policy is defined in the KS. You must specify the ISAKMP policy, the ISAKMP authentication information, and the GDOI server IP address. Here is a sample configuration that defines an ISAKMP policy and authenticates using pre-shared-keys:

https://t.me/learningnets

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp key CISCO address 150.1.1.1
```

The next part is defining a GDOI group and the KSIP address. After this, you may create a crypto-map that references this group and apply it to a physical interface. Notice the use of the local address with the crypto-map; this is the address that will be used for ISAKMP authentication. The data plane traffic will continue to use its original IP addressing:

```
crypto gdoi group GETVPN_GROUP_GM
  identity number 1234
  server address ipv4 150.1.1.1
!
crypto map GETVPN_MAP local-address Loopback0
crypto map GETVPN_MAP 100 gdoi
  set group GETVPN_GROUP_GM
!
interface FastEthernet0/0
  crypto map GETVPN_MAP
```

Configuration

```
R1:
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
crypto ipsec transform-set GETVPN esp-3des esp-md5-hmac
!
crypto ipsec profile GETVPN_PROFILE
 set transform-set GETVPN
!
```

https://t.me/learningnets

```
crypto key generate rsa general-keys label GETVPN_KEYS modulus 1024
access-list 100 permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
!
crypto gdoi group GETVPN_GROUP
  identity number 123
  server local
   rekey authentication mypubkey rsa GETVPN_KEYS
   rekey transport unicast
   address ipv4 136.1.18.1
   sa ipsec 1
    profile GETVPN_PROFILE
    match address ipv4 100
```

R2:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 136.1.18.1
crypto gdoi group GETVPN_GROUP_GM
 identity number 123
  server address ipv4 136.1.18.1
!
crypto map GETVPN_MAP local-address GigabitEthernet0/0.28
crypto map GETVPN_MAP 10 gdoi
  set group GETVPN_GROUP_GM
!
interface GigabitEthernet0/0.28
  crypto map GETVPN_MAP
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 136.1.18.1
crypto gdoi group GETVPN_GROUP_GM
 identity number 123
```

https://t.me/learningnets

```
  server address ipv4 136.1.18.1
!
crypto map GETVPN_MAP local-address FastEthernet0/0.38
crypto map GETVPN_MAP 10 gdoi
  set group GETVPN_GROUP_GM
!
interface FastEthernet0/0.38
  crypto map GETVPN_MAP
```

Verification

The following two messages on the GM console indicate the start and finish of the registration message with the KS:

```
%GDOI-5-GM_REKEY_TRANS_2_UNI: Group GETVPN_GROUP_GM transitioned to Unicast Rekey.
%CRYPTO-5-GM_REGSTER: Start registration to KS 136.1.18.1 for group GETVPN_GROUP_GM using address 136.1.38.3
%GDOI-5-GM_REGS_COMPL: Registration to KS 136.1.18.1 complete for group GETVPN_GROUP_GM using address 136.1.38.3
%GDOI-5-GM_INSTALL_POLICIES_SUCCESS: SUCCESS: Installation of Reg/Rekey policies from KS 136.1.18.1 for group GETVPN
```

Verify that both GMs have registered with KS:

```
R1#show crypto gdoi ks members

Group Member Information :

Number of rekeys sent for group GETVPN_GROUP : 0
Group Member ID     : 136.1.28.2  GM Version: 1.0.3
Group ID            : 123
Group Name          : GETVPN_GROUP
Key Server ID       : 136.1.18.1
 Rekeys sent        : 0
 Rekeys retries     : 0
 Rekey Acks Rcvd    : 0
 Rekey Acks missed  : 0

 Sent seq num : 0       0       0       0
Rcvd seq num :  0       0       0       0
Group Member ID     : 136.1.38.3  GM Version: 1.0.1
Group ID            : 123
Group Name          : GETVPN_GROUP
Key Server ID       : 136.1.18.1

 Rekeys sent        : 0
 Rekeys retries     : 0
```

https://t.me/learningnets

```
 Rekey Acks Rcvd   : 0
 Rekey Acks missed : 0

 Sent seq num : 0       0       0       0
Rcvd seq num :  0       0       0       0
```

Verify the KEK/TEK settings on the KS:

```
R1#show crypto gdoi ks policy
Key Server Policy:
For group GETVPN_GROUP (handle: 2147483650) server 136.1.18.1 (handle: 2147483650):

  # of teks : 1  Seq num : 0 KEK POLICY (transport type : Unicast)
    spi : 0xB82B8E09081C62CFD9B5FB8D434459B0
    management alg     : disabled    encrypt alg       : 3DES
    crypto iv length   : 8           key size          : 24
    orig life(sec): 86400       remaining life(sec): 85379
    sig hash algorithm : enabled     sig key length    : 162
    sig size           : 128       sig key name      : GETVPN_KEYS
 TEK POLICY (encaps : ENCAPS_TUNNEL)
    spi                : 0x7010693A
    access-list        : 100 transform          : esp-3des esp-md5-hmac

    alg key size       : 24            sig key size          : 16
    orig life(sec)     : 3600          remaining life(sec)   : 2580
    tek life(sec)      : 3600          elapsed time(sec)     : 1020
    override life (sec): 0             antireplay window size: 64
```

Verify that both GMs have successfully registered with the KS, downloaded the encryption access-list, and received the KEK/TEK:

```
R2#show crypto gdoi
GROUP INFORMATION

    Group Name               : GETVPN_GROUP_GM
    Group Identity           : 123
    Crypto Path              : ipv4
    Key Management Path      : ipv4
    Rekeys received          : 0
    IPSec SA Direction       : Both
 Group Server list       : 136.1.18.1
 Group member            : 136.1.28.2       vrf: None
       Version            : 1.0.3  Registration status   : Registered
       Registered with    : 136.1.18.1
       Re-registers in    : 2718 sec
```

https://t.me/learningnets

```
        Succeeded registration: 1
        Attempted registration: 1
        Last rekey from       : 0.0.0.0
        Last rekey seq num    : 0
        Unicast rekey received: 0
        Rekey ACKs sent       : 0
        Rekey Received        : never
        allowable rekey cipher: any
        allowable rekey hash  : any
        allowable transformtag: any ESP

    Rekeys cumulative
        Total received        : 0
        After latest register : 0
        Rekey Acks sents      : 0
 ACL Downloaded From KS 136.1.18.1:
 access-list  permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
 KEK POLICY:
    Rekey Transport Type     : Unicast
    Lifetime (secs)          : 85647 Encrypt Algorithm        : 3DES
    Key Size                 : 192      Sig Hash Algorithm       : HMAC_AUTH_SHA
    Sig Key Length (bits)    : 1024
 TEK POLICY for the current KS-Policy ACEs Downloaded:
  GigabitEthernet0/0.28:
    IPsec SA:
        spi: 0x7010693A(1880123706) transform: esp-3des esp-md5-hmac
        sa timing:remaining key lifetime (sec): (2848)
        Anti-Replay : Disabled
!
!R3#show crypto gdoi
GROUP INFORMATION

    Group Name               : GETVPN_GROUP_GM
    Group Identity           : 123
    Rekeys received          : 0
    IPSec SA Direction       : Both
 Group Server list        : 136.1.18.1
 Group member             : 136.1.38.3        vrf: None
 Registration status    : Registered
        Registered with       : 136.1.18.1
        Re-registers in       : 3046 sec
        Succeeded registration: 1
        Attempted registration: 3
        Last rekey from       : 0.0.0.0
        Last rekey seq num    : 0
        Unicast rekey received: 0
```

https://t.me/learningnets

```
        Rekey ACKs sent       : 0
        Rekey Received        : never
        allowable rekey cipher: any
        allowable rekey hash  : any
        allowable transformtag: any ESP

    Rekeys cumulative
        Total received        : 0
        After latest register : 0
        Rekey Acks sents      : 0
 ACL Downloaded From KS 136.1.18.1:
 access-list  permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
 KEK POLICY:
    Rekey Transport Type     : Unicast
    Lifetime (secs)          : 86034 Encrypt Algorithm          : 3DES
    Key Size                 : 192      Sig Hash Algorithm        : HMAC_AUTH_SHA
    Sig Key Length (bits)    : 1024
 TEK POLICY for the current KS-Policy ACEs Downloaded:
  FastEthernet0/0.38:
    IPsec SA:
        spi: 0x7010693A(1880123706) transform: esp-3des esp-md5-hmac

        sa timing:remaining key lifetime (sec): (3151)
        Anti-Replay :  Disabled
```

Verify that IPsec SAs are established on both GMs:

https://t.me/learningnets

```
R2#show crypto gdoi ipsec sa


SA created for group GETVPN_GROUP_GM:
  GigabitEthernet0/0.28:
    protocol = ip local ident  = 150.1.0.0/16, port = 0
remote ident = 150.1.0.0/16, port = 0

      direction: Both, replay: Disabled
!
!R3#show crypto gdoi ipsec sa


SA created for group GETVPN_GROUP_GM:
  FastEthernet0/0.38:
    protocol = ip local ident  = 150.1.0.0/16, port = 0
remote ident = 150.1.0.0/16, port = 0

      direction: Both, replay: Disabled
```

Verify that traffic between GM Loopback0 addresses is encrypted:

```
R2#ping 150.1.3.3 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.2.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
!
!R3#show crypto ipsec sa

interface: FastEthernet0/0.38
    Crypto map tag: GETVPN_MAP, local addr 136.1.38.3

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   remote ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   current_peer 0.0.0.0 port 848
     PERMIT, flags={origin_is_acl,} #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
#pkts decaps: 5, #pkts decrypt: 5, #pkts verify: 5
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0
```

https://t.me/learningnets

```
     local crypto endpt.: 136.1.38.3, remote crypto endpt.: 0.0.0.0
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.38
     current outbound spi: 0x7010693A(1880123706)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0x7010693A(1880123706) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2047, flow_id: NETGX:47, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (2288)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0x7010693A(1880123706) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2048, flow_id: NETGX:48, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (2288)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 DMVPN and GETVPN Integration

Task

- Configure a DMVPN Phase2 cloud between R2 and R3:
  - R2 is the hub and R3 is the spoke.
  - Use EIGRP AS 123 over the DMVPN cloud and advertise VLAN 22 and VLAN 33.
- Ensure that traffic over the DMVPN cloud is GET VPN protected.
  - R1 should be the KS and R2/R3 the GMs.
  - Use 3DES/MD5 encryption for both IPSec phases.
  - Use a pre-shared-key of **CISCO** for authentication of GMs with the KS.
  - Use unicast for rekeying.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

DMVPN has a scaling problem in terms of SAs: essentially, for each spoke-to-spoke dynamic tunnel, a pair of two IPsec SAs is created, whereas on the hub, a pair of two IPsec SAs is created for each spoke. GET VPN fixes this problem by using a pair of two IPsec SAs for each entry defined in the encryption access-list. For this reason, we can use DMVPN without encryption (no IPsec profile applied at tunnel level) and configure GET VPN to encrypt traffic between GRE tunnel headends. Another advantage to using GET VPN to encrypt DMVPN traffic is that there is no

https://t.me/learningnets

need to negotiate a spoke-to-spoke IPsec tunnel on demand, so traffic no longer passes through the hub. All traffic is routed over the DMVPN tunnel, and because it is GRE encapsulated, it will be automatically encrypted.

From the configuration perspective, you must simply ensure that there are no IPsec policies applied directly to DMVPN (GET VPN will be in charge with IPsec). Configure a non-IPsec aware DMVPN cloud. Note that the encryption ACL on the KS should allow GRE traffic between tunnel headends; the proxy ACL can use IP statetements instead of GRE statements to ensure all IP traffic between tunnel headends is encrypted.

Configuration

R1:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
crypto ipsec transform-set GETVPN esp-3des esp-md5-hmac
!
crypto ipsec profile GETVPN_PROFILE
 set transform-set GETVPN
!
crypto key generate rsa general-keys label GETVPN_KEYS modulus 1024 exportable
access-list 100 permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
!
crypto gdoi group GETVPN_GROUP
  identity number 123
  server local
   rekey authentication mypubkey rsa GETVPN_KEYS
   rekey transport unicast
   address ipv4 136.1.18.1
   sa ipsec 1
    profile GETVPN_PROFILE
    match address ipv4 100
```

R2:

```
interface Tunnel0
 ip address 100.100.100.2 255.255.255.0
```

https://t.me/learningnets

```
 no ip split-horizon eigrp 123
 no ip next-hop-self eigrp 123
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map multicast dynamic
!
router eigrp 123
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.22.0 0.0.0.255
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 136.1.18.1
crypto gdoi group GETVPN_GROUP_GM
 identity number 123
  server address ipv4 136.1.18.1
!
crypto map GETVPN_MAP local-address GigabitEthernet0/0.28
crypto map GETVPN_MAP 10 gdoi
  set group GETVPN_GROUP_GM
!
interface GigabitEthernet0/0.28
  crypto map GETVPN_MAP
```

R3:

```
interface Tunnel0
 ip address 100.100.100.3 255.255.255.0
 tunnel source Loopback0
 tunnel mode gre multipoint
 ip nhrp network-id 123
 ip nhrp authentication CISCO
 ip nhrp map 100.100.100.2 150.1.2.2
 ip nhrp map multicast 150.1.2.2
 ip nhrp nhs 100.100.100.2
!
router eigrp 123
```

https://t.me/learningnets

```
 no auto-summary
 network 100.100.100.0 0.0.0.255
 network 136.1.33.0 0.0.0.255
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 136.1.18.1
crypto gdoi group GETVPN_GROUP_GM
 identity number 123
  server address ipv4 136.1.18.1
!
crypto map GETVPN_MAP local-address FastEthernet0/0.38
crypto map GETVPN_MAP 10 gdoi
  set group GETVPN_GROUP_GM
!
interface FastEthernet0/0.38
  crypto map GETVPN_MAP
```

Verification

Verify the dynamic NHRP mapping on the hub and EIGRP adjacency:

https://t.me/learningnets

```
R2#show ip nhrp
100.100.100.3/32 via 100.100.100.3

   Tunnel0 created 00:23:48, expire 01:36:11
   Type: dynamic, Flags: unique registered NBMA address: 150.1.3.3

!
!R2#show ip eigrp 123 neighbors

EIGRP-IPv4 Neighbors for AS(123)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   100.100.100.3           Tu0                      14 00:24:33    7  1470  0  3

!
!R2#show ip route eigrp 123 | b Gateway

Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 7 subnets, 2 masks
D        136.1.33.0/24 [90/26882560] via 100.100.100.3, 00:24:54, Tunnel0
```

Verify that traffic over the DMVPN cloud is encrypted/decrypted through GET VPN:

```
R3#clear crypto sa counters
!
!R3#show crypto ipsec sa

interface: FastEthernet0/0.38
    Crypto map tag: GETVPN_MAP, local addr 136.1.38.3

   protected vrf: (none) local  ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
remote ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   current_peer 0.0.0.0 port 848
     PERMIT, flags={origin_is_acl,} #pkts encaps: 2, #pkts encrypt: 2, #pkts digest: 2
#pkts decaps: 2, #pkts decrypt: 2, #pkts verify: 2
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.38.3, remote crypto endpt.: 0.0.0.0
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.38
```

https://t.me/learningnets

```
     current outbound spi: 0xAFEDF469(2951607401)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0xC0973554(3231135060) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2031, flow_id: NETGX:31, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (3497)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE

     inbound ah sas:

     inbound pcp sas:
 outbound esp sas:
      spi: 0xC0973554(3231135060) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2032, flow_id: NETGX:32, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (3497)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE

     outbound ah sas:

     outbound pcp sas:
!
!R3#ping 150.1.2.2 source loopback0 repeat 100

Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 150.1.2.2, timeout is 2 seconds:
Packet sent with a source address of 150.1.3.3
!!!
```

```
    Crypto map tag: GETVPN_MAP, local addr 136.1.38.3

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   remote ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   current_peer 0.0.0.0 port 848
     PERMIT, flags={origin_is_acl,} #pkts encaps: 208, #pkts encrypt: 208, #pkts digest: 208
#pkts decaps: 208, #pkts decrypt: 208, #pkts verify: 208
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.38.3, remote crypto endpt.: 0.0.0.0
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.38
     current outbound spi: 0xC0973554(3231135060)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0xC0973554(3231135060) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2031, flow_id: NETGX:31, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (3258)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0xC0973554(3231135060) transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2032, flow_id: NETGX:32, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (3258)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv1 GETVPN Key Server Redundancy

Task

- Configure GET VPN between R1, R2, and R3 as follows:
  - R1 should be the primary KS and R2 the secondary KS.
  - R3 should be the GM.
  - Protect traffic between /16 subnet of Loopback0 interfaces.
  - Use 3DES/MD5 encryption for both IPSec phases.
  - Use a pre-shared-key of **CISCO** for authentication of GMs with the KS.
  - Use unicast for rekeying.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Because of the importance of KS, GET VPN supports multiple KS's serving the same group. This is called COOP KS, and the member KS's use a special protocol to negotiate and synchronize with each other. Among a group of KS's, one is elected as primary and others are designated as secondary. The primary server is responsible for re-keying, but GMs may register to any KS in the group and obtain current keys. This allows for load distribution in addition to the redundancy. The primary KS sends periodic GDOI messages to the secondaries; when a certain number of messages are missing, the re-election process is started and a new primary is elected. The GMs could be configured with multiple KS addresses; if the

https://t.me/learningnets

one configured as primary fails and the KS detects it, the next in sequence is tried for registration. However, in general, a KS failure is transparent to the GM if it already has registered and obtained KEK/TEK. The secondary KS taking the role of the primary will transparently continue to notify the GM of the re-keys. The initial ISAKMP SA established between the GM and the KS is no longer used after GDOI SA has been established, so the client will not detect the loss of the original server.

Here is a brief outline of the COOP KS configuration steps:

1. Generate RSA keys in the primary KS (as required for re-keys) and export both private and public keys to all the COOP KS's. This is required if the primary KS goes down; the rey-keys sent by the newly elected primary KS will still be properly verified by the GM.
2. Election between the KS's is based on the highest-priority value configured. If they are the same, it is based on highest IP address. Configuring priorities for selecting the primary KS is recommended for easy setup and troubleshooting.
3. Periodic ISAKMP keepalive (dead peer detection, or DPD) must be configured on the COOP KS's so that the primary KS can track the secondary KS state.
4. Re-key configuration, policies, and anti-replay configurations must be the same for all KS's. This is not done automatically in the current phase of GET VPN; you must ensure that the values match.

As noted, all secondary KS's must use the same key-pair to be able to sign the keys. This requires that you export the keys at the primary KS and import them at the secondary. You may export/import keys via a TFTP server or terminal. The commands to do this via a terminal session are:

```
crypto key export rsa GETVPN_KEYS pem terminal 3des CISCO1234
crypto key import rsa GETVPN_KEYS pem exportable terminal CISCO1234
```

Note that the pass-phrase at the end of every command is used to encrypt the keys. The first command produces a PEM output on the screen. You must paste it to the input of the second command. Do not forget to specify the proper RSA key-pair ID in these commands. You configure general redundancy settings using the GDOI policy configuration and the following commands:

https://t.me/learningnets

```
crypto gdoi group GETVPN_GROUP
 server local
  address ipv4 150.1.1.1
   redundancy
     local priority 100
      peer address ipv4 150.1.2.2
```

Remember that GDOI protocol is protected by an ISAKMP SA. For KS-to-KS communication, this SA should always remain up, because there is no GDOI data SA. For an SA to come up, both KS's must authenticate each other, using either pre-shard-keys or RSA signatures. Make sure that you have configured proper authentication information on every KS for every other KS. A final note worth mentioning about ISAKMP is that when you enable ISAKMP keepalives (DPD), you should use the periodic keyword, because it will enable the KS to detect the loss of bidirectional communication faster. The GM configuration is minimal; it requires only adding new KS addresses and configuring proper authentication, such as adding pre-shared-keys for the secondary KS's.

Note that although some IOS versions support the KS and GM role on the same router, it is not officially supported and working on all IOS versions. Some adjustments can be made by running the GM role in a custom VRF and KS role in the default VRF, because at this moment the KS role is not VRF aware.

Configuration

```
R1:
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp keepalive 10 periodic
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
crypto ipsec transform-set GETVPN esp-3des esp-md5-hmac
!
crypto ipsec profile GETVPN_PROFILE
 set transform-set GETVPN
!
crypto key generate rsa general-keys label GETVPN_KEYS modulus 1024 exportable
access-list 100 permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
```

https://t.me/learningnets

```
!
crypto gdoi group GETVPN_GROUP
  identity number 123
  server local
   rekey authentication mypubkey rsa GETVPN_KEYS
   rekey transport unicast
   address ipv4 136.1.18.1
   redundancy
    local priority 100
    peer address ipv4 136.1.28.2
   sa ipsec 1
    profile GETVPN_PROFILE
    match address ipv4 100
!
crypto key export rsa GETVPN_KEYS pem terminal 3des CISCO1234
```

R2:

```
crypto key import rsa GETVPN_KEYS pem terminal CISCO1234
!
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
crypto isakmp keepalive 10 periodic
!
crypto isakmp key CISCO address 0.0.0.0 0.0.0.0
crypto ipsec transform-set GETVPN esp-3des esp-md5-hmac
!
crypto ipsec profile GETVPN_PROFILE
 set transform-set GETVPN
!
access-list 100 permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
!
crypto gdoi group GETVPN_GROUP
  identity number 123
  server local
   rekey authentication mypubkey rsa GETVPN_KEYS
   rekey transport unicast
   address ipv4 136.1.28.2
   sa ipsec 1
    profile GETVPN_PROFILE
    match address ipv4 100
   redundancy
```

https://t.me/learningnets

```
    local priority 50
    peer address ipv4 136.1.18.1
```

R3:

```
crypto isakmp policy 100
 authentication pre-share
 encryption 3des
 hash md5
 group 2
!
!
crypto isakmp key CISCO address 136.1.18.1
crypto isakmp key 0 CISCO address 136.1.28.2
!
crypto gdoi group GETVPN_GROUP_GM
 identity number 123
  server address ipv4 136.1.18.1
  server address ipv4 136.1.28.2
!
crypto map GETVPN_MAP local-address FastEthernet0/0.38
crypto map GETVPN_MAP 10 gdoi
  set group GETVPN_GROUP_GM
!
interface FastEthernet0/0.38
  crypto map GETVPN_MAP
```

Verification

Verify that both key servers have established an ISAKMP SA and negotiated the primary KS:

```
R1#show crypto gdoi ks coop
Crypto Gdoi Group Name :GETVPN_GROUP
        Group handle: 2147483650, Local Key Server handle: 2147483650
 Local Address: 136.1.18.1
Local Priority: 100
Local KS Role: Primary    , Local KS Status: Alive
        Local KS version: 1.0.3
        Primary Timers:
                Primary Refresh Policy Time: 20
                Remaining Time: 3
```

https://t.me/learningnets

```
                Antireplay Sequence Number: 149

        Peer Sessions:
        Session 1:
                Server handle: 2147483651 Peer Address: 136.1.28.2
                Peer Version: 1.0.3  Peer Priority: 50
               Peer KS Role: Secondary , Peer KS Status: Alive
                Antireplay Sequence Number: 4
 IKE status: Established
                Counters:
                    Ann msgs sent: 146
                    Ann msgs sent with reply request: 2
                    Ann msgs recv: 2
                    Ann msgs recv with reply request: 1
                    Packet sent drops: 1
                    Packet Recv drops: 1
                    Total bytes sent: 75931
                    Total bytes recv: 659
!
!R2#show crypto gdoi ks coop
Crypto Gdoi Group Name :GETVPN_GROUP
        Group handle: 2147483650, Local Key Server handle: 2147483650
 Local Address: 136.1.28.2
Local Priority: 50
Local KS Role: Secondary , Local KS Status: Alive
        Local KS version: 1.0.3
        Secondary Timers:
                Sec Primary Periodic Time: 30
                Remaining Time: 20, Retries: 0
                Invalid ANN PST recvd: 0
                New GM Temporary Blocking Enforced?: No
                Antireplay Sequence Number: 6

        Peer Sessions:
        Session 1:
                Server handle: 2147483651 Peer Address: 136.1.18.1
                Peer Version: 1.0.3  Peer Priority: 100
Peer KS Role: Primary   , Peer KS Status: Alive
                Antireplay Sequence Number: 157
 IKE status: Established

                Counters:
                    Ann msgs sent: 2
                    Ann msgs sent with reply request: 2
                    Ann msgs recv: 154
                    Ann msgs recv with reply request: 2
```

https://t.me/learningnets

```
                    Packet sent drops: 2
                    Packet Recv drops: 0
                    Total bytes sent: 715
                    Total bytes recv: 80307
```

Verify that GM is successfully registered with its primary KS:

```
R3#show crypto gdoi gm

Group Member Information For Group GETVPN_GROUP_GM:
    IPSec SA Direction       : Both ACL Received From KS      : gdoi_group_GETVPN_GROUP_GM_temp_acl

    Group member             : 136.1.38.3       vrf: None
       Registration status   : Registered Registered with        : 136.1.18.1
       Re-registers in       : 3480 sec
       Succeeded registration: 1
       Attempted registration: 1 Last rekey from        : 136.1.18.1

       Last rekey seq num    : 1
       Unicast rekey received: 1
       Rekey ACKs sent       : 1
       Rekey Rcvd(hh:mm:ss)  : 00:00:01
```

Shut down the primary KS VLAN 18 interface so that the secondary KS takes over the primary role and performs re-keying:

```
R1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.R1(config)#interface gigabitEthernet0/0.18
R1(config-subif)#shutdown
```

After DPD messages expire the ISAKMP SA, the secondary KS should become the primary KS and log the following messages to indicate this:

```
%GDOI-5-COOP_KS_TRANS_TO_PRI: KS 136.1.28.2 in group GETVPN_GROUP transitioned to Primary (Previous Primary = 136.1.
%GDOI-3-COOP_KS_UNREACH: Cooperative KS 136.1.18.1 Unreachable in group GETVPN_GROUP. IKE SA Status = Failed to esta
```

Verify that R2 is the primary KS:

```
R2#show crypto gdoi ks coop
Crypto Gdoi Group Name :GETVPN_GROUP
        Group handle: 2147483650, Local Key Server handle: 2147483650
```

https://t.me/learningnets

```
Local Address: 136.1.28.2

        Local Priority: 50        Local KS Role: Primary    , Local KS Status: Alive
        Local KS version: 1.0.3
        Primary Timers:
                Primary Refresh Policy Time: 20
                Remaining Time: 6
                Antireplay Sequence Number: 13

        Peer Sessions:
        Session 1:
                Server handle: 2147483651
                Peer Address: 136.1.18.1
                Peer Version: 1.0.3
                Peer Priority: Unknown  Peer KS Role: Primary    , Peer KS Status: Dead
                Antireplay Sequence Number: 169
IKE status: Failed

                Counters:
                    Ann msgs sent: 2
                    Ann msgs sent with reply request: 2
                    Ann msgs recv: 166
                    Ann msgs recv with reply request: 2
                    Packet sent drops: 9
                    Packet Recv drops: 0
                    Total bytes sent: 715
                    Total bytes recv: 87187
```

Traffic between Loopback0 subnets is identified by the protected ACL, but with the current setup there is only one GM and two KS's. Because the KS's do not perform encryption/decryption (there are no IPsec SAs), traffic between Loopback0 subnets of R3-R2 and R3-R1 will not be successful: R3 will send encrypted packets, but these will be dropped by the KS, so traffic is unidirectional.

```
R3#ping 150.1.2.2 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.2.2, timeout is 2 seconds:
Packet sent with a source address of 150.1.3.3 .....
Success rate is 0 percent (0/5)
!
!R3#show crypto ipsec sa

interface: FastEthernet0/0.38
    Crypto map tag: GETVPN_MAP, local addr 136.1.38.3
```

https://t.me/learningnets

```
   protected vrf: (none)
   local  ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   remote ident (addr/mask/prot/port): (150.1.0.0/255.255.0.0/0/0)
   current_peer 0.0.0.0 port 848
     PERMIT, flags={origin_is_acl,} #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
#pkts decaps: 0, #pkts decrypt: 0, #pkts verify: 0
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.38.3, remote crypto endpt.: 0.0.0.0
     path mtu 1500, ip mtu 1500, ip mtu idb FastEthernet0/0.38
     current outbound spi: 0xA15012E6(2706379494)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0xA15012E6(2706379494)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2097, flow_id: NETGX:97, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (2972)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0xA15012E6(2706379494)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2098, flow_id: NETGX:98, sibling_flags 80000040, crypto map: GETVPN_MAP
        sa timing: remaining key lifetime (sec): (2972)
        Kilobyte Volume Rekey has been disabled
        IV size: 8 bytes
        replay detection support: N Status: ACTIVE


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

ASA Basic Clientless SSL VPN

Task

- Configure ASA4 to permit WebVPN connections on its VLAN100 interface using the port number 443:
  - Create a URL-List entry named **CISCO** pointing to the URL **http://www.cisco.com**.
  - Filter WebVPN connections and only allow users to connect on TCP port 80 to sites in **.com** domains.
  - Configure R1 as a DNS server and the firewall to use it for WebVPN.
  - Create a DNS entry in R1 for **www.cisco.com** resolving to R1's Loopback0 address.
- Authenticate sessions locally with the username **WEBVPN** and the password **CISCO1234**.
- Ensure that this user is only allowed to use the configured WEBVPN group.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.

Overview

SSL VPN is a technology that allows tunneling IP packets over an SSL connection (TCP protected by SSL). Using regular SSL connections on the default port 443 allows for better compatibility and firewall/NAT traversal. An SSL VPN connection can be either client-based (requires a software installed on the client side) or clientless (requires only a compatible web browser on the client side).

Clientless SSL VPN technology (aka WebVPN) allows remote users to access corporate resources solely using their web browsers. This feature does not require

https://t.me/learningnets

any client installation on the user's side. The ASA firewall is used by default as a proxy for HTTP/HTTPs connections initiated by the user. The user connects to the ASA firewall using an HTTPS connection and authenticates with a username and password, digital certificate, or both. The firewall opens a special (customizable) portal page to the user, which emulates a browser with a URL address bar. The user may enter URLs for company resources, and the firewall resolves them using a configured DNS server and downloads the requested pages. Optionally, the firewall may apply a URL filter to restrict access to certain corporate resources, or even disallow URL entry completely, providing the user with a list of static bookmarks.

In addition to simple HTTP/HTTPs browsing, the user may access Microsoft file shares or use the SMTPs/POP3s/IMAPs proxy for email access. If you have TCP-based applications requiring custom TCP ports access, WebVPN lets you download a special Java applet that implements port-forwarding (more about this in a separate scenario). A more advanced feature is the Smart-Tunnel option, which allows automatic port-redirection for any application written using Winsock library (covered in a separate scenario).

Aside from TCP-based connectivity, the firewall may allow users to download and run application-specific plugins, such as RDP, VNC, or Telnet applets. These plugins let the user initiate a connection using the specified protocol to any destination on the protected network.

The following are the basic WebVPN configuration steps. First define some global properties, and then configure a group-policy and a connection-profile (tunnel-group):

Step 1:

Enable WebVPN on the required interfaces by entering webvpn configuration mode and using the command `enable <INTERFACE_NAMEIF>` . You may need to define a port number used for WebVPN connection with the command `port <port-number>` , which defaults to 443. Notice that this no longer affects ASDM access on the same port, which it did in previous code versions. ASDM and clientless SSL VPN can run concurrently on the same port, because ASDM uses a dedicated URL in the format of **https://IP_ADDRESS/admin**.

You may define HTTP and HTTPs proxy servers used by the firewall to hop-off incoming connections to the internal resources. The command `http(s)-proxy <IP> <PORT> username <Proxy-User> password <Proxy-Password>` defines the proxy and the credentials used for proxy authentication. You may specify the URL to be excluded from proxying by using the command `http(s)-proxy <IP> <PORT> exclude <URL-Spec>` . When defining the URL-Spec, you may use meta-characters:

https://t.me/learningnets

- `?` for any symbol
- `*` for any string
- `[ ]` to define a range of characters
- `![ ]` to define a negative range of characters (not in the list)

Another useful command is **`tunnel-group-list enable`**, which allows the user to select the tunnel-group to be used for login and authentication from the VPN portal. An alternative to this is to create a specific URL for each connection-profile and ask the user to use it. WebVPN cannot classify incoming connections and match a connection profile automatically, unless you provide a hint for it or certificate authentication is being used. This is why, without the explicit/implicit group selection, the firewall will revert to using the default remote-access tunnel-group of **DefaultWEBVPNGroup** for any incoming connection.

Step 2:

Define a DNS server group. This step is required for WebVPN because the firewall acts as a proxy server and needs to resolve the names in the URLs presented by the users (unless they use IP addressing). To define a DNS server group, you must enable DNS lookups on the interface-facing DNS servers using the command **`dns domain-lookup <INTERFACE_NAMEIF>`**. After this, you may define a server-group using the command **`dns server-group <NAME>`**. Define at least one DNS server IP address and the default domain under server-group configuration mode.

Step 3:

https://t.me/learningnets

Define a group policy for WebVPN connection. The group policy command is generic, but you may define specific WebVPN attributes under the webvpn configuration mode of the group-policy. Here you may define functions enabled by this policy, such as url-entry and file-browsing. The default values are inherited from the `DfltGrpPolicy` group-policy. A policy setting of interest here is the filter command, which applies a webtype access-list to the connections requested by users. A webtype access-list permits or denies connections based on URL strings in the format `{permit|deny} url [method]://<pattern>` or `{permit|deny} tcp <SUBNET> <MASK> <PORT>` . Basically, you can filter based on destination URL and protocol or destination subnet and TCP port. The URL method could be http, https, cifs, nfs, pop3, imap4, smtp, and ftp only. The pattern specifies the URL, and you may use the same meta-characters defined in the proxy section for the pattern. For example, the URL may look like `http://www.ci*.com:80[12]?` . In addition to the method/pattern entries, you can also use the regular ACL entries, such as `permit tcp any host 2.2.2.2 eq 80` . The webtype access-list is for webvpn what the vpn filter access-list is for client-based VPN (SSL or IKE).

Step 4:

Define a connection-profile (tunnel-group) for WebVPN users. The tunnel-group type should be `remote-access` , and you must associate the previously defined group-policy with the tunnel-group. Finally, you should define some basic WebVPN attributes. The first one is the authentication method for remote clients, username/password or digital certificates or both. The `group-alias <NAME>` command defines a name to be used on the login page for the tunnel-group selection. The DNS group created earlier can be associated with the tunnel group using the command `dns-group <NAME>` under the `tunnel-group <NAME> webvpn-attributes` definition.

According to the scenario requirements, we define a group-lock attribute to limit the user to the group `WEBVPN` only. Additionally, we create a bookmark list (URL-List) to be used with the WebVPN portal. We must use the ASDM for this; version 8.x does not support the legacy `url-list` command and bookmarks are now stored as XML files, so you must use the GUI/ASDM to define them. Notice the webtype access-list created for this task. It allows connecting to sites in the `.com` domain on port number 80 only. To verify the scenario, we configure R1 as a DNS server and define some names (in **.com** and **.net** domains) there.

Configuration

ASA4:

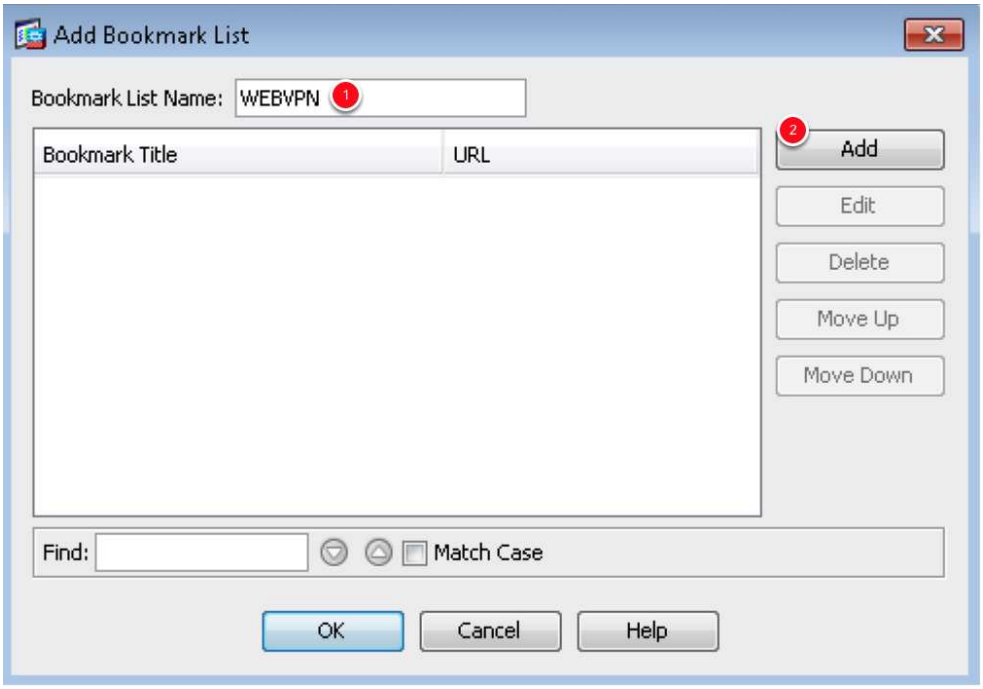
```
webvpn
```

https://t.me/learningnets

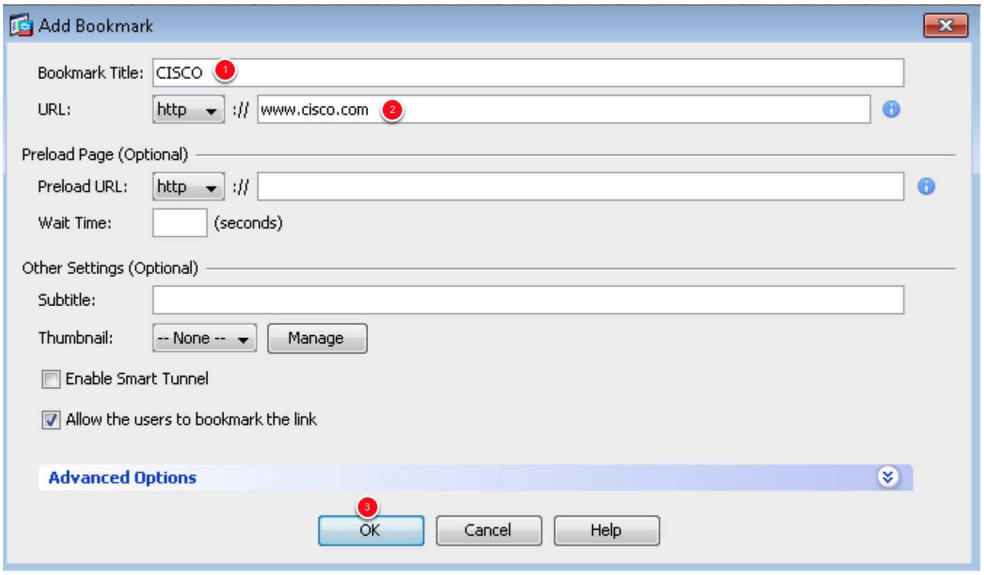
```
 port 443
 enable VLAN100
 tunnel-group-list enable
!
access-list WEBACCESS webtype permit url http://*.com:80
!
dns domain-lookup VLAN18
dns server-group WEBVPN
 name-server 136.1.18.1
 domain-name cisco.com
!
group-policy WEBVPN internal
group-policy WEBVPN attributes
 vpn-tunnel-protocol ssl-clientless
 webvpn
  url-entry enable
  filter value WEBACCESS
!
tunnel-group WEBVPN type remote-access
tunnel-group WEBVPN general-attributes
 default-group-policy WEBVPN
!
tunnel-group WEBVPN webvpn-attributes
 group-alias WEBVPN enable
 authentication aaa
 dns-group WEBVPN
!
username WEBVPN password CISCO1234
username WEBVPN attributes
 group-lock value WEBVPN
!
http server enable
http 0 0 VLAN100
```

To configure the URL-List, from the **TEST-PC-A** connect to the ASA using ASDM; the same username configured for WebVPN can be used. Navigate in the GUI as follows: **Configuration** -> **Remote Access VPN** -> **Clientless SSL VPN Access** -> **Portal** -> **Bookmarks**. Assign the list a name in the **Bookmark List Name** field.

https://t.me/learningnets



Click **Add**, give the entry a name in the **Bookmark Title** field, and enter **www.cisco.com** in the **URL** field.



Assign the URL-List to the group-policy from CLI:

ASA4:

```
group-policy WEBVPN attributes
 webvpn
  url-list value WEBVPN
```

R1:

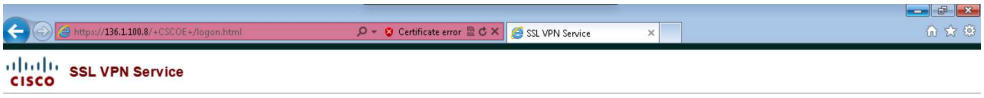
```
ip dns server
```

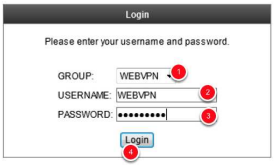
https://t.me/learningnets

```
ip host cisco.com 150.1.1.1
ip host www.cisco.com 150.1.1.1
ip host www.google.com 150.1.1.1
ip host www.ripe.net 150.1.1.1
```

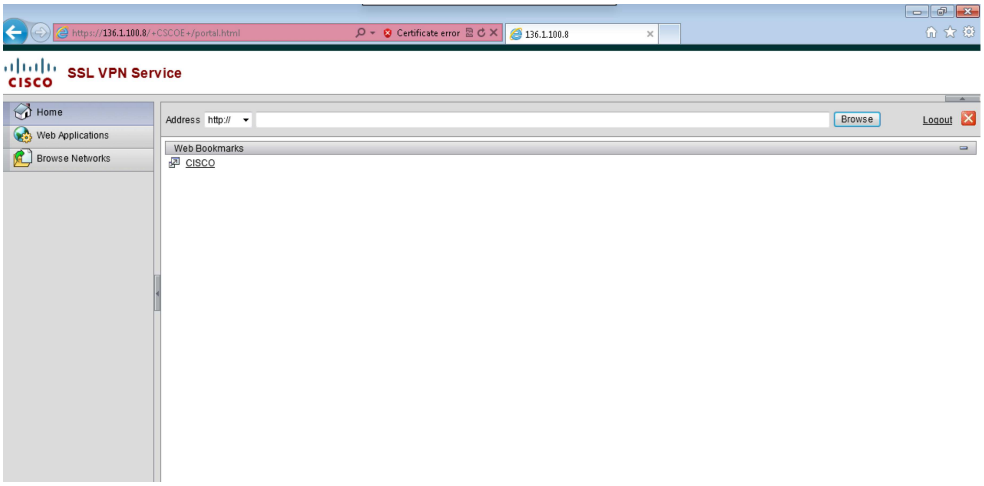
Verification

From **TEST-PC-A**, use Internet Explorer to browse to **https://136.1.100.8**; you will see the drop-down box that allows you to select the connection-profile. However, there is only one created. Authenticate using the configured username and password:



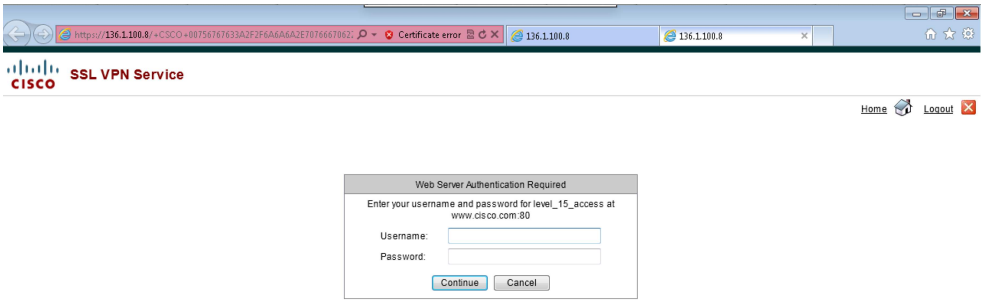


Upon successful authentication, you should be presented with the following page:

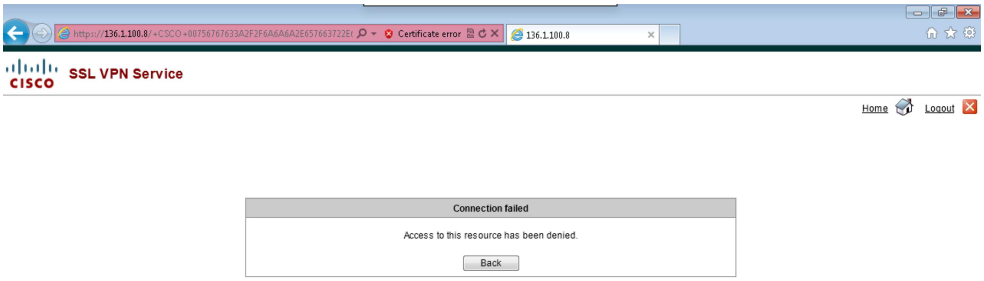


Click the box with an arrow to the right of the bookmark to open the hyperlink in a new page. This will redirect you to R1's HTTP server page, as configured, asking you for authentication:

https://t.me/learningnets



Try typing **www.google.com** in the Address field of the VPN portal and click **Browse**; it should again direct you to the same page of R1's HTTP server, because the **.com** domain is allowed by the webtype access-list. Try typing **www.ripe.net** in the Address field of the VPN portal and click **Browse**; it should redirect you to the same page of R1's HTTP server, but because the **.net** domain is denied by the webtype access-list, you get an error message:



Check the VPN session on the ASA:

```
ASA4# show vpn-sessiondb webvpn filter name WEBVPN
 Session Type: WebVPN
 Username     : WEBVPN
                Index        : 17
Public IP    : 136.1.100.100 Protocol     : Clientless
License      : AnyConnect Premium
Encryption   : RC4                   Hashing      : SHA1
Bytes Tx     : 135601                Bytes Rx     : 19501
Group Policy : WEBVPN                Tunnel Group : WEBVPN

Login Time   : 17:34:23 UTC Thu May 16 2013
Duration     : 0h:00m:34s
Inactivity   : 0h:00m:00s
```

https://t.me/learningnets

```
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

ASA Clientless SSL VPN Port Forwarding

Task

- Configure ASA4 to permit WebVPN connections on its VLAN100 interface using the port number 443:
  - A remote user connecting to the local TCP port 20023 is redirected to R1's Loopback0 on port 23.
  - The applet should be automatically downloaded upon user login.
  - Define a custom name for the downloaded applet.
  - Filter WebVPN connections and only allow users to connect on TCP port 23.
- Authenticate sessions locally with the username **WEBVPN** and the password **CISCO1234**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

WebVPN port forwarding, also called **Application Access**, allows various TCP-based applications to work over the secure SSL connection. The firewall pushes a special Java applet to the client's PC. This applet could be pre-configured to listen on the selected ports of the remote PC and forward local connections to these ports over the WebVPN connection to the servers on the private network. For example, you may configure the applet to listen on ports 20143, 20110, and 20025 and forward connections made to the localhost on these ports to the respective

https://t.me/learningnets

corporate servers. Port forwarding is supported on all major operating platforms, including Microsoft Windows, Linux, and Mac OS. You need administrative privileges to start the port-forwarding on the local machine, and you cannot use certificate-based authentication with the Java applet. Notice that port-forwarding only supports applications using static TCP ports; any applications that open secondary data channels such as FTP are not supported. Configuring port-forwarding consists of the following steps:

1. Define a port-forwarding list under the global webpvn configuration mode using the command **`port-forward <LIST-NAME> <local-port> <inside-server> <server-port> <description>`** . You may define multiple statements under one list.
2. Associate the port-forwarding list with the group-policy for the respective tunnel-group. Use the following commands under the webvpn section of the group-policy: **`port-forward {enable|auto-start} <LIST-NAME>`** and **`port-forward name [Descriptive name for the Application Access]`** . Notice the difference between the **`enable`** and **`auto-start`** options. In the former, you must activate application access manually. In the latter, the applet is started automatically upon user logon.
3. Ensure that your webvpn filter permits the connections specified in the port-forwarding configuration. You only need this step if you apply filtering to the connections of WebVPN clients.

Configuration

ASA4:

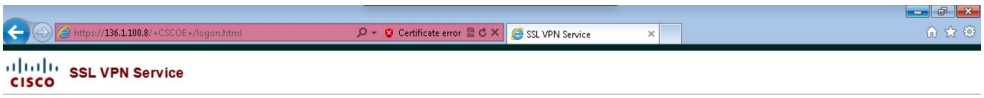
```
webvpn
 port 443
 enable VLAN100
 tunnel-group-list enable
 port-forward TELNET_R1 20023 150.1.1.1 telnet
!
access-list WEBACCESS webtype permit tcp any eq 23
!
group-policy WEBVPN internal
group-policy WEBVPN attributes
 vpn-tunnel-protocol ssl-clientless
 webvpn
  port-forward name “Port Forwarding”
  port-forward auto-start TELNET_R1
  filter value WEBACCESS
!
tunnel-group WEBVPN type remote-access
```

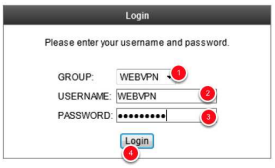
https://t.me/learningnets

```
tunnel-group WEBVPN general-attributes
 default-group-policy WEBVPN
!
tunnel-group WEBVPN webvpn-attributes
 group-alias WEBVPN enable
 authentication aaa
!
username WEBVPN password CISCO1234
```

Verification

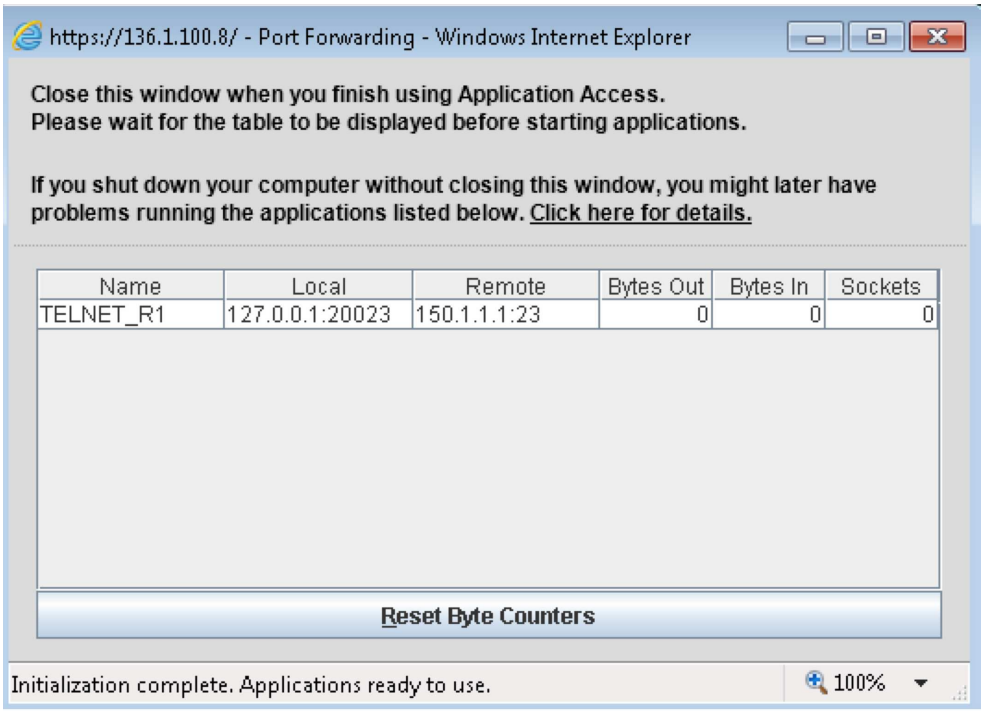
From **TEST-PC-A**, use Internet Explorer to browse to **https://136.1.100.8**; you will see the drop-down box that allows you to select the connection-profile. However, there is only one created. Authenticate using the configured username and password:



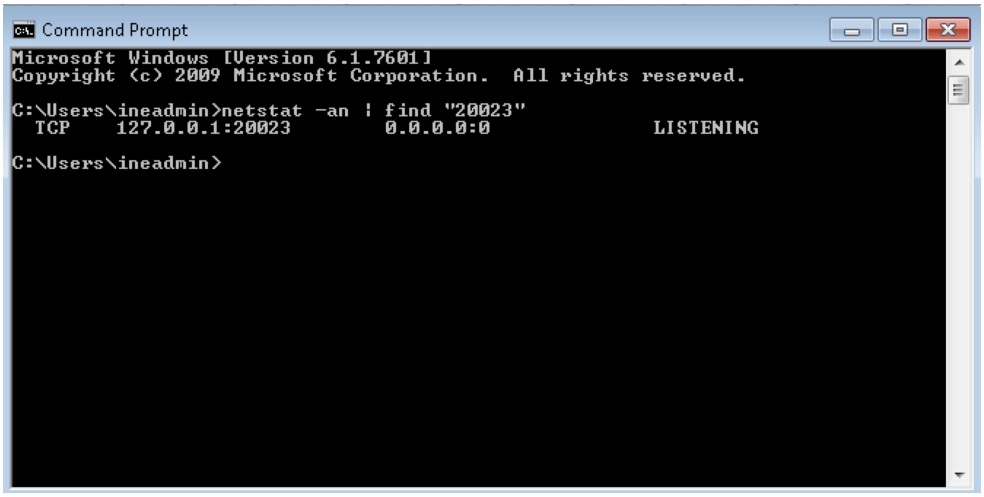


Upon successful connection, you will see the Java applet automatically downloaded and started (do not close the window, or port forwarding configuration will stop functioning):

https://t.me/learningnets

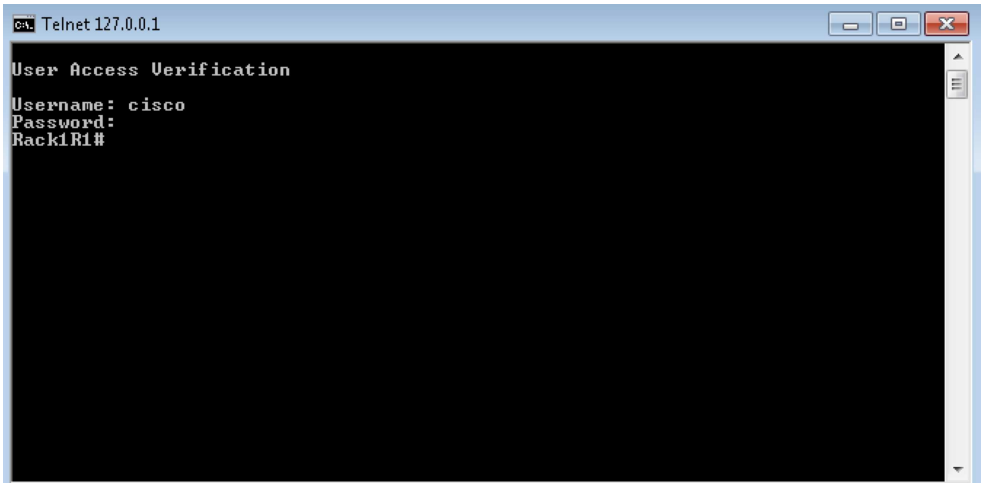


Verify that the Java applet started a service listening on TCP port 20023 on **TEST-PC-A**:



From a command prompt, type **telnet 127.0.0.1 20023**. You should be redirected to R1's telnet server; authenticate using username/password **cisco/cisco**:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

ASA Clientless SSL VPN Smart Tunnel

Task

- Configure ASA4 to permit WebVPN connections on its VLAN100 interface using the port number 443:
  - Configure the ASA so that telnet traffic is transparently proxied across the SSL VPN tunnel.
  - The applet should be automatically downloaded upon user login.
- Authenticate sessions locally with the username **WEBVPN** and the password **CISCO1234**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The Smart Tunnel feature is a replacement for the legacy Port Forwarding. Instead of relying on the statically configured port-forwarding rules implemented in a Java applet, the technology implements transparent TCP proxy for applications written using Winsock2 library. A special application is pushed down to the client and used as an application wrapper. This wrapper intercepts TCP library calls and redirects them across the ASA firewall as through a proxy. The operations are completely transparent to the end-user application, which is not aware of any additional security proxy in the path.

https://t.me/learningnets

You have the option of limiting the applications eligible to use the Smart Tunnel feature. In the firewall, you can define a smart-tunnel application list using the global webvpn command `smart-tunnel list <NAME> <AppName> <AppPath> [platform {windows|mac}] [hash <sha1-hash>]`. Here, `<AppName>` is the unique application identifier string, `<AppPath>` is the path to the application executable in the client's PC, and hash defines a SHA1 40-character's hash value used to verify the user's application integrity. The application path is matched against the full application path starting from the right. Thus, you may simply specify the application name and extension, such as `telnet.exe`, to match the program found in any directory. For extra security, you may use the full application path, such as `c:\program files\cisco\app\program.exe`, but this will disallow transparent proxy function for the same application located at a different path.

You may assign the smart-tunnel application list to the group-policy using the command `smart-tunnel {enable|auto-start} <LIST-NAME>`. As in the port forwarding case, the `enable` option lets the user activate the feature manually, and `auto-start` activates the feature automatically.

Configuration

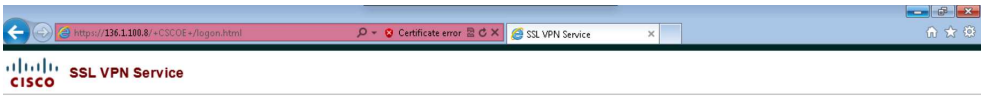
https://t.me/learningnets

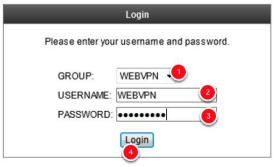
ASA4:

```
webvpn
 port 443
 enable VLAN100
 tunnel-group-list enable
 smart-tunnel list SMART_TUNNEL TELNET telnet.exe
!
group-policy WEBVPN internal
group-policy WEBVPN attributes
 vpn-tunnel-protocol ssl-clientless
 webvpn
  smart-tunnel auto-start SMART_TUNNEL
!
tunnel-group WEBVPN type remote-access
tunnel-group WEBVPN general-attributes
 default-group-policy WEBVPN
!
tunnel-group WEBVPN webvpn-attributes
 group-alias WEBVPN enable
 authentication aaa
!
username WEBVPN password CISCO1234
```

Verification

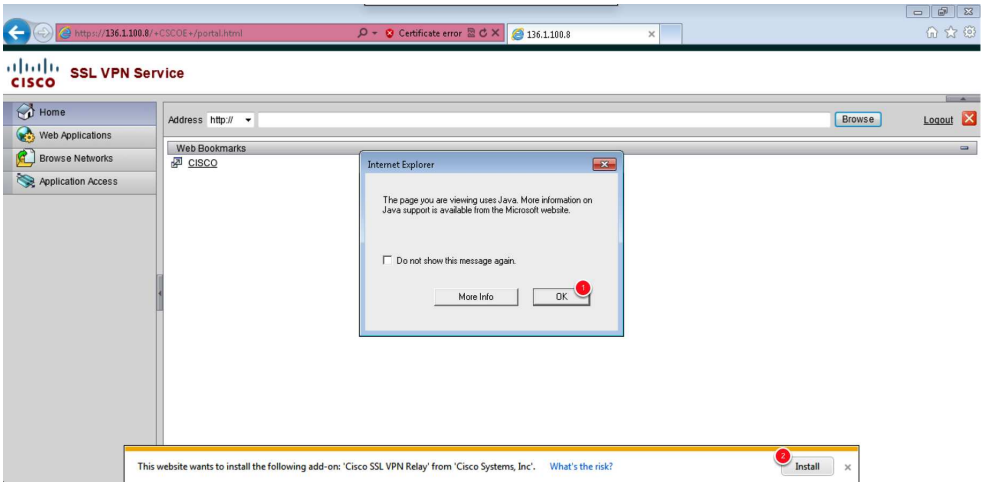
From **TEST-PC-A**, use Internet Explorer to browse to **https://136.1.100.8**; you will see the drop-down box that allows you to select the connection-profile. However, there is only one created. Authenticate using the configured username and password:



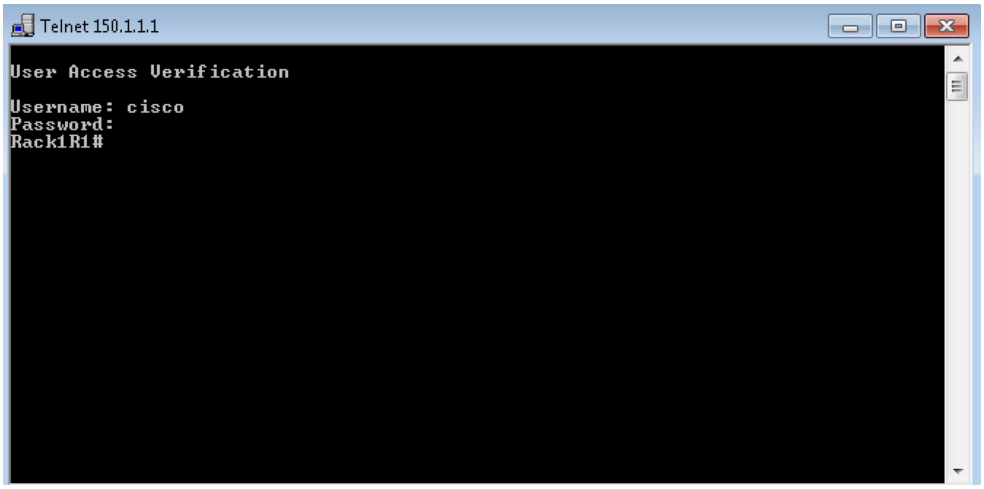


https://t.me/learningnets

Upon successful connection, the browser will prompt you to install Cisco SSL VPN Relay, the smart-tunnel applet:

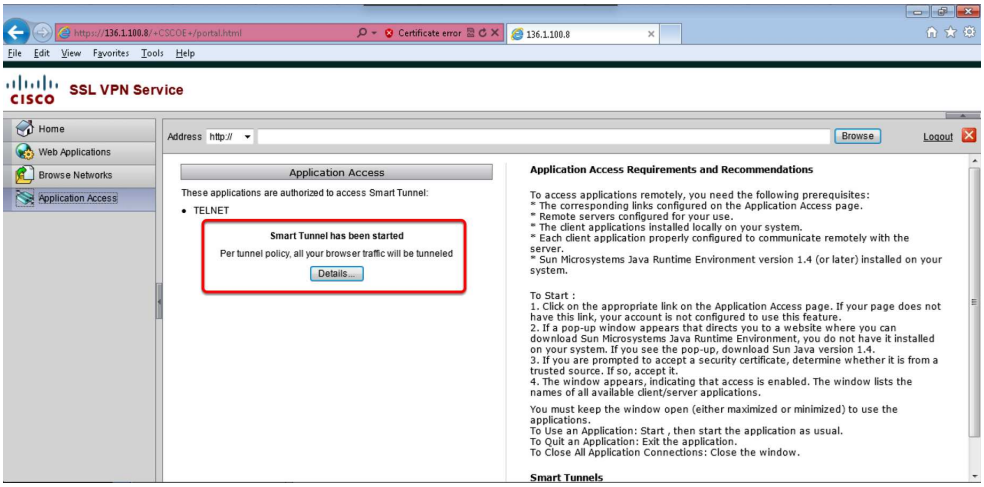


When the add-on is installed, you may be required to log out and log back in. Then open a command prompt, type **telnet.exe**, and in the new window type **open 150.1.1.1**; you should be successfully connected to R1. Authenticate using username/password **cisco**/**cisco**:



To ensure that Smart Tunnel is enabled, in the SSL VPN portal browse to the **Application Access** tab and verify that it says **Smart Tunnel has been started**:

https://t.me/learningnets



If you run a **`debug ip tcp transactions`** on R1, you will notice that ASA4 proxies the TCP telnet session; it actually initiates the session in the name of the client:

```
R1#debug ip tcp transactions
TCP special event debugging is on
Reserved port 0 in Transport Port Agent for TCP IP type 0
TCP0: state was LISTEN -> SYNRCVD [23 -> 136.1.18.8(57417)]
TCP: tcb 31654704 connection to 136.1.18.8:57417, peer MSS 1380, MSS is 516
TCP: Selective ack is disabled from the CLI
TCP: sending SYN, seq 1075989519, ack 2119415753
TCP0: Connection to 136.1.18.8:57417, advertising MSS 1380
TCP0: state was SYNRCVD -> ESTAB [23 -> 136.1.18.8(57417)]
TCB31654704 setting property TCP_TOS (11) 2AB0AAD4
TCB31654704 setting property TCP_RTRANSTMO (3
Rack1R1#1) 3003A850
TCB31654704 setting property TCP_GIVEUP (35) 3003A854
TCP388: bad seg from 136.1.18.8 -- outside window: port 23 seq 2119415774 ack 1075989578 rcvnxt 2119415784 rcvwnd 40
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

ASA AnyConnect SSL VPN with AAA

Task

- Configure a connection-profile named **SSLVPN** on the ASA4 firewall.
- A user named **SSLVPN** with the password of **CISCO1234** should be allowed to log in to this group when connecting to the ASA via HTTPS.
- As soon as the user logs in, AnyConnect VPN software should be pushed back to the user's PC.
- Allocate the connecting user an IP address from the pool 20.0.0.0/24 and ensure that the user is allowed to reach only the subnet 136.1.18.0/24 through the tunnel.
- Enable the protocol that compensates for the negative effect of the TCP protocol.
- Ensure that the amount of information sent over the VPN link is reduced to the minimum.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Cisco's implementation of client-based SSL VPN is called AnyConnect Secure Mobility Client and adds new functionality (compared with the Cisco IKEv1 VPN Client), such as automatic client and profile download and multiple OS support. Generally, you don't have to download and install a VPN client to connect to the security gateway because the client is downloaded automatically and the connection

https://t.me/learningnets

is established upon successful authentication. The initial connection is being made using any web browsers and the proper URL. After the client is downloaded, it obtains an IP address from the server, configures split-tunneling, and inherits all settings from the security gateway.

Configuring SSL VPN in the ASA firewall is similar to previous WebVPN configuration: you define connection-profiles (tunnel-groups), group-policies, and VPN users. First, you must enable the feature on the required firewall interface (usually the outside) and specify the port used for SSL VPN and the AnyConnect image location for the required operating system of the client. Here is a minimal configuration:

```
webvpn
 port 443
 enable outside
 tunnel-group-list enable
 anyconnect image flash:/anyconnect-win-2.3.0254-k9.pkg
 anyconnect enable
```

Note the **`tunnel-group-list command`** . As in the WebVPN section, it allows the connecting user to select the connection-profile from the AnyConnect interface. This is very similar to a normal WebVPN configuration, except for two **`anyconnect`** commands that load the SSL VPN client into memory and enable AnyConnect functionality globally.

When configuring a group policy for SSL VPN clients, specify the **`vpn-tunnel-protocol ssl-client`** attribute. This replaces the legacy syntax of **`svc enable`** and **`svc required`** and allows the user's browser to download the VPN client automatically upon successful connection. Another difference from the regular WebVPN configuration is configuring split-tunnel policy under the group policy. We normally do this for EzVPN policies, but SSL VPN group policy is a hybrid of regular WebVPN and EzVPN settings.

All other AnyConnect settings are defined under the webvpn attribute group of the group-policy. One particularly interesting feature is the DTLS (datagram TLS) protocol, which is the version of TLS that does not require the use of TCP, because it is based on UDP. Using TCP as underlying transport has negative drawbacks resulting from the TCP congestion management behavior. This has a poor effect on VoIP quality and some other real-time applications. Using DTLS allows for two simultaneous connections (both secure), one TLS and one DTLS. DTLS is automatically enabled by default when SSL is enabled (client or clientless). Another useful option may be the use of connection compression, based on the deflate algorithm. Although it is only helpful when browsing large amounts of textual

https://t.me/learningnets

information, you may need it sometimes. Remember the extra load on the firewall when you enable compression.

Before you define the tunnel-group, create an address pool for the SSL VPN clients. Use the global `ip local pool` command for this. The pool is assigned to the tunnel-group under general-attributes using the command `address-pool` . In addition to this, you may want to enable redistribution of the static routes into your IGP, to make sure all network devices may reach dynamic RRI entries. All other tunnel-group settings are similar to the ones you use with regular WebVPN groups, such as group-alias and authentication aaa (if you don’t use certificate-based authentication).

Configuration

ASA4:

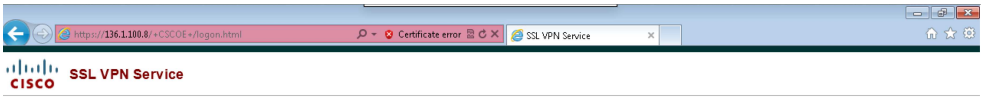
```
webvpn
 port 443
 enable VLAN100
 tunnel-group-list enable
 anyconnect image disk0:/anyconnect-win-3.0.11042-k9.pkg
 anyconnect enable
!
access-list SPLIT_TUNNEL standard permit 136.1.18.0 255.255.255.0
!
group-policy SSLVPN internal
group-policy SSLVPN attributes
 vpn-tunnel-protocol ssl-client
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value SPLIT_TUNNEL
 webvpn
  anyconnect ssl dtls enable
  anyconnect keep-installer installed
  anyconnect ssl compression deflate
!
ip local pool ANYCONNECT_POOL 20.0.0.1-20.0.0.254
!
router eigrp 111
 redistribute static metric 10000 100 255 1 1500
!
tunnel-group SSLVPN type remote-access
tunnel-group SSLVPN general-attributes
 default-group-policy SSLVPN
 address-pool ANYCONNECT_POOL
!
```

https://t.me/learningnets

```
tunnel-group SSLVPN webvpn-attributes
 group-alias SSLVPN enable
 authentication aaa
!
username SSLVPN password CISCO1234
```

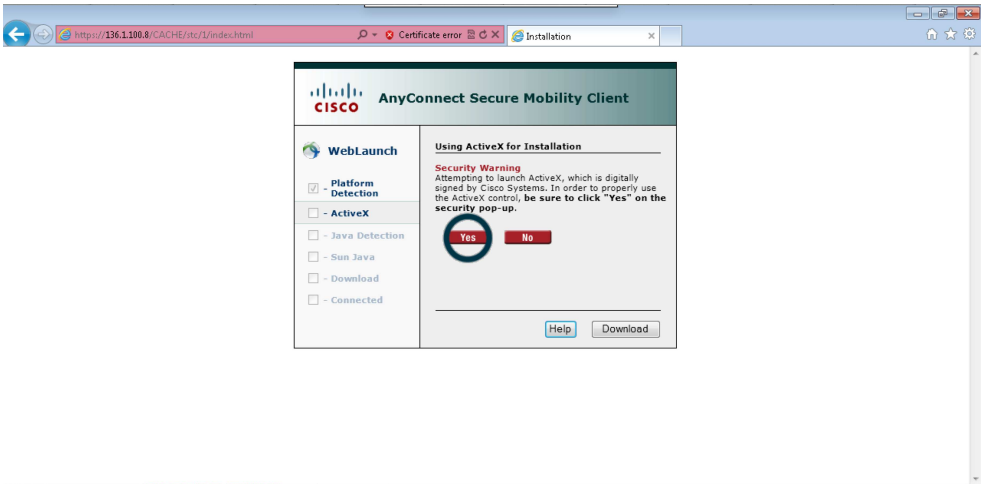
Verification

From **TEST-PC-A**, use Internet Explorer to browse to **https://136.1.100.8**; you will see the drop-down box that allows you to select the connection-profile. However, there is only one created. Authenticate using the configured username and password:



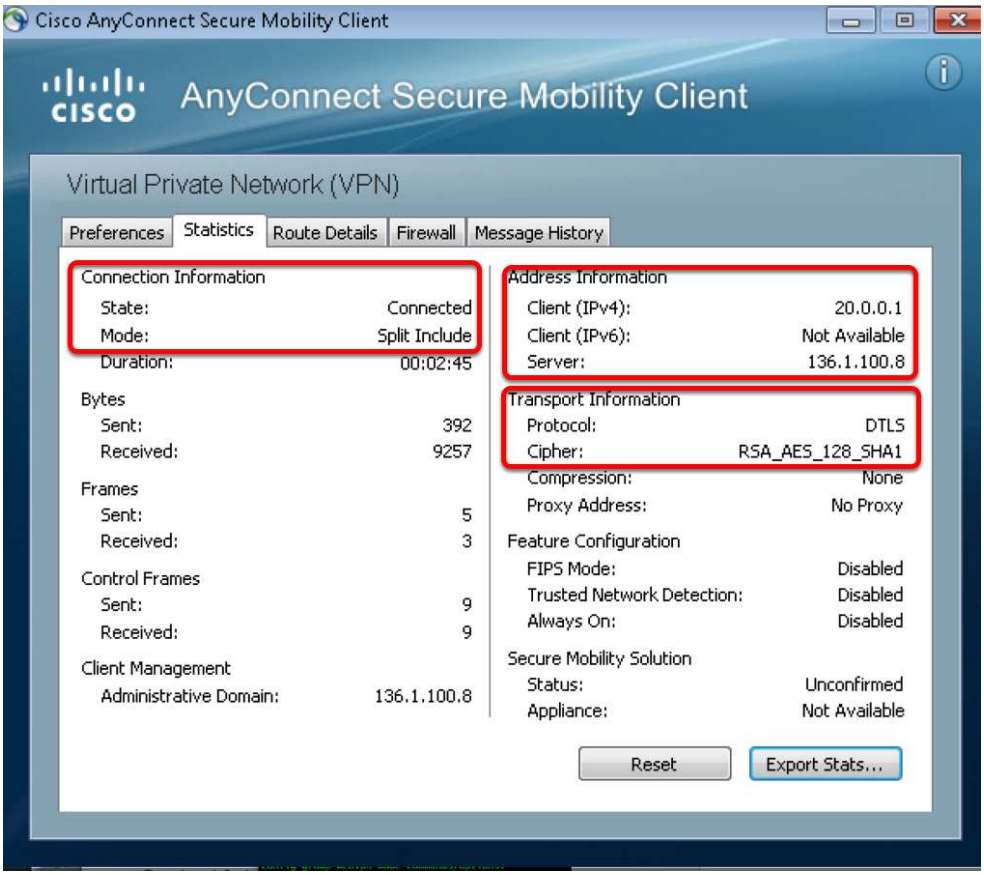


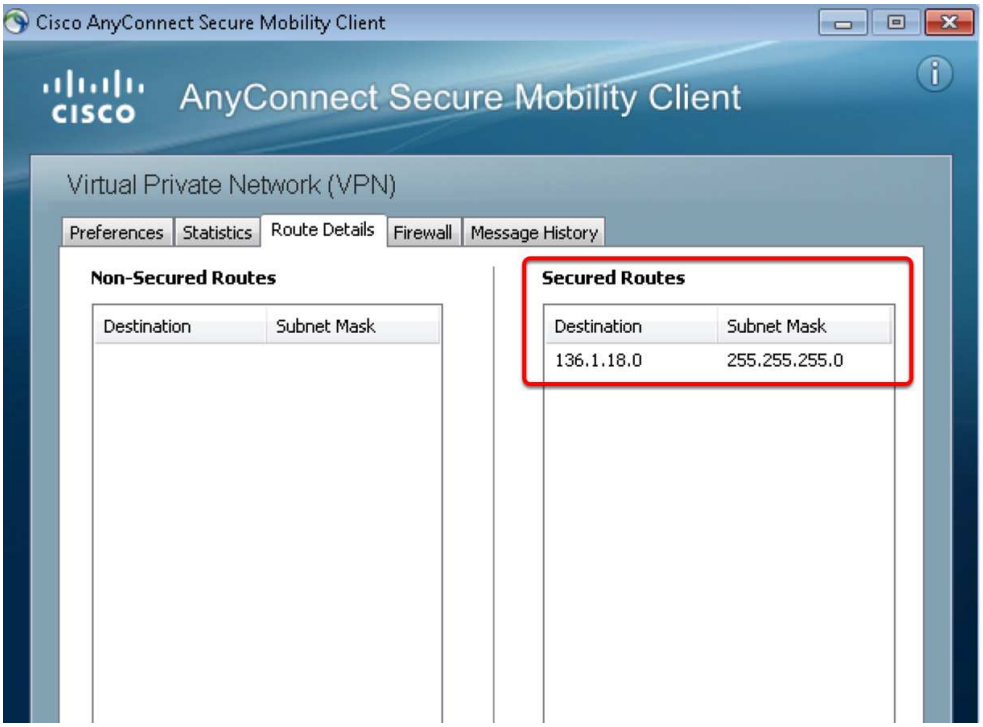
Upon successful connection, you may be prompted to download and install AnyConnect Secure Mobility Client if not already installed on the **TEST-PC-A**:



https://t.me/learningnets

Connection should be successful. Right-click the AnyConnect icon, select **Open Anyconnect**, click **Advanced**, and go to the **Statistics** tab to see the assigned IP address and negotiated algorithms. Go to the **Route Details** tab to verify split-tunneling configuration:





https://t.me/learningnets

Verify that the connection shows as established on the ASA as well:

```
ASA4# show vpn-sessiondb anyconnect filter name SSLVPN

Session Type: AnyConnect
Username     : SSLVPN
               Index        : 18 Assigned IP  : 20.0.0.1                Public IP    : 136.1.100.100
Protocol     : AnyConnect-Parent SSL-Tunnel DTLS-Tunnel
License      : AnyConnect Premium
Encryption   : RC4 AES128             Hashing      : none SHA1
Bytes Tx     : 9244                   Bytes Rx     : 993
Group Policy : SSLVPN                 Tunnel Group : SSLVPN

Login Time   : 17:40:07 UTC Thu May 16 2013
Duration     : 0h:04m:03s
Inactivity   : 0h:00m:00s
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none
```

Verify the RRI EIGRP redistribution:

```
R1#show ip route eigrp | i D EX
D EX     20.0.0.1 [170/284160] via 136.1.18.8, 00:04:49, GigabitEthernet0/0.18
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

ASA AnyConnect SSL VPN with PKI

Task

- Configure a connection-profile named **SSLVPN** on the ASA4 firewall to be used for AnyConnect sessions:
  - Use both certificate and username/password authentication.
  - The username should be automatically completed from the OU field of the certificate.
  - The username should not be visible to the connecting AnyConnect user.
  - Any incoming AnyConnect session with an OU of **ccielab** should be automatically mapped to the SSLVPN connection profile.
  - Allocate the connecting user an IP address from the pool 20.0.0.0/24.
- Use ASA4 as the CA server, enrolling the user **test-user**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

First we must activate the CA server on ASA and enroll the user; it is most important to specify the OU value of **ccielab**. Modify the connection profile to use double authentication (both certificate and username/password) using the command `authentication certificate aaa` , auto-complete the username value from the certificate OU field using the command `username-from-certificate OU` , and hide the

https://t.me/learningnets

username in the AnyConnect window using the command **`pre-fill-username ssl-client hide .`**

For an incoming IKE or SSL VPN session to be automatically mapped to a connection profile based on certificate attributes, we must first enable the functionality with the command **`tunnel-group-map enable rules .`** Then we must configure a certificate map that matches on required certificate fields using the command **`crypto ca certificate map <NAME> <SQ_NR> .`** Finally, activate the map at webvpn level using the command **`certificate-group-map <MAP_NAME> <SQ_NR> <CONNECTION_PROFILE> .`**

Configuration

ASA4:

```
crypto ca server
 no shutdown
!
crypto ca server user-db add test-user dn CN=test-user.ine.com,OU=ccielab
crypto ca server user-db allow user test-user
!
username ccielab password CISCO1234
!
webvpn
 port 443
 enable VLAN100
 anyconnect image disk0:/anyconnect-win-3.0.11042-k9.pkg
 anyconnect enable
 certificate-group-map SSLVPN_MAP 10 SSLVPN
!
group-policy SSLVPN internal
group-policy SSLVPN attributes
 vpn-tunnel-protocol ssl-client
 split-tunnel-policy tunnelall
 webvpn
  anyconnect keep-installer installed
!
tunnel-group SSLVPN type remote-access
tunnel-group SSLVPN general-attributes
 default-group-policy SSLVPN
 address-pool ANYCONNECT_POOL
 username-from-certificate OU
!
tunnel-group SSLVPN webvpn-attributes
```

https://t.me/learningnets

```
 authentication certificate aaa
 pre-fill-username ssl-client hide
!
tunnel-group-map enable rules
!
crypto ca certificate map SSLVPN_MAP 10
 subject-name attr ou eq ccielab
!
ip local pool ANYCONNECT_POOL 20.0.0.1-20.0.0.254
!
router eigrp 111
 redistribute static metric 10000 100 255 1 1500
```

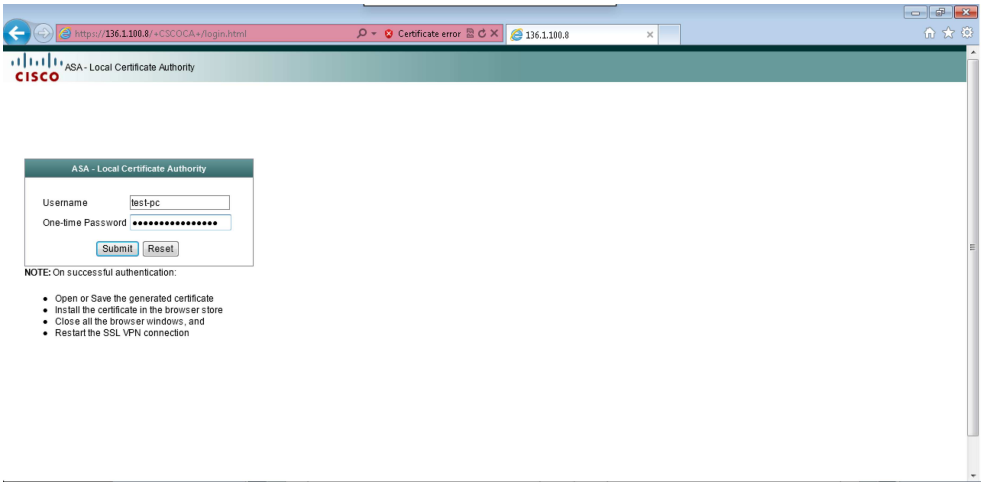
Verification

Find the OTP issued by the ASA for the **test-pc** user:

```
ASA4(config)# crypto ca server user-db show-otp user test-pc
Username: test-pcOTP: C0D8B0972050157B


Enrollment Allowed Until: 17:59:53 UTC Sun May 19 2013
```

From **TEST-PC-A**, browse to **https://136.1.100.8/+CSCOCA+/enroll.html** to download the client certificate, and authenticate with the **test-pc** user and the above OTP:



https://t.me/learningnets

Install the client certificate and CA certificate using the procedure described in the PKI section. When asked for the password that protects the private-key, enter the above OTP. After the certificate is installed, ensure that it is valid by modifying the Windows system time accordingly. Using AnyConnect, start the session to the ASA; you will receive a prompt to enter the password for the user **ccielab**:

https://t.me/learningnets

Image not readable or empty
/uploads/workbooks/images/originals/52b367708d7758183a45221eb413de74b14da739.png

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IOS AnyConnect SSL VPN with AAA

Task

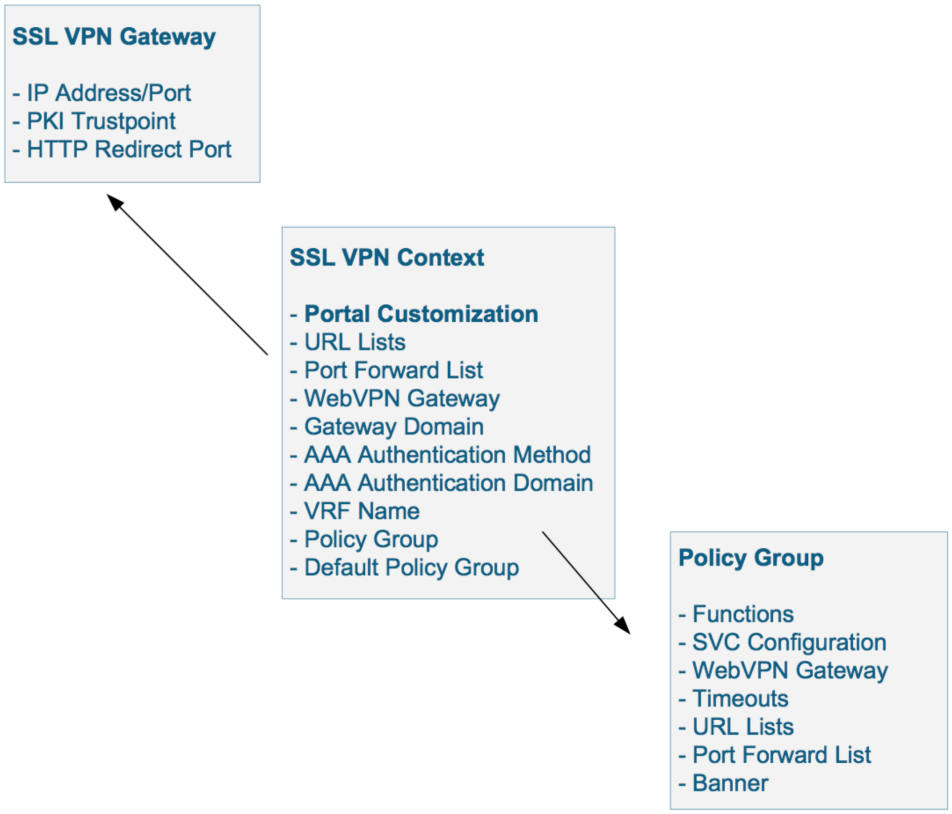
- Configure R3 as AnyConnect VPN server:
  - Users should connect using the URL **http://150.1.3.3/SSLVPN**.
  - Allocate the IP addresses using the local pool 21.0.0.0/24, and make sure traffic is only tunneled to the subnet 136.1.33.0/24.
  - Authenticate users against the local database and create a new user named **SSLUSER** in the local database. Ensure that this user is not allowed to log in under any other VPN context.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

AnyConnect Secure Mobility Client is supported in both IOS and ASA platforms. The configuration syntax is pretty different between the two platforms. In this scenario, we'll discuss the IOS configuration steps. Here is a visual outline of SSL VPN configuration steps in IOS router:

https://t.me/learningnets



The first step is to load the VPN client image in the flash memory and make the router aware of it. This is accomplished with the command `webvpn install svc flash:/<IMAGE_FILE_NAME>` , using the image name appropriate for the client operating system. There could be up to three different images loaded in the router with the sequence numbers configured, such as an image for Win32, Linux, and MacOSX. Next, enable HTTP secure server using the command `ip http secure-server .`

Let's start the configuration overview of the SSL VPN gateway component. This is the point where VPN clients connect to and establish a secure tunnel. The gateway defines the local IP address and port, redirection policy (HTTP to HTTPS), and the local PKI trustpoint to be used for SSL certificate presentation and signing. If you don't want to use any special PKI trustpoint, don't specify any and the one generated for the local HTTPS server will be picked up.

https://t.me/learningnets

The next logical concept is SSL VPN context, the object that defines an isolated VPN group. Typically, visual portal customization and WebVPN policy settings are associated with the contexts: AAA lists for authentication, URL lists, port-forward lists, colors, banners. A context must be associated with a parent gateway that will be used to access it. When someone accesses the gateway, they can select the context to log in to and authenticate themself. Instead of selecting the context, a domain name mapped to that context could be specified in the URL, using the syntax `https://<IP_ADDRESS>/DOMAIN` . The domain name is a part of the context configuration.

Authentication settings are an important part of the context configuration. Typically, you select the AAA lists for SSL VPN authentication and AAA authentication domain if you're implementing the group-lock functionality. The domain name will be appended to the user's name before it is sent to the AAA subsystem. For example, if the username is **USER** and the domain is **DOMAIN**, the string sent to the AAA subsystem will be **USER@DOMAIN**. You should configure usernames in this format, even if they are configured locally.

Every context defines a policy that essentially defines settings enforced for the connecting users. Some policy settings, mostly related to WebVPN, such as port-forwarding, are defined under the context. However, a more flexible way to configure policy is by using policy groups. Additionally, AnyConnect settings such as address pool are only configurable under a policy-group. There could be multiple policy groups configured under the context, but only one group is defined as the default group. Others could be associated with users by using external RADIUS server and special attribute defining the policy group name.

Finally, if you are wondering about routing issues, SSL VPN server automatically injects static routes into the local routing table for every endpoint connected, so don't forget to redistribute those into IGP if you need end-to-end connectivity, or provide other means of propagating the routing information.

Note

> Due to a bug in the running IOS code which randomly appears, this task may not work as expected.

Configuration

R3:

```
webvpn install svc flash:/anyconnect-win-3.0.11042-k9.pkg
```

https://t.me/learningnets

```
!
aaa new-model
aaa authentication login SSLVPN local
username SSLUSER@SSLVPN password CISCO
!
ip http server
ip http secure-server
ip local pool ANYCONNECT_POOL 21.0.0.1 21.0.0.254
!
webvpn gateway SSLVPN_GATEWAY
 ip interface Loopback0 port 443
 http-redirect port 80
 inservice
!
webvpn context SSLVPN
 title "AnyConnect VPN on R3"
!
policy group DEFAULT_POLICY
 functions svc-required
 svc address-pool ANYCONNECT_POOL
 svc split include 136.1.33.0 255.255.255.0
default-group-policy DEFAULT_POLICY
aaa authentication list SSLVPN
aaa authentication domain @SSLVPN
gateway SSLVPN_GATEWAY domain SSLVPN
inservice
```

Verification

First, verify configuration:

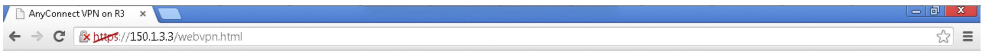
```
R3#show webvpn install status svc
SSLVPN Package SSL-VPN-Client version installed:
CISCO STC win2k+
3,0,11042
Hostscan Version 3.0.11042
Wed 03/23/2011 16:54:44.04
!
!R3#show webvpn gateway SSLVPN_GATEWAY
 Admin Status: up
Operation Status: up
Error and Event Logging: DisabledIP: 150.1.3.3, port: 443
 HTTP Redirect port: 80
SSL Trustpoint: TP-self-signed-3052748620
FVRF Name not configured
```

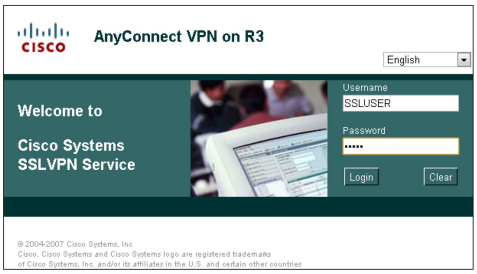
https://t.me/learningnets

```
!
!R3#show webvpn context
Context Name: SSLVPN
Admin Status: up
Operation Status: up
Error and Event Logging: Disabled
CSD Status: Disabled
Certificate authentication type: All attributes (like CRL) are verified AAA Authentication List: SSLVPN
AAA Authorization List not configured
AAA Accounting List not configured AAA Authentication Domain: @SSLVPN
Authentication mode: AAA authentication
Default Group Policy: DEFAULT_POLICY
Associated WebVPN Gateway: SSLVPN_GATEWAY Domain Name: SSLVPN

Maximum Users Allowed: 1000 (default)
NAT Address not configured
VRF Name not configured
Virtual Template not configured
```

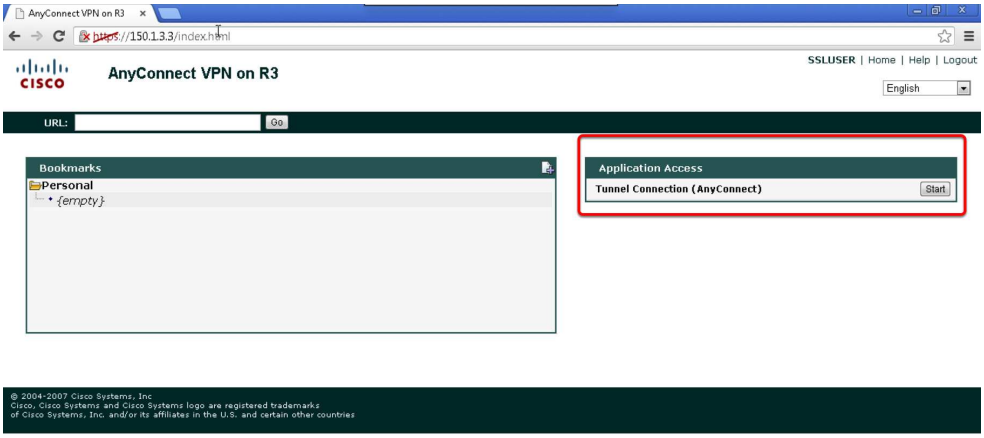
From **TEST-PC-A**, navigate to **http://150.1.3.3/SSLVPN** using Internet Explorer and authenticate using the username/password **SSLUSER/CISCO**:





After successful authentication, the VPN portal appears. Click **Start** under **Application Access** to start downloading and connecting with AnyConnect, if this is not already installed on **TEST-PC-A**:

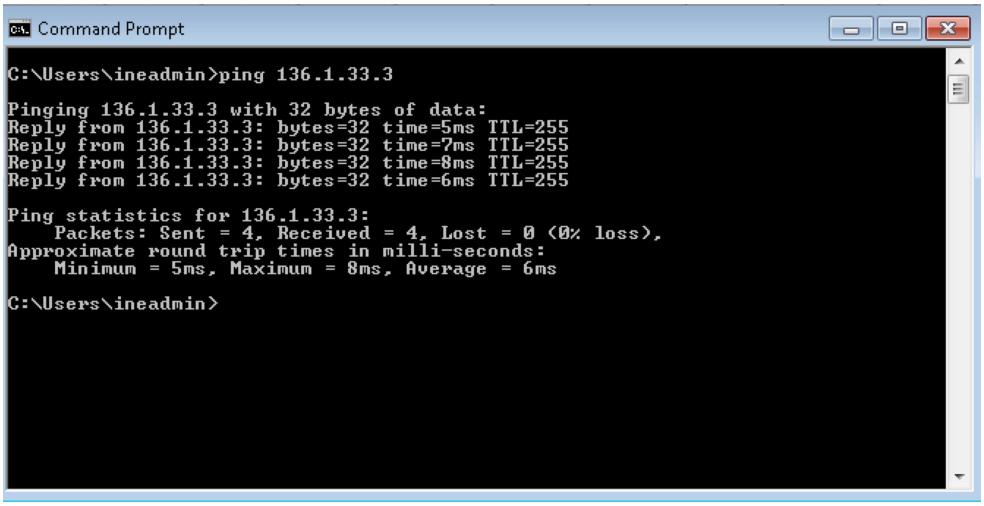
https://t.me/learningnets



Verify that the connection was successfully established on the router:

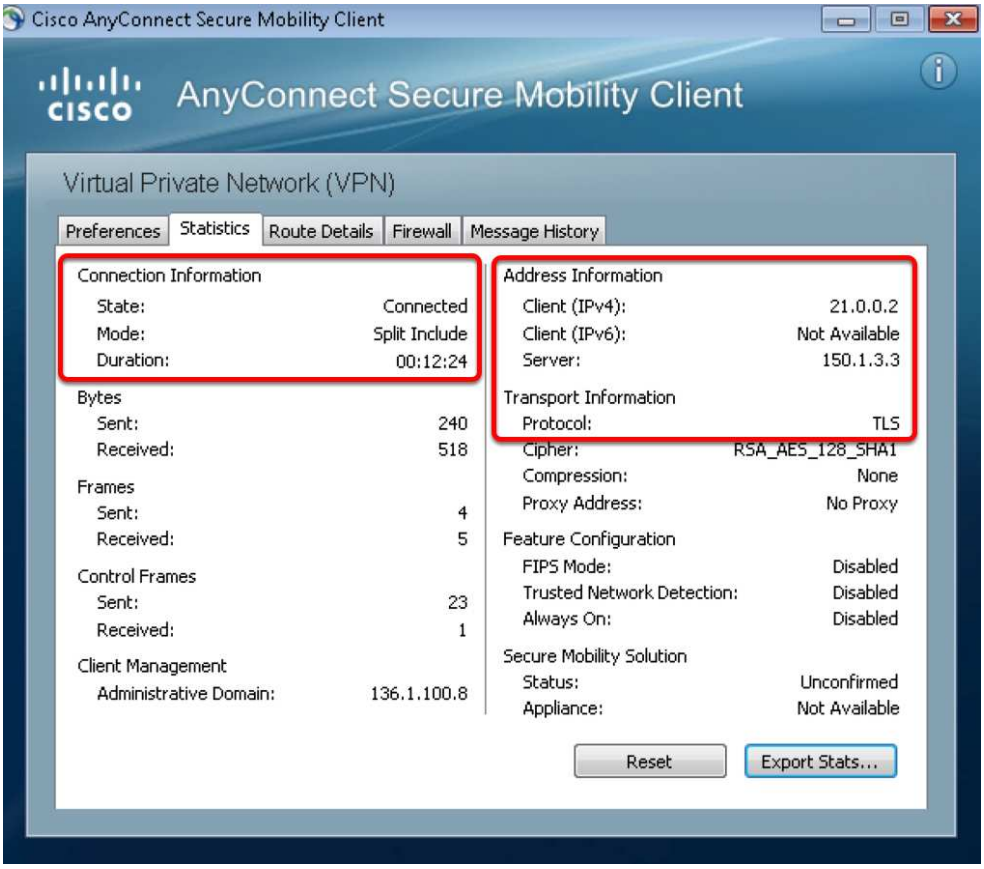
```
R3#show webvpn session context SSLVPN
WebVPN context name: SSLVPN
Client_Login_Name   Client_IP_Address   No_of_Connections   Created   Last_Used
SSLUSER             136.1.100.100               2         00:10:50  00:08:24
```

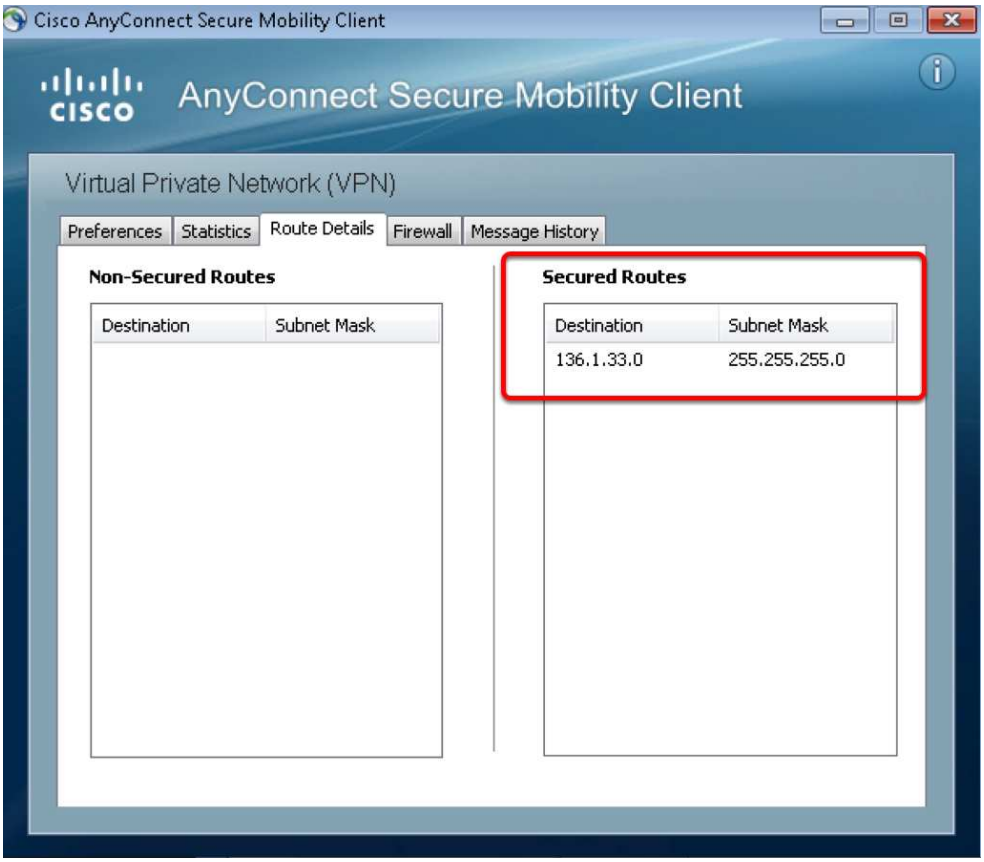
From **TEST-PC-A**, verify that you have connectivity with VLAN 33 subnet:



From **TEST-PC-A**, right-click the AnyConnect icon, click **Open AnyConnect**, and in the **Advanced** section, verify the **Statistics** and **Route Details** tabs:

https://t.me/learningnets





https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IOS Basic Clientless SSL VPN

Task

- Configure R3 to act as a WebVPN proxy for clientless SSL VPN users.
  - The users should access the WebVPN services using the URL **http://150.1.3.3/WEBVPN**.
  - Authenticate users against the local database and create the user **WEBVPN** with the password **CISCO**.
  - Filter WebVPN connections, and only allow users to connect on port 80 to any website in domains **.net** and **.com** during weekdays (Mon-Fri, 9:00am-6:00pm).
- Pre-configure the following two URLs in a list named **Important Links**:
  - **http://www.cisco.com**
  - **http://www.google.com**
- Use the local router as DNS server and make sure that these hostnames resolve to the IP address of 136.1.33.3.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

WebVPN technology was the precursor to the complete AnyConnect VPN technology. Today it is referred to as clientless SSL VPN because it does not

https://t.me/learningnets

require any additional software on the client's PC. The IOS router is used as a proxy for HTTP/HTTPs connections initiated by the user. The user connects to the firewall using a secure HTTP connection and logs in using a valid name/password pair. The router opens a special (customizable) portal page to the user, which emulates a browser with a URL address bar. The user may enter URLs for company resources, and the router resolves them using a configured DNS server or local hostnames and downloads the requested pages. Optionally, IOS may apply a URL filter to restrict access to certain corporate resources, or even disallow URL entry entirely, providing the user with a list of static bookmarks.

Notice that the IOS router does not allow configuration of DNS servers per WebVPN context; the globally configured DNS settings are used instead. Therefore, make sure that you have configured a proper list of external DNS servers using the command `ip name-server <IP_ADDRESS>` and enabled DNS lookups using the command `ip domain-lookup` . Alternatively, you may configure all names locally using the command `ip host` .

In addition to simple HTTP/HTTPs browsing, the user may access Microsoft file shares via CIFS protocol or use the SMTPs/POP3s/IMAPs proxy for email access. If you have a TCP-based application requiring custom TCP ports access, WebVPN lets you download a special Java applet that implements port-forwarding; this is called **thin-client** mode. IOS does not support plug-ins as the ASA, but Smart Tunnels are supported in IOS.

Configuration of the WebVPN features follows the same path as a regular SSL VPN configuration. However, the policy may now need to omit the `functions svc-required` setting, so the AnyConnect does not start every time a user connects. You may still use policy-groups for scalable policy configuration, or use global context settings to specify URL lists, CIFS file shares, and everything else needed. The authentication settings are the same as in the AnyConnect task.

Configuration

R3:

```
aaa new-model
aaa authentication login WEBVPN local
username WEBVPN@WEBVPN password CISCO
!
ip host www.cisco.com 136.1.33.3
ip host www.google.com 136.1.33.3
ip name-server 150.1.3.3
ip domain-lookup
!
```

https://t.me/learningnets

```
ip http server
ip http secure-server
!
webvpn gateway SSLVPN_GATEWAY
 ip interface Loopback0 port 443
 http-redirect port 80
 inservice
!
webvpn context WEBVPN
 title "Clientless SSL VPN at R3"
 url-list "MYURLS"
   heading "Important Links"
   url-text "Cisco" url-value "http://www.cisco.com"
   url-text "Google" url-value "http://www.google.com"
!
time-range "WEEKDAYS"
 periodic weekdays 9:00 to 18:00
!
acl "WEBACL"
 permit url "http://*.com" time-range WEEKDAYS
 permit url "http://*.net" time-range WEEKDAYS
 !
policy group WEBVPN_POLICY
 url-list "MYURLS"
 acl "WEBACL"
default-group-policy WEBVPN_POLICY
aaa authentication list WEBVPN
aaa authentication domain @WEBVPN
gateway SSLVPN_GATEWAY domain WEBVPN
inservice
```

Verification

First, verify configuration:

```
R3#show webvpn context WEBVPN
 Admin Status: up
 Operation Status: up
Error and Event Logging: Disabled
CSD Status: Disabled
Certificate authentication type: All attributes (like CRL) are verified AAA Authentication List: WEBVPN
AAA Authorization List not configured
AAA Accounting List not configured
AAA Authentication Domain: @WEBVPN Authentication mode: AAA authentication
```

https://t.me/learningnets

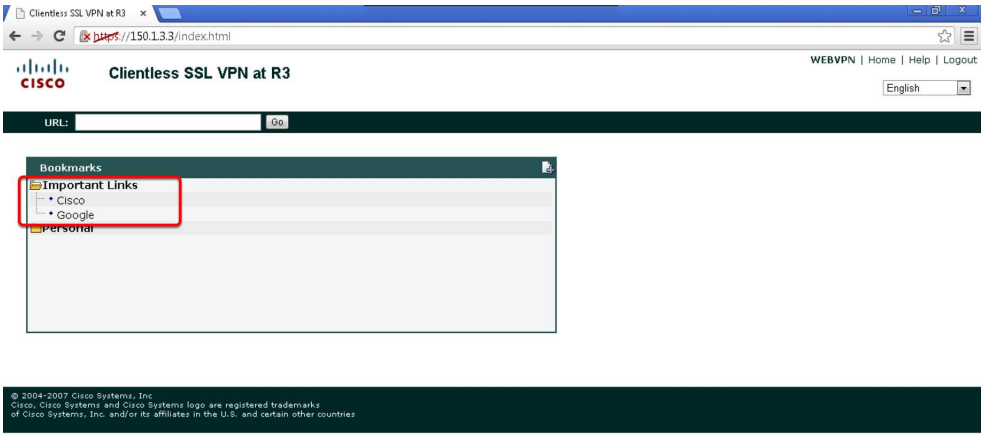
```
Default Group Policy: WEBVPN_POLICY Associated WebVPN Gateway: SSLVPN_GATEWAY
Domain Name: WEBVPN

Maximum Users Allowed: 1000 (default)
NAT Address not configured
VRF Name not configured
Virtual Template not configured
```

From **TEST-PC-A** using Internet Explorer, connect to **http://150.1.3.3/WEBVPN** and authenticate using the username/password **WEBVPN/CISCO**:



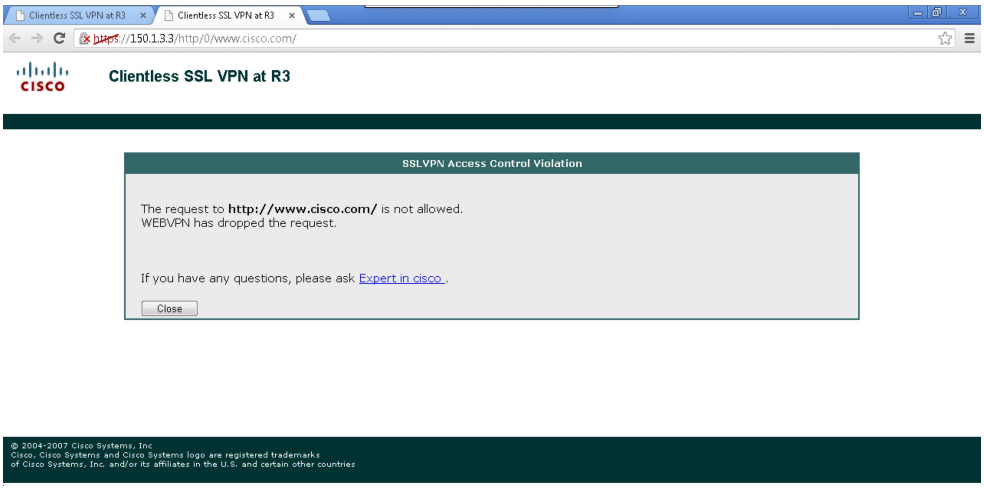
When successfully authenticated, the portal webpage should be presented, and it should contain the two pre-configured URLs:



Because the clock is outside of the configure time-range, browsing the pre-configured links should be disallowed:

https://t.me/learningnets

```
R3#show clock
23:08:32.881 UTC Thu May 16 2013
```



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IOS Clientless SSL VPN Port Forwarding

Task

- Configure R3 to act as a WebVPN proxy for clientless SSL VPN users.
  - The users should be accessing the WebVPN services using the URL **http://150.1.3.3/WEBVPN**.
  - Authenticate users against the local database and create the user **WEBVPN** with the password **CISCO**.
  - When the client logs in, the thin client is automatically started and maps the local port 2080 to port 80 at 136.1.33.3.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **SSL VPN** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Port forwarding, also known as SSL VPN **thin client**, allows remote clients to use TCP applications in a secure manner over the HTTPS connection established with the WebVPN gateway. This solution supports only the applications that do not open dynamic secondary ports. For example, FTP or H.323 signaling will never work using the thin client. Note that all connections made using the thin client application are not subject to the WebVPN access list checks like on the ASA, so you don't need extra entries to permit tunneled connections.

https://t.me/learningnets

Configuration

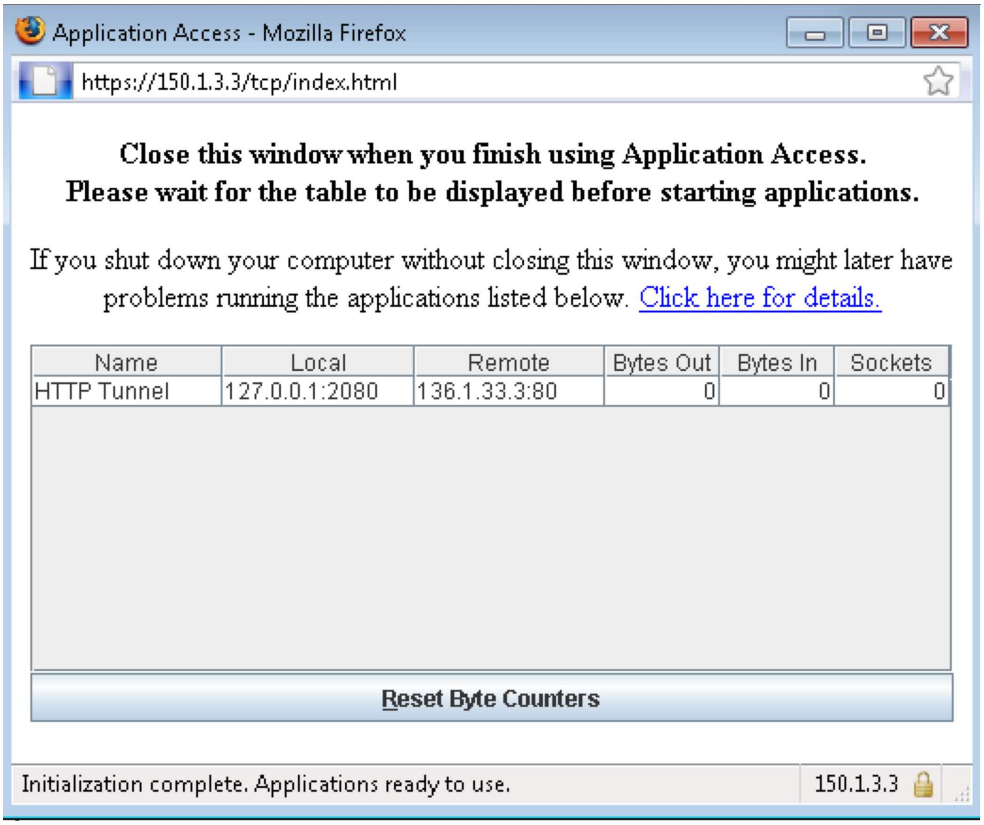
R3:

```
aaa new-model
aaa authentication login WEBVPN local
username WEBVPN@WEBVPN password CISCO
!
ip http server
ip http secure-server
!
webvpn gateway SSLVPN_GATEWAY
 ip interface Loopback0 port 443
 http-redirect port 80
 inservice
!
webvpn context WEBVPN
 port-forward "PORT_FORWARD"
  local-port 2080 remote-server "136.1.33.3" remote-port 80 description "HTTP Tunnel"
 !
policy group WEBVPN_POLICY
 port-forward "PORT_FORWARD" auto-download
default-group-policy WEBVPN_POLICY
aaa authentication list WEBVPN
aaa authentication domain @WEBVPN
gateway SSLVPN_GATEWAY domain WEBVPN
inservice
```

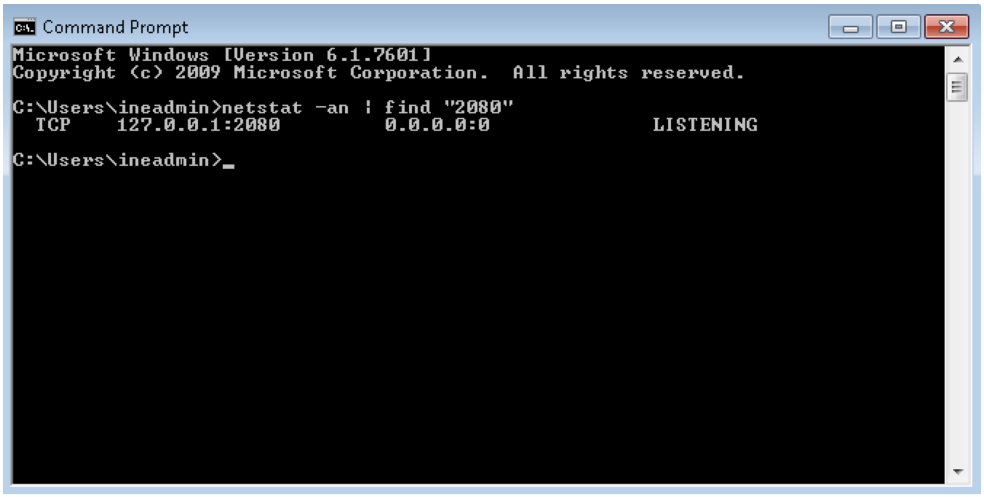
Verification

From **TEST-PC-A**, connect with a web browser to **http://150.1.3.3/WEBVPN** and authenticate using the username/password **WEBVPN/CISCO**. When authenticated and presented with the webvpn portal, the java applet for port forwarding should automatically start:

https://t.me/learningnets

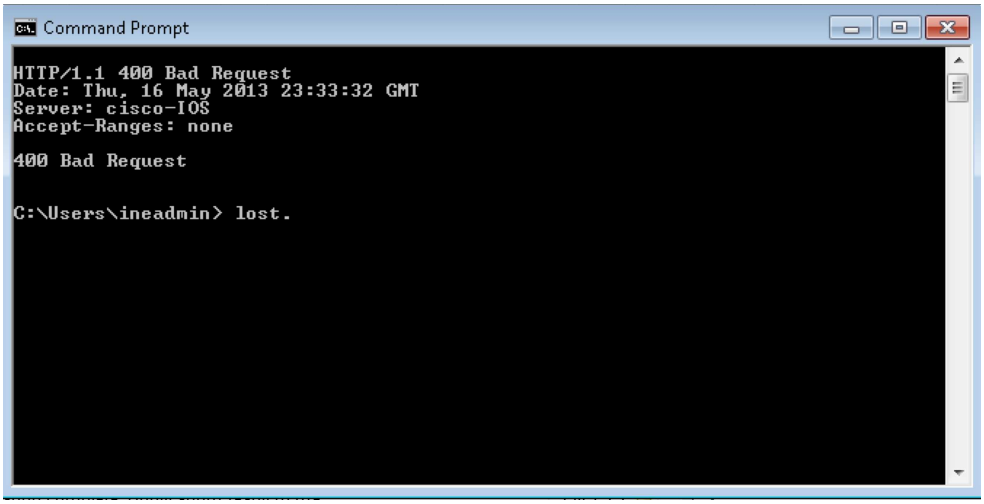


Verify that the applet opened port 2080 and started listening on the local PC:



Using telnet, connect on 127.0.0.1 port 2080; you should reach R3's HTTP service:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv2 L2L Between IOS and ASA with PSK

Task

- Configure IKEv2 IPsec between ASA4 and R2:
  - Negotiate AES-256 and 3DES for encryption, SHA256 and MD5 for integrity, DH groups 5 and 2.
  - Protect traffic between R2 and R1/R3 Loopback0 subnets, negotiate AES-256 and 3DES for encryption, SHA1 and MD5 for integrity.
  - R2 should authenticate with a PSK of **r2psk** and ASA4 with a PSK of **asa4psk**.
  - Limit the number of SAs to two and ensure that R2 generates a log message when the tunnel is UP/DOWN.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

ASA can use different algorithms for PRF and integrity (this is configurable), whereas on the IOS router the PRF algorithm is nonconfigurable and inherited from the integrity algorithm. ASA does not have IKEv2 smart defaults, but as soon as an IKEv2 policy (which is the IKEv2 proposal from IOS) is created, the following default values are used, unless modified:

https://t.me/learningnets

- Encryption: 3des
- Integrity: sha1
- DH group: 2
- PRF: sha1
- Lifetime: 86400 seconds

Based on the above statements and task requirements, PRF must also be configured on the ASA to match the integrity algorithms; otherwise IKEv2 negotiation will fail. Additionally, if multiple values are configured for each option under the same IKEv2 policy (for example NULL, 3DES and AES-256 are specified for encryption in this exact order), the order of configuration is not relevant, because the ASA will automatically convert it to ensure that the most secure algorithms are preferred (it will become AES-256, 3DES, and NULL); IOS works differently and keeps the configuration order. For a newly created IKEv2 ipsec proposal on the ASA (which is the ipsec transform-set from IOS), if no encryption and integrity algorithms are specified, the following default values are used (the same as for the IKEv2 policy):

- Encryption: 3des
- Integrity: sha1

Additionally, if multiple values are configured for each option under the same IKEv2 ipsec proposal (for example NULL, 3DES and AES-256 are specified for encryption in this exact order), the order of configuration is not relevant, because the ASA will automatically convert it to ensure that the most secure algorithms are preferred (it will become AES-256, 3DES, and NULL); IOS works differently and keeps the configuration order. From the configuration perspective on the ASA, IKEv2 LAN-to-LAN tunnel follows the same steps as IKEv1 LAN-to-LAN tunnel, the only difference being that the IKEv1 keyword is replaced with the IKEv2 keyword in the command structure; additionally, the IKEv1 transform-set is replaced by IKEv2 ipsec-proposal and because IKEv2 allows for asymmetric authentication, under the tunnel-group both local and remote authentication methods must be configured. ASA supports only crypto-map-based VPN with both IKEv1 and IKEv2.

> On IOS, if the IKEv2 profile is not attached to the crypto-map or the IPsec-profile, IKEv2 session cannot be initiated, so the router can act only as a responder.

Configuration

https://t.me/learningnets

R2:

```
crypto ikev2 proposal R2_ASA4_PROPOSAL
 encryption aes-cbc-256 3des
 integrity sha256 md5
 group 5 2
!
crypto ikev2 limit max-sa 2
crypto logging ikev2
!
!
crypto ikev2 keyring R2_ASA4_KEYRING
 peer ASA4
  address 136.1.28.8
  pre-shared-key local r2psk
  pre-shared-key remote asa4psk
!
!
crypto ikev2 profile R2_ASA4_PROFILE
 match identity remote address 136.1.28.8
 authentication local pre-share
 authentication remote pre-share
 keyring local R2_ASA4_KEYRING
!
!
ip access-list extended R2_ASA4_ACL
 permit ip host 150.1.2.2 host 150.1.1.1
 permit ip host 150.1.2.2 host 150.1.3.3
!
!
crypto ipsec transform-set AES_SHA esp-aes 256 esp-sha-hmac
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto map VPN 100 ipsec-isakmp
 set peer 136.1.28.8
 match address R2_ASA4_ACL
 set transform-set AES_SHA 3DES_MD5
 set ikev2-profile  R2_ASA4_PROFILE
!
interface GigabitEthernet0/0.28
 crypto map VPN
```

ASA4:

```
crypto ikev2 enable VLAN28
```

https://t.me/learningnets

```
crypto ikev2 policy 100
 encryption aes-256 3des
 integrity sha256 md5
 group 5 2
 prf sha256 md5
!
crypto ikev2 limit max-sa 2
!
!
access-list R2_ASA4_ACL permit ip host 150.1.1.1 host 150.1.2.2
access-list R2_ASA4_ACL permit ip host 150.1.3.3 host 150.1.2.2
!
!
crypto ipsec ikev2 ipsec-proposal R2_ASA4_PROPOSAL
 protocol esp encryption aes-256 3des
 protocol esp integrity sha-1 md5
!
!
tunnel-group 136.1.28.2 type ipsec-l2l
tunnel-group 136.1.28.2 ipsec-attributes
 ikev2 local-authentication pre-shared-key asa4psk
 ikev2 remote-authentication pre-shared-key r2psk
!
!
crypto map VPN 100 set peer 136.1.28.2
crypto map VPN 100 match address R2_ASA4_ACL
crypto map VPN 100 set ikev2 ipsec-proposal R2_ASA4_PROPOSAL
crypto map VPN interface VLAN28
```

Verification

Initiate traffic to trigger IKEv2 tunnel negotiation:

https://t.me/learningnets

```
R2#ping 150.1.1.1 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 150.1.2.2

%IKEV2-5-OSAL_INITIATE_TUNNEL: Received request to establish an IPsec tunnel; local traffic selector = Address Range
%IKEV2-5-SA_UP: SA UP
.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Verify the negotiated IKEv2 attributes, and note the default IKEv2 ID of IP address used by both ASA and IOS and that a single child-SA was negotiated:

```
R2#show crypto ikev2 sa
 IPv4 Crypto IKEv2  SA

Tunnel-id Local                 Remote                fvrf/ivrf            Status
1         136.1.28.2/500        136.1.28.8/500        none/none            READY
      Encr: AES-CBC, keysize: 256, Hash: SHA256, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/588 sec

 IPv6 Crypto IKEv2  SA
!
!R2#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session
Session-id:8, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local                 Remote                fvrf/ivrf            Status 1
136.1.28.2/500        136.1.28.8/500        none/none            READY
   Encr: AES-CBC, keysize: 256, Hash: SHA256, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/549 sec
      CE id: 1008, Session-id: 8
      Status Description: Negotiation done
      Local spi: A1875CECD9DB32D7       Remote spi: AC16CAFCADA59464 Local id: 136.1.28.2
Remote id: 136.1.28.8
      Local req msg id:  2              Remote req msg id:  49
      Local next msg id: 2              Remote next msg id: 49
      Local req queued:  2              Remote req queued:  49
      Local window:      5              Remote window:      1
      DPD configured for 0 seconds, retry 0
      NAT-T is not detected
```

https://t.me/learningnets

```
      Cisco Trust Security SGT is disabled
      Initiator of SA : YesChild sa: local selector  150.1.2.2/0 - 150.1.2.2/65535
remote selector 150.1.1.1/0 - 150.1.1.1/65535
          ESP spi in/out: 0x70442BFA/0x231204AA
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0   Encr: AES-CBC, keysize: 256, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel

 IPv6 Crypto IKEv2 Session
!
!ASA4# show crypto ikev2 sa detail

IKEv2 SAs:

Session-id:15, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id                 Local                Remote     Status         Role
32487981         136.1.28.8/500         136.1.28.2/500      READY    RESPONDER
 Encr: AES-CBC, keysize: 256, Hash: SHA256, DH Grp:5, Auth sign: PSK, Auth verify: PSK

      Life/Active Time: 86400/223 sec
      Session-id: 15
      Status Description: Negotiation done
      Local spi: AC16CAFCADA59464       Remote spi: A1875CECD9DB32D7 Local id: 136.1.28.8
 Remote id: 136.1.28.2
      Local req mess id: 16             Remote req mess id: 2
      Local next mess id: 16            Remote next mess id: 2
      Local req queued: 16              Remote req queued: 2
      Local window: 1                   Remote window: 5
      DPD configured for 10 seconds, retry 2
      NAT-T is not detected  Child sa: local selector  150.1.1.1/0 - 150.1.1.1/65535
 remote selector 150.1.2.2/0 - 150.1.2.2/65535
          ESP spi in/out: 0x231204aa/0x70442bfa
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0   Encr: AES-CBC, keysize: 256, esp_hmac: SHA96

          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel
```

Generate traffic to negotiate the second child-SA and verify it (ASA shows child-SA as number of IPsec tunnels), and confirm that packets are encrypted/decrypted:

```
R2#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
```

https://t.me/learningnets

```
Packet sent with a source address of 150.1.2.2

%IKEV2-5-OSAL_INITIATE_TUNNEL: Received request to establish an IPsec tunnel; local traffic selector = Address Range
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
!
!ASA4# show vpn-sessiondb detail l2l

Session Type: LAN-to-LAN Detailed

Connection   : 136.1.28.2
Index        : 35                    IP Addr      : 136.1.28.2
Protocol     : IKEv2 IPsec
Encryption   : AES256                Hashing      : SHA256 SHA1
Bytes Tx     : 800                   Bytes Rx     : 800
Login Time   : 15:02:12 UTC Tue Mar 4 2014
Duration     : 0h:16m:27s IKEv2 Tunnels: 1
IPsec Tunnels: 2

IKEv2:
  Tunnel ID    : 35.1
  UDP Src Port : 500                   UDP Dst Port : 500 Rem Auth Mode: preSharedKeys
Loc Auth Mode: preSharedKeys
  Encryption   : AES256                Hashing      : SHA256
  Rekey Int (T): 86400 Seconds         Rekey Left(T): 85413 Seconds
  PRF          : SHA256                D/H Group    : 5
  Filter Name  :
  IPv6 Filter  :
IPsec:
  Tunnel ID    : 35.2 Local Addr   : 150.1.1.1/255.255.255.255/0/0
Remote Addr  : 150.1.2.2/255.255.255.255/0/0
Encryption   : AES256                Hashing      : SHA1

  Encapsulation: Tunnel
  Rekey Int (T): 28800 Seconds         Rekey Left(T): 27813 Seconds
  Rekey Int (D): 4608000 K-Bytes       Rekey Left(D): 4608000 K-Bytes
  Idle Time Out: 30 Minutes            Idle TO Left : 13 Minutes
  Bytes Tx     : 400                   Bytes Rx     : 400
Pkts Tx      : 4                     Pkts Rx      : 4

IPsec:
  Tunnel ID    : 35.3 Local Addr   : 150.1.3.3/255.255.255.255/0/0
Remote Addr  : 150.1.2.2/255.255.255.255/0/0
Encryption   : AES256                Hashing      : SHA1

  Encapsulation: Tunnel
```

https://t.me/learningnets

```
  Rekey Int (T): 28800 Seconds          Rekey Left(T): 28784 Seconds
  Rekey Int (D): 4608000 K-Bytes        Rekey Left(D): 4608000 K-Bytes
  Idle Time Out: 30 Minutes             Idle TO Left : 29 Minutes
  Bytes Tx     : 400                    Bytes Rx     : 400
Pkts Tx       : 4                      Pkts Rx      : 4


NAC:
  Reval Int (T): 0 Seconds              Reval Left(T): 0 Seconds
  SQ Int (T)   : 0 Seconds              EoU Age(T)   : 987 Seconds
  Hold Left (T): 0 Seconds              Posture Token:
  Redirect URL :
!
!R2#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session
Session-id:8, Status:UP-ACTIVE, IKE count:1, CHILD count:2

Tunnel-id Local                 Remote                fvrf/ivrf            Status
1         136.1.28.2/500        136.1.28.8/500        none/none            READY
      Encr: AES-CBC, keysize: 256, Hash: SHA256, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/1057 sec
      CE id: 1008, Session-id: 8
      Status Description: Negotiation done
      Local spi: A1875CECD9DB32D7       Remote spi: AC16CAFCADA59464
      Local id: 136.1.28.2
      Remote id: 136.1.28.8
      Local req msg id:  3              Remote req msg id:  100
      Local next msg id: 3              Remote next msg id: 100
      Local req queued:  3              Remote req queued:  100
      Local window:      5              Remote window:      1
      DPD configured for 0 seconds, retry 0
      NAT-T is not detected
      Cisco Trust Security SGT is disabled
      Initiator of SA : YesChild sa: local selector  150.1.2.2/0 - 150.1.2.2/65535
remote selector 150.1.3.3/0 - 150.1.3.3/65535
          ESP spi in/out: 0xF8696BF0/0x9FD83EB0
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0   Encr: AES-CBC, keysize: 256, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel
Child sa: local selector  150.1.2.2/0 - 150.1.2.2/65535
remote selector 150.1.1.1/0 - 150.1.1.1/65535
          ESP spi in/out: 0x70442BFA/0x231204AA
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0   Encr: AES-CBC, keysize: 256, esp_hmac: SHA96
```

https://t.me/learningnets

```
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel


 IPv6 Crypto IKEv2 Session
```

Alternatively, the legacy **`show crypto ipsec sa`** commands can be used to visualize only IPsec tunnel-negotiated parameters (child-SA):

```
ASA4# show crypto ipsec sa

interface: VLAN28
    Crypto map tag: VPN, seq num: 100, local addr: 136.1.28.8
 access-list R2_ASA4_ACL extended permit ip host 150.1.1.1 host 150.1.2.2
 local ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
 remote ident (addr/mask/prot/port): (150.1.2.2/255.255.255.255/0/0)
      current_peer: 136.1.28.2
 #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
 #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 4, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 136.1.28.8/500, remote crypto endpt.: 136.1.28.2/500
      path mtu 1500, ipsec overhead 74, media mtu 1500
      current outbound spi: 70442BFA
      current inbound spi : 231204AA
 inbound esp sas:
      spi: 0x231204AA (588383402) transform: esp-aes-256 esp-sha-hmac no compression

         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 143360, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (3962879/27532)
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x0000001F outbound esp sas:
      spi: 0x70442BFA (1883515898) transform: esp-aes-256 esp-sha-hmac no compression
         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 143360, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4285439/27532)
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x00000001
```

https://t.me/learningnets

```
    Crypto map tag: VPN, seq num: 100, local addr: 136.1.28.8
 access-list R2_ASA4_ACL extended permit ip host 150.1.3.3 host 150.1.2.2
  local ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
remote ident (addr/mask/prot/port): (150.1.2.2/255.255.255.255/0/0)
      current_peer: 136.1.28.2
 #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 4, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0

      local crypto endpt.: 136.1.28.8/500, remote crypto endpt.: 136.1.28.2/500
      path mtu 1500, ipsec overhead 74, media mtu 1500
      current outbound spi: F8696BF0
      current inbound spi : 9FD83EB0
 inbound esp sas:
      spi: 0x9FD83EB0 (2681749168) transform: esp-aes-256 esp-sha-hmac no compression
         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 143360, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (3962879/28501)
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x0000001F outbound esp sas:
      spi: 0xF8696BF0 (4167658480) transform: esp-aes-256 esp-sha-hmac no compression

         in use settings ={L2L, Tunnel, }
         slot: 0, conn_id: 143360, crypto-map: VPN
         sa timing: remaining key lifetime (kB/sec): (4285439/28498)
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x00000001
!
!R2#show crypto ipsec sa


interface: GigabitEthernet0/0.28
    Crypto map tag: VPN, local addr 136.1.28.2

   protected vrf: (none) local  ident (addr/mask/prot/port): (150.1.2.2/255.255.255.255/0/0)
remote ident (addr/mask/prot/port): (150.1.3.3/255.255.255.255/0/0)
```

https://t.me/learningnets

```
   current_peer 136.1.28.8 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.28.2, remote crypto endpt.: 136.1.28.8
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.28
     current outbound spi: 0x9FD83EB0(2681749168)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0xF8696BF0(4167658480) transform: esp-256-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2022, flow_id: Onboard VPN:22, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4170300/3230)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0x9FD83EB0(2681749168) transform: esp-256-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2021, flow_id: Onboard VPN:21, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4170300/3230)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     outbound ah sas:

     outbound pcp sas:

   protected vrf: (none) local  ident (addr/mask/prot/port): (150.1.2.2/255.255.255.255/0/0)
remote ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   current_peer 136.1.28.8 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
#pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.28.2, remote crypto endpt.: 136.1.28.8
```

https://t.me/learningnets

```
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.28
     current outbound spi: 0x231204AA(588383402)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0x70442BFA(1883515898) transform: esp-256-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2020, flow_id: Onboard VPN:20, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4283288/2259)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0x231204AA(588383402) transform: esp-256-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2019, flow_id: Onboard VPN:19, sibling_flags 80000040, crypto map: VPN
        sa timing: remaining key lifetime (k/sec): (4283288/2259)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv2 AnyConnect on ASA with EAP

Task

- Configure ASA4 for IKEv2 remote-access VPN on its VLAN38 interface:
  - Negotiate AES-256 and 3DES for encryption, SHA1 and MD5 for integrity, DH groups 5 and 2.
  - Assign IP addresses from 136.1.11.0/24 subnet and ensure that only traffic for VLAN28 is sent through the tunnel.
- Test functionality from **TEST-PC-A** with the username/password pair **ANYCONNECT/ikev2password**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

When using AnyConnect with IKEv2, the following should be mentioned:

- By default, AnyConnect uses SSL for connecting to the VPN gateway, so to use IKEv2 it needs a client profile configured for IKEv2; the AnyConnect client profile can be automatically provisioned by the ASA after first SSL successful connection, or manually provisioned by copying the XML file into the correct path, depending on the operating system of the client.
- The IKEv2 control-plane session (IKE_SA_INIT) is initiated over UDP 500, but after

https://t.me/learningnets

this is complete, EAP authentication is performed over UDP 4500.

- The IKEv2 data-plane session always negotiates NAT-T, even though it is not behind a NAT device; so after it is connected, client traffic through the IKEv2 tunnel is always encapsulated into UDP 4500.
- Although EAP was selected so that standard authentication is performed with IKEv2 remote access clients (EAP-MD5, EAP-GTC, EAP-MSCHAPv2, EAP-TLS, EAP-PEAP), ASA supports only a proprietary EAP-AnyConnect method that requires username/password based-authentication for the client; IOS also supports IKE-RSA authentication instead of EAP, so the client is authenticated using only certificates.
- Disabling the EAP-Anyconnect method removes the ability of the ASA to use features such as session timeout, idle timeout, disconnected timeout, split tunneling, split DNS, and MSIE proxy configuration.
- IKEv2 remote access on ASA uses a asymmetric authentication, where the ASA is authenticated through certificates and the client through username/password; optionally, the client can use double authentication, both certificates and username/password.
- IKEv2 IKE-ID is configurable only when the VPN gateway is an IOS router; for ASA firewall, it uses a nonconfigurable value of ***$AnyConnectClient$***.

Because of the special IKE-ID value of ***$AnyConnectClient$***, ASA has the following options for binding an incoming IKEv2 remote access session to a tunnel-group:

- Double authentication is implemented (EAP-Anyconnect and digital certificate) for the client; certificate-maps are configured on the ASA and functionality is enabled with the command `tunnel-group-map enable rules .`
- Group-alias is configured under the tunnel-group, and alias listing is enabled under webvpn with the command `tunnel-group-list enable` so that the client can select the tunnel-group when prompted for authentication.
- AnyConnect client profile is configured to connect to a specific tunnel-group.

Because the Anyconnect XML profile for IKEv2 connection is not already provisioned on the client, the below configuration allows for both IKEv2 and SSL connections; the first connection will be SSL to download the Anyconnect XML profile, and the second connection will be IKEv2. From a configuration perspective, the following steps are required:

1. Enroll the ASA with a PKI to request an identity certificate that will be used to authenticate the VPN gateway to the Anyconnect client.

https://t.me/learningnets

2. Activate the certificate to be used for IKEv2 remote access authentication; depending on the running AnyConnect version, a bug may require you to also configure the certificate for SSL session.
3. Enable IKEv2, and configure IKEv2 policy and IPsec proposal.
4. Define the pool to be used for client IP address allocation, and optionally define the split-tunneling policy.
5. Configure the tunnel-group and group-policy.
6. Define the dynamic crypto-map, enable RRI, bind it to a static crypto-map, and attach the static crypto-map to the interface used to terminate IKEv2 sessions.
7. If the Anyconnect XML profile is not provisioned on the client, enable and configure webvpn.
8. Redistribute the RRI routes in the internal IGP.

IKEv2 does not support pre-shared-key authentication for remote-access VPNs, only for LAN-to-LAN. When IKEv2 is used for remote-access VPNs, the VPN gateway must always authenticate with digital certificates, whereas the client has several options: username/password, certificates, or both but using an EAP method. So in IKEv2, there is no X-AUTH as there is in IKEv1; it is being replaced by an EAP type of authentication. EAP authentication takes place in the IKE_AUTH phase as additional messages and is only used to authenticate the initiator (client) to the responder (VPN gateway); with this extra step, the number of messages in the IKE_AUTH phase is increased from just 2 to somewhere between 12 and 16, depending on the EAP method used. The MODE_CONFIG from IKEv1, which was known as Phase 1.5, is being replaced by CONFIGURATION_PAYLOAD and it is based on IKEv2 informational messages, taking place within IKE_AUTH phase. With CONFIGURATION_PAYLOAD, the client requests configuration attributes from the VPN gateway (such as IPv4/IPv6 addressing information) using CFG_REQUEST messages and receives response from the VPN gateway through CFG_REPLY messages. The real distinction from IKEv1 MODE_CONFIG is that in IKEv2, the client can also push configuration attributes to the VPN gateway (like IPV4 subnets locally attached) through CFG_SET messages.

Configuration

ASA4:

```
crypto key generate rsa general-keys label IKEV2 modulus 1024 noconfirm
crypto ca trustpoint SELF
 enrollment self
 subject-name CN=ASA4.inelab.local
 keypair IKEV2
```

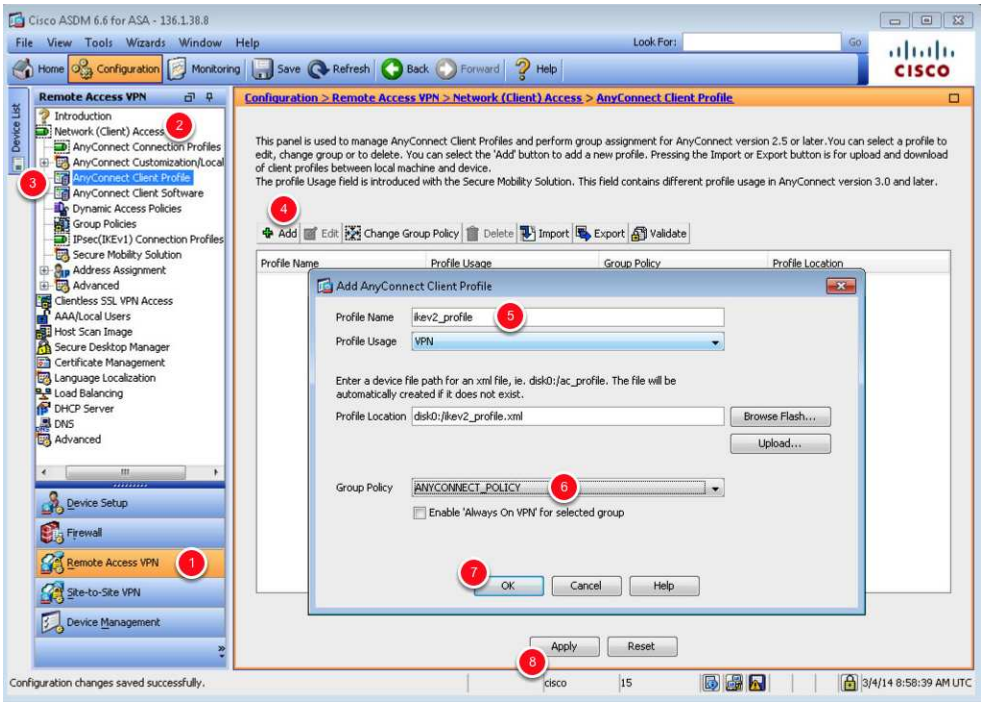
https://t.me/learningnets

```
!
crypto ca enroll SELF noconfirm
crypto ikev2 remote-access trustpoint SELF
ssl trust-point SELF VLAN38
!
!
crypto ikev2 enable VLAN38
crypto ikev2 policy 100
 encryption aes-256 3des
 integrity sha md5
 group 5 2
!
!
access-list ANYCONNECT_SPLIT_ACL permit 136.1.28.0 255.255.255.0
ip local pool ANYCONNECT_POOL 136.1.11.1-136.1.11.254 mask 255.255.255.0
username ANYCONNECT password ikev2password
!
!
crypto ipsec ikev2 ipsec-proposal ANYCONNECT_PROPOSAL
 protocol esp encryption aes-256 3des
 protocol esp integrity sha-1 md5
!
!
group-policy ANYCONNECT_POLICY internal
group-policy ANYCONNECT_POLICY attributes
 vpn-tunnel-protocol ikev2 ssl-client
 split-tunnel-policy tunnelspecified
 split-tunnel-network-list value ANYCONNECT_SPLIT_ACL
 webvpn
  anyconnect profiles value ikev2_profile type user
!
!
tunnel-group ANYCONNECT_TUNNEL type remote-access
tunnel-group ANYCONNECT_TUNNEL general-attributes
 address-pool ANYCONNECT_POOL
 default-group-policy ANYCONNECT_POLICY
tunnel-group ANYCONNECT_TUNNEL webvpn-attributes
 group-alias ANYCONNECT_IKEV2
 authentication aaa
!
!
crypto dynamic-map ANYCONNECT_MAP 100 set ikev2 ipsec-proposal ANYCONNECT_PROPOSAL
crypto dynamic-map ANYCONNECT_MAP 100 set reverse-route
crypto map VPN 1000 ipsec-isakmp dynamic ANYCONNECT_MAP
crypto map VPN interface VLAN38
!
```

https://t.me/learningnets

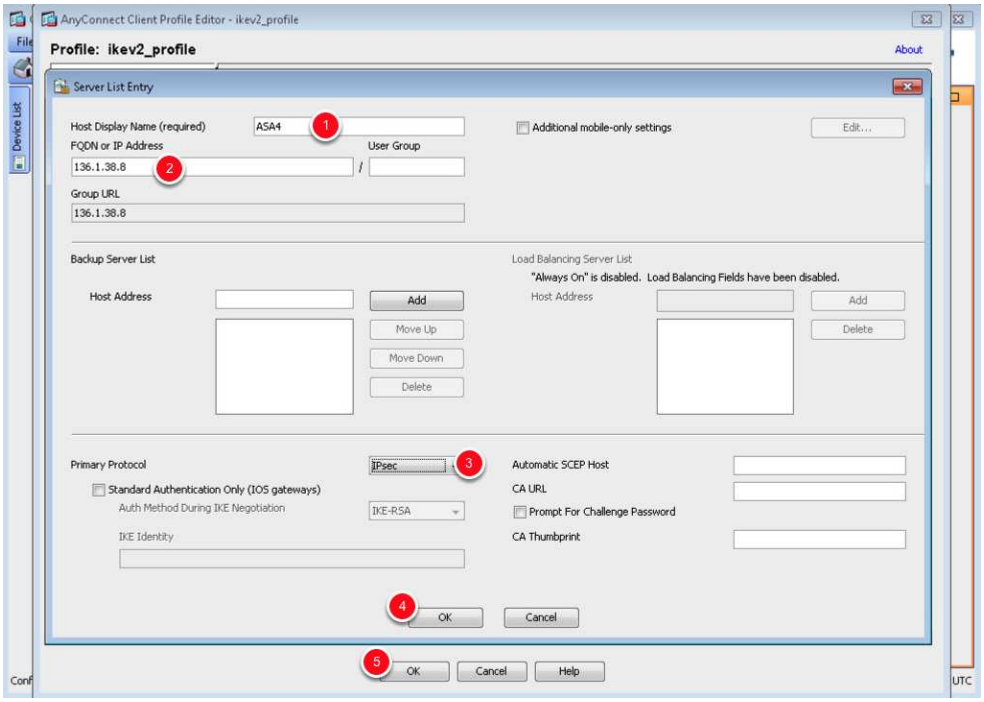
```
!
webvpn
 enable VLAN38
 anyconnect image disk0:/anyconnect-win-3.0.11042-k9.pkg
 anyconnect profiles ikev2_profile disk0:/ikev2_profile.xml
 anyconnect enable
 tunnel-group-list enable
!
!
router eigrp 111
 redistribute static
```

Using ASDM, connect to the ASA and create an AnyConnect profile, and assign it to the **ANYCONNECT_POLICY** group-policy. This requires enabling the HTTP server and allowing access from VLAN 3.



Edit the client profile, navigate to the **Server List** option, and click **Add**; give the profile a name such as **ASA4**, which will appear in the AnyConnect window, and specify the IP address or FQDN of the ASA (because DNS is not implemented, we will be using the IP address) and specify IPsec (which means IKEv2) as the primary protocol:

https://t.me/learningnets



Verification

Verify the ASA self-signed certificate:

```
ASA4# show crypto ca certificates
 Certificate
Status: Available

  Certificate Serial Number: 0e911453
  Certificate Usage: General Purpose
  Public Key Type: RSA (1024 bits)
  Signature Algorithm: SHA1 with RSA Encryption
  Issuer Name:
    hostname=ASA4.ine.com
    cn=ASA4.inelab.local
  Subject Name:
    hostname=ASA4.ine.com
    cn=ASA4.inelab.local
  Validity Date:
    start date: 17:31:36 UTC Mar 3 2014
    end   date: 17:31:36 UTC Feb 29 2024
  Associated Trustpoints: SELF
```

Initiate the AnyConnect session from **TEST-PC-A**; it will connect using SSL and download the client profile, which will force IKEv2 as the preferred protocol:

https://t.me/learningnets

```
ASA4#show vpn-sessiondb detail anyconnect


Session Type: AnyConnect Detailed

Username     : ANYCONNECT             Index        : 25
Assigned IP  : 136.1.11.1             Public IP    : 136.1.3.51
Protocol     : AnyConnect-Parent SSL-Tunnel DTLS-Tunnel
License      : AnyConnect Premium
Encryption   : RC4 AES128             Hashing      : none SHA1
Bytes Tx     : 12109                  Bytes Rx     : 1116
Pkts Tx      : 11                     Pkts Rx      : 6
Pkts Tx Drop : 0                      Pkts Rx Drop : 0
Group Policy : ANYCONNECT_POLICY      Tunnel Group : ANYCONNECT_TUNNEL
Login Time   : 09:13:33 UTC Tue Mar 4 2014
Duration     : 0h:02m:24s
Inactivity   : 0h:00m:00s
NAC Result   : Unknown
VLAN Mapping : N/A                    VLAN         : none
AnyConnect-Parent Tunnels: 1
SSL-Tunnel Tunnels: 1
DTLS-Tunnel Tunnels: 1
AnyConnect-Parent:
  Tunnel ID    : 25.1
  Public IP    : 136.1.3.51
  Encryption   : none                   TCP Src Port : 1128
  TCP Dst Port : 443                    Auth Mode    : userPassword
  Idle Time Out: 30 Minutes             Idle TO Left : 28 Minutes
  Client Type  : AnyConnect
  Client Ver   : 3.0.11042
  Bytes Tx     : 7230                   Bytes Rx     : 871
  Pkts Tx      : 7                      Pkts Rx      : 2
  Pkts Tx Drop : 0                      Pkts Rx Drop : 0
SSL-Tunnel:
  Tunnel ID    : 25.2 Assigned IP  : 136.1.11.1             Public IP    : 136.1.3.51
Encryption   : RC4                    Hashing      : SHA1
                  Encapsulation: TLSv1.0                TCP Src Port : 1134

  TCP Dst Port : 443                    Auth Mode    : userPassword
  Idle Time Out: 30 Minutes             Idle TO Left : 28 Minutes
  Client Type  : SSL VPN Client
  Client Ver   : Cisco AnyConnect VPN Agent for Windows 3.0.11042
  Bytes Tx     : 4879                   Bytes Rx     : 245
  Pkts Tx      : 4                      Pkts Rx      : 4
  Pkts Tx Drop : 0                      Pkts Rx Drop : 0
DTLS-Tunnel:
```

https://t.me/learningnets

```
  Tunnel ID    : 25.3 Assigned IP  : 136.1.11.1            Public IP    : 136.1.3.51
Encryption   : AES128                Hashing      : SHA1
                  Encapsulation: DTLSv1.0              UDP Src Port : 56306

  UDP Dst Port : 443                   Auth Mode    : userPassword
  Idle Time Out: 30 Minutes            Idle TO Left : 28 Minutes
  Client Type  : DTLS VPN Client
  Client Ver   : 3.0.11042
  Bytes Tx     : 0                     Bytes Rx     : 0
  Pkts Tx      : 0                     Pkts Rx      : 0
  Pkts Tx Drop : 0                     Pkts Rx Drop : 0

NAC:
  Reval Int (T): 0 Seconds             Reval Left(T): 0 Seconds
  SQ Int (T)   : 0 Seconds             EoU Age(T)   : 148 Seconds
  Hold Left (T): 0 Seconds             Posture Token:
  Redirect URL :
```

Disconnect the session and connect again using the **ASA4** entry, which represents the downloaded client profile. Verify the IKEv2 session, and notice the non-configurable IKEv2 IKE-ID used by AnyConnect and that NAT-T was enabled although there is no NAT device between the client and the VPN gateway:

```
ASA4# show crypto ikev2 sa detail


IKEv2 SAs:

Session-id:8, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id                Local                Remote     Status         Role
14706621       136.1.38.8/4500      136.1.3.51/63090      READY    RESPONDER
 Encr: AES-CBC, keysize: 256, Hash: SHA96, DH Grp:5, Auth sign: RSA, Auth verify: EAP
      Life/Active Time: 86400/682 sec
      Session-id: 8
      Status Description: Negotiation done
 Local spi: DA5BC919D4C9CC20       Remote spi: 8A295C4CD35D5F81
 Local id: hostname=ASA4.ine.com,cn=ASA4.inelab.local
 Remote id: *$AnyConnectClient$*
      Local req mess id: 21            Remote req mess id: 16
      Local next mess id: 21           Remote next mess id: 16
      Local req queued: 21             Remote req queued: 16
      Local window: 1                  Remote window: 1
      DPD configured for 10 seconds, retry 2 NAT-T is detected  outside
```

https://t.me/learningnets

```
Assigned host addr: 136.1.11.1

Child sa: local selector  0.0.0.0/0 - 255.255.255.255/65535
          remote selector 136.1.11.1/0 - 136.1.11.1/65535
          ESP spi in/out: 0x5e55752c/0x247fda58
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0
          Encr: AES-CBC, keysize: 256, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel
```

Verify that ASA injected a static route for the Anyconnect client assigned IP address, which was redistributed into EIGRP:

```
ASA4# show route VLAN38



Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set
S     136.1.11.1 255.255.255.255 [1/0] via 136.1.38.3, VLAN38
D    136.1.3.0 255.255.255.0 [90/28416] via 136.1.38.3, 7:04:23, VLAN38
C    136.1.38.0 255.255.255.0 is directly connected, VLAN38
D    150.1.3.3 255.255.255.255 [90/130816] via 136.1.38.3, 7:04:23, VLAN38
!
!R2#show ip route 136.1.11.1
Routing entry for 136.1.11.1/32 Known via "eigrp 111", distance 170, metric 28416, type external
  Redistributing via eigrp 111
  Last update from 136.1.28.8 on GigabitEthernet0/0.28, 00:14:56 ago
  Routing Descriptor Blocks: * 136.1.28.8, from 136.1.28.8, 00:14:56 ago, via GigabitEthernet0/0.28

      Route metric is 28416, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Verify connectivity between R2 and the AnyConnect client, and check the encrypted/decrypted packet counters:

https://t.me/learningnets

```
R2#ping 136.1.11.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.11.1, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/8 ms
!
!ASA4# show crypto ipsec sa

interface: VLAN38
    Crypto map tag: ANYCONNECT_MAP, seq num: 100, local addr: 136.1.38.8

      local ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
      remote ident (addr/mask/prot/port): (136.1.11.1/255.255.255.255/0/0)
current_peer: 136.1.3.51, username: ANYCONNECT
 dynamic allocated peer ip: 136.1.11.1
 #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
 #pkts decaps: 7, #pkts decrypt: 7, #pkts verify: 7
      #pkts compressed: 0, #pkts decompressed: 0
      #pkts not compressed: 5, #pkts comp failed: 0, #pkts decomp failed: 0
      #pre-frag successes: 0, #pre-frag failures: 0, #fragments created: 0
      #PMTUs sent: 0, #PMTUs rcvd: 0, #decapsulated frgs needing reassembly: 0
      #send errors: 0, #recv errors: 0
 local crypto endpt.: 136.1.38.8/4500, remote crypto endpt.: 136.1.3.51/63090
      path mtu 1300, ipsec overhead 82, media mtu 1500
      current outbound spi: 247FDA58
      current inbound spi : 5E55752C
 inbound esp sas:
      spi: 0x5E55752C (1582658860)
         transform: esp-aes-256 esp-sha-hmac no compression
in use settings ={RA, Tunnel,  NAT-T-Encaps, }
         slot: 0, conn_id: 90112, crypto-map: ANYCONNECT_MAP
         sa timing: remaining key lifetime (sec): 27856
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
          0x00000000 0x000000FF outbound esp sas:
      spi: 0x247FDA58 (612358744)
         transform: esp-aes-256 esp-sha-hmac no compression
in use settings ={RA, Tunnel,  NAT-T-Encaps, }

         slot: 0, conn_id: 90112, crypto-map: ANYCONNECT_MAP
         sa timing: remaining key lifetime (sec): 27856
         IV size: 16 bytes
         replay detection support: Y
         Anti replay bitmap:
```

https://t.me/learningnets

```
          0x00000000 0x00000001
```

Verify the AnyConnect session details, and notice the non-configurable local authentication of RSA-certificate and remote authentication of username/password, as well as NAT-T being used for data-plane traffic:

```
ASA4# show vpn-sessiondb detail anyconnect


Session Type: AnyConnect Detailed
Username     : ANYCONNECT              Index        : 22
Assigned IP  : 136.1.11.1              Public IP    : 136.1.3.51
Protocol     : IKEv2 IPsecOverNatT AnyConnect-Parent
License      : AnyConnect Premium Encryption   : AES256                 Hashing      : none SHA1
Bytes Tx     : 500                     Bytes Rx     : 595
Pkts Tx      : 5                       Pkts Rx      : 7
Pkts Tx Drop : 0                       Pkts Rx Drop : 0
Group Policy : ANYCONNECT_POLICY       Tunnel Group : ANYCONNECT_TUNNEL
Login Time   : 21:13:14 UTC Mon Mar 3 2014
Duration     : 0h:16m:01s
Inactivity   : 0h:00m:00s
NAC Result   : Unknown
VLAN Mapping : N/A                     VLAN         : none
IKEv2 Tunnels: 1
IPsecOverNatT Tunnels: 1
AnyConnect-Parent Tunnels: 1
AnyConnect-Parent:
  Tunnel ID    : 22.1
  Public IP    : 136.1.3.51
  Encryption   : none                    Auth Mode    : userPassword
  Idle Time Out: 30 Minutes              Idle TO Left : 13 Minutes
  Client Type  : AnyConnect
  Client Ver   : 3.0.11042
IKEv2:
  Tunnel ID    : 22.2 UDP Src Port : 63090                  UDP Dst Port : 4500
Rem Auth Mode: userPassword
Loc Auth Mode: rsaCertificate
Encryption   : AES256                  Hashing      : SHA1
  Rekey Int (T): 86400 Seconds           Rekey Left(T): 85438 Seconds
  PRF          : SHA1                    D/H Group    : 5
  Filter Name  :
  Client OS    : Windows
IPsecOverNatT:
  Tunnel ID    : 22.3
  Local Addr   : 0.0.0.0/0.0.0.0/0/0
```

https://t.me/learningnets

```
  Remote Addr  : 136.1.11.1/255.255.255.255/0/0
Encryption    : AES256                  Hashing      : SHA1

  Encapsulation: Tunnel
  Rekey Int (T): 28800 Seconds          Rekey Left(T): 27838 Seconds
  Rekey Int (D): 4608000 K-Bytes        Rekey Left(D): 4608000 K-Bytes
  Idle Time Out: 30 Minutes             Idle TO Left : 29 Minutes
  Bytes Tx     : 500                    Bytes Rx     : 595
  Pkts Tx      : 5                      Pkts Rx      : 7

NAC:
  Reval Int (T): 0 Seconds              Reval Left(T): 0 Seconds
  SQ Int (T)   : 0 Seconds              EoU Age(T)   : 964 Seconds
  Hold Left (T): 0 Seconds              Posture Token:
  Redirect URL :
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv2 L2L Between IOS Routers with SVTI and PSK

Task

- Configure IKEv2 IPsec between R1 and R2 as follows:
  - Use the default IKEv2/IPsec configurations.
  - Enable PSK authentication with the string **ikev2psk**.
  - Enable EIGRP in AS 12 over the IPsec tunnel to protect traffic between VLAN 11 and VLAN 22.
  - On the tunnel, use IPv4 addressing in the format 136.1.12.Y/24, where **Y** is the router number.
- Make sure to match the following output on both routers:

```
R1#show crypto ipsec sa | i ident

    local  ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
    remote ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
```

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

In IPsec VPNs, IKE or IKEv1 (defined in RFC2409) were used to establish a secure

https://t.me/learningnets

communication tunnel/channel between two VPN endpoints (derive keying material), so that IPsec negotiation was protected and finally user traffic was protected by the IPsec tunnel. Looking at the chain of events, moving from IKEv1 to IKEv2 only affects how building the initial protected channel takes place; IPsec by itself does not change and the final scope to protect user traffic remains the same. Because IKEv2 is not backward compatible with IKEv1, this is the first sign that from a design/architectural level IKEv2 protocol is very different from IKEv1. This does not mean that you cannot run both at the same time on a Cisco device; if the IOS code and licensing support it, you can enable and use both at the same time, but between any two VPN endpoints there will be no IKE version negotiation supported. Both must be configured for IKEv1 or IKEv2. The main reason for this incompatibility is that IKEv2 is not a new IKE version bringing enhancements to IKEv1 or a patch to the current IKEv1; it's a completely redesigned protocol but it keeps the scope of IKEv1.

IKEv2 was initially defined in RFC4306 but updated in RFC5996. It still uses UDP 500 and UDP 4500, however:

- Features such as DPD, NAT-T, MODE-CONFIG, and relevant specifications from ISAKMP RFC2408, IKE RFC2409, and DOI RFC2407 are now standardized and specified within common RFC5996.
- IKEv1 supports PSK, RSA-SIG, and RSA-ENCR as authentication methods; IKEv2 supports PSK, RSA-SIG, and EAP.
- IKEv1 requires symmetric authentication (both peers using the same method), whereas IKEv2 allows for asymmetric authentication (for example, one side with RSA-SIG and one wide with PSK, or even different PSK).
- IKEv1 policies require authentication type to be specified, whereas IKEv2 policies are agnostic to authentication (IOS calls it proposals now), which allows for IKEv2 ID to be protected at all times.
- IKEv1 does not have a built-in Anti-DOS function, whereas IKEv2 does through anti-clogging cookies.
- IKEv1 requires re-authentication for IKE rekeying, whereas IKEv2 does not.
- IKEv1 uses the same algorithm for integrity and deriving keying material; IKEv2 uses PRF (Pseudo Random Function) for deriving keying material, which can be different than integrity algorithm (only ASA supports it).
- IKEv1 informational/notification messages are not acknowledged, whereas in IKEv2 they must be acknowledged.

IKEv2 negotiation has been simplified:

- Negotiation is shorter. In most cases it requires only four messages for both IKEv2

https://t.me/learningnets

and IPsec tunnel to be established.
- It uses only four message types: IKE_SA_INIT, IKE_AUTH, CREATE_CHILD_SA, and Informational.
- Phase 1 from IKEv1, which has two functional modes (Main and Aggressive), is known in IKEv2 as IKE_SA_INIT and has a single functional mode requiring two messages to be exchanged.
- Within a single policy (known as proposal on IOS and policy on ASA), multiple encryption/integrity/PRF/DH groups can be specified in an **OR** fashion.
- After IKE_SA_INIT derives the keying material, mutual authentication is performed through IKE_AUTH, which requires two messages to be exchanged (except for remote-access with EAP authentication).
- Phase 2 from IKEv1 (Quick Mode) is known in IKEv2 as CREATE_CHILD_SA; however, the first CHILD_SA is derived in IKE_AUTH, so this step is optional and used for additional SA creation or rekeying.

There is not an exact one-to-one mapping between IKEv1 Phase1 and IKEv2 IKE_SA_INIT, or between IKEv1 Phase2 and IKEv2 CREATE_CHILD_SA. A more appropriate mapping between IKEv1/IKEv2 messages would be:

- IKE_SA_INIT in IKEv2 is the equivalent of MM1,MM2,MM3, and MM4 in IKEv1.
- IKE_AUTH in IKEv2 is the equivalent of MM5, MM6, and part of QM in IKEv1.
- CREATE_CHILD_SA in IKEv2 is the equivalent of MM or QM in IKEv1.

IKEv2 supports most IPsec VPN types as IKEv1:

- Legacy crypto-map (static and dynamic)
- DVTI and SVTI
- GRE over IPsec
- DMVPN

At this moment, GET VPN still only supports IKEv1, and the EzVPN server/remote functionality from IKEv1 has been replaced by FlexVPN server/client in IKEv2. However, FlexVPN is a completely new framework for IKEv2 configuration on IOS. From the configuration perspective, we must learn some new terms on IOS:

- **IKEv2 proposal** defines the integrity/encryption/DH group settings and replaces the IKEv1 policies; IKEv2 proposals use names instead of priority numbers, and within each proposal multiple algorithms of each type can be defined, the order of configuration being important. For example, if **3DES** and **AES128** are configured in this order, **3DES** will be negotiated first with the remote peer.

https://t.me/learningnets

- **IKEv2 policy** did not exist in IKEv1, and its scope is that you can now restrict which IKEv2 proposals will be negotiated with each VPN peer based on several factors: the local VPN IP address, the FVRF on which the VPN is terminated, and the IP address of the remote VPN peer (not yet implemented). In IKEv1, with each new IPsec VPN, all IKEv1 globally configured policies were used in negotiation.

Additionally, in IKEv1 IPsec VPNs, ISAKMP/IKEv1 profiles were optional but recommended, but in IKEv2 IPsec VPNs, IKEv2 profiles are required; Cisco essentially enforced this to make sure that each IPsec VPN properly matches the intended configuration when multiple VPN types are configured on the same device. Under the IKEv2 profile, configuration of the following parameters is mandatory:

- Matching IKEv2 identity of the remote peer, which can be an IPv4/IPv6 address, DN (if certificates are used), FQDN, email, or key-id; only the email IKEv2 ID did not exist in IKEv1.
- Specifying the local and remote authentication types.
- Specifying the PKI trustpoint (which defines the certificate to be used or the PKI chain to validate remote peer certificate) if local and/or remote authentication type is RSA-SIG.
- Specifying the IKEv2 keyring (which defines the PSK for local and/or remote authentication) if local and/or remote authentication type is PSK; globally defined PSKs are no longer configurable.

Regardless of the authentication type, by default (unless modified in the IKEv2 profile), IOS will always use an IKEv2 ID of IPv4/IPv6 address. The intention is that IKEv2 profile must be attached to a crypto map or ipsec profile to be active; however, based on the IOS code that you are running, this may be required only on the initiator side. To keep the configuration consistent and to allow both VPN gateways to play the initiator role, make sure to apply the IKEv2 profile on both sides. Just like in IKEv1, there are also some default configurations with IKEv2, which Cisco calls *smart defaults*:

- One default IKEv2 proposal
- One default IKEv2 policy that uses the default IKEv2 proposal
- One default IPsec transform-set
- One default IPsec profile that uses the default IPsec transform-set

Configuration

R1:

```
crypto ikev2 keyring R1_R2_KEYRING
 peer R2
```

https://t.me/learningnets

```
  address 136.1.28.2
  pre-shared-key local ikev2psk
  pre-shared-key remote ikev2psk
!
!
crypto ikev2 profile R1_R2_PROFILE
 match identity remote address 136.1.28.2
 authentication local pre-share
 authentication remote pre-share
 keyring local R1_R2_KEYRING
!
!
crypto ipsec profile default
 set ikev2-profile R1_R2_PROFILE
!
!
interface Tunnel12
 ip address 136.1.12.1 255.255.255.0
 tunnel source 136.1.18.1
 tunnel destination 136.1.28.2
 tunnel protection ipsec profile default
 tunnel mode ipsec ipv4
!
!
router eigrp 12
 no auto-summary
 network 136.1.12.0 0.0.0.255
 network 136.1.11.0 0.0.0.255
```

R2:

```
crypto ikev2 keyring R1_R2_KEYRING
 peer R1
  address 136.1.18.1
  pre-shared-key local ikev2psk
  pre-shared-key remote ikev2psk
!
!
crypto ikev2 profile R1_R2_PROFILE
 match identity remote address 136.1.18.1
 authentication local pre-share
 authentication remote pre-share
 keyring local R1_R2_KEYRING
!
!
```

https://t.me/learningnets

```
crypto ipsec profile default
 set ikev2-profile R1_R2_PROFILE
!
!
interface Tunnel12
 ip address 136.1.12.2 255.255.255.0
 tunnel source 136.1.28.2
 tunnel destination 136.1.18.1
 tunnel protection ipsec profile default
 tunnel mode ipsec ipv4
!
!
router eigrp 12
 no auto-summary
 network 136.1.12.0 0.0.0.255
 network 136.1.22.0 0.0.0.255
```

Verification

Verify the smart defaults that are used by our configuration:

```
R1#show crypto ikev2 proposal default
 IKEv2 proposal: default
      Encryption : AES-CBC-256 AES-CBC-192 AES-CBC-128
      Integrity  : SHA512 SHA384 SHA256 SHA96 MD596
      PRF        : SHA512 SHA384 SHA256 SHA1 MD5
      DH Group   : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
!
!R1#show crypto ikev2 policy default
 IKEv2 policy : default
       Match fvrf : any
       Match address local : any  Proposal    : default
!
!R1#show crypto ipsec transform-set default

{ esp-aes esp-sha-hmac  }
   will negotiate = { Transport,  },
!
!R1#show crypto ipsec profile default
IPSEC profile default IKEV2 profile R1_R2_PROFILE
        Security association lifetime: 4608000 kilobytes/3600 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N Transform sets={
 default:  { esp-aes esp-sha-hmac  } ,
```

https://t.me/learningnets

```
        }
```

Verify the configured IKEv2-profile, which is applied under the default ipsec profile:

```
R1#show crypto ikev2 profile



IKEv2 profile: R1_R2_PROFILE
 Ref Count: 4
 Match criteria:
  Fvrf: global
  Local address/interface: none
  Identities:  address 136.1.28.2 255.255.255.255
  Certificate maps: none
 Local identity: none
 Remote identity: none Local authentication method: pre-share
 Remote authentication method(s): pre-share
 EAP options: none Keyring: R1_R2_KEYRING

 Trustpoint(s): none
 Lifetime: 86400 seconds
 DPD: disabled
 NAT-keepalive: disabled
 Ivrf: none
 Virtual-template: none
 AAA EAP authentication mlist: none
 AAA Accounting: none
 AAA group authorization: none
 AAA user authorization: none
```

Verify the IKEv2 tunnel; note that the first algorithms from the default IKEv2 proposal have been negotiated:

https://t.me/learningnets

```
R1#show crypto ikev2 sa

 IPv4 Crypto IKEv2  SA

Tunnel-id Local                 Remote                fvrf/ivrf            Status
3         136.1.18.1/500        136.1.28.2/500        none/none            READY

      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/55528 sec

 IPv6 Crypto IKEv2  SA
```

Verify that routing is functional, and check IPv4 connectivity:

https://t.me/learningnets

```
R1#show ip eigrp 12 neighbors

EIGRP-IPv4 Neighbors for AS(12)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   136.1.12.2              Tu12                     13 15:26:06    6  1470  0  27

!
!R1#show ip route eigrp 12

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 10 subnets, 2 masks
D        136.1.22.0/24 [90/26882560] via 136.1.12.2, 15:26:11, Tunnel12

!
!R1#ping 136.1.22.2 source gigabitEthernet0/0.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.22.2, timeout is 2 seconds:
Packet sent with a source address of 136.1.11.1 !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Verify the IPsec tunnel (child-SA); note the proxy-acl, which is specific to SVTI, and that algorithms from the default ipsec transform-set have been negotiated:

```
R1#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session

Session-id:17, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local                 Remote                fvrf/ivrf            Status
1         136.1.18.1/500        136.1.28.2/500        none/none            READY
```

https://t.me/learningnets

```
      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/8 sec
      CE id: 1030, Session-id: 17
      Status Description: Negotiation done
      Local spi: 95BD1DE7169D75EC       Remote spi: 292FE513E3E8C352
      Local id: 136.1.18.1
      Remote id: 136.1.28.2
      Local req msg id:  2              Remote req msg id:  0
      Local next msg id: 2              Remote next msg id: 0
      Local req queued:  2              Remote req queued:  0
      Local window:      5              Remote window:      5
      DPD configured for 0 seconds, retry 0
      NAT-T is not detected
      Cisco Trust Security SGT is disabled
      Initiator of SA : YesChild sa: local selector  0.0.0.0/0 - 255.255.255.255/65535
remote selector 0.0.0.0/0 - 255.255.255.255/65535
          ESP spi in/out: 0x7463A74C/0x2AB22960
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0   Encr: AES-CBC, keysize: 128, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel

 IPv6 Crypto IKEv2 Session
!
!R1#show crypto ipsec sa peer 136.1.28.2

interface: Tunnel12
    Crypto map tag: Tunnel12-head-0, local addr 136.1.18.1

   protected vrf: (none)local  ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
remote ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
   current_peer 136.1.28.2 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 17, #pkts encrypt: 17, #pkts digest: 17
    #pkts decaps: 17, #pkts decrypt: 17, #pkts verify: 17
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.1, remote crypto endpt.: 136.1.28.2
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.18
     current outbound spi: 0x2AB22960(716319072)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0x7463A74C(1952687948) transform: esp-aes esp-sha-hmac ,
```

https://t.me/learningnets

```
        in use settings ={Tunnel, }
        conn id: 2092, flow_id: Onboard VPN:92, sibling_flags 80000040, crypto map: Tunnel12-head-0
        sa timing: remaining key lifetime (k/sec): (4182480/3551)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

    inbound ah sas:

    inbound pcp sas:
outbound esp sas:
     spi: 0x2AB22960(716319072) transform: esp-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2091, flow_id: Onboard VPN:91, sibling_flags 80000040, crypto map: Tunnel12-head-0
        sa timing: remaining key lifetime (k/sec): (4182480/3551)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


    outbound ah sas:

    outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv2 L2L Between IOS Routers with DVTI and PKI

Task

- Configure IKEv2 IPsec between R1 and R2 as follows:
  - Disable any default IKEv2/IPsec settings which relate to your configuration.
  - Use a crypto-map based configuration on R1 and a DVTI configuration on R2 to encrypt traffic between loopback0 subnets.
  - Use certificate authentication; R1 is pre-configured as a CA server from initial configurations.
  - Use 3DES, MD5, and DH group 2 as negotiated algorithms.
  - All SAs should have a lifetime of 2 hours and should be deleted within 2 minutes after detecting peer failure.
- Make sure to match the following output by using an IKEv2 ID of FQDN:

```
R1#show crypto ikev2 session detailed | i inelab.local

      Local id: R1.dvti.inelab.local
      Remote id: R2.dvti.inelab.local
```

A Note on Task Initial Configuration Files: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.

If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

https://t.me/learningnets

Overview

All or just part of the smart defaults used by IKEv2/IPsec can be disabled so as not to be used by the router in any negotiations (optionally, smart-default settings can also by modified). For example, to disable the default IPsec transform-set, use the command `no crypto ipsec transform-set default` ; to re-enable, it use the command `default crypto ipsec transform-set .` To understand which default or non-default settings are being used and when, we must understand how the router will process an inbound/outbound IKEv2 negotiation:

1. First it will scan globally configured IKEv2 policies and choose the one that matches the VPN session being negotiated (available conditions to be configured are local IP address and FVRF, and the remote IP address in the future); the default IKEv2 policy is used only when no user-configured IKEv2 policy exists.
2. When the IKEv2 policy is matched, the negotiation with remote peer is done through the IKEv2 proposal used by the IKEv2 policy; basically, globally configured IKEv2 proposals have no value and are not used by the router in negotiations unless applied under a IKEv2 policy. The default IKEv2 proposal can be used even though user-configured IKEv2 proposals exist, but it must be applied under an IKEv2 policy.
3. After negotiation is successful, the IKE_SA_INIT phase is finished and keying material is derived; the VPN initiator selects the IKEv2 profile attached to the crypto map or ipsec profile and starts the IKE_AUTH phase, and it sends its IKEv2 ID to the remote peer.
4. At the VPN responder, the IKEv2 profile is matched based on the received remote IKEv2 identity, and from here the local/remote authentication method is detected.
5. Within the IKE-AUTH phase, the first CHILD_SA is also negotiated using the parameters from the crypto map or ipsec profile.

The default IPsec transform-set, unless disabled, will be used in all cases where no user-configured transform-set has been applied at the crypto map or ipsec profile level. The default IPsec profile is not used unless manually applied at the DVTI/SVTI level using the command `tunnel protection ipsec profile default .`

In IKEv2, lifetime of IKEv2 SA and IPsec SA is not negotiated, whereas in IKEv1 the IKEv1 SA lifetime was negotiated to the lowest configured value between the peers, because this setting was configured in the IKEv1 policy (which now is the IKEv2 proposal). With IKEv2 proposal, there is no lifetime configuration, and therefore no negotiation. To ensure that both VPN peers use the same lifetime values:

https://t.me/learningnets

- Configure the IKEv2 SA lifetime under the IKEv2 profile on both VPN peers; the default value is 24 hours.
- Configure the IPsec SA lifetime under the crypto map or ipsec profile on both VPN peers; the IPsec SA lifetime can also be globally configured, which applies to all IPsec SAs (the default value is 1 hour).

DPD is now standard in IKEv2, but it is disabled on IOS by default. It is configurable under the IKEv2 profile; for DPD to be negotiated, both peers must have it enabled. When you enable DPD:

- Define the interval for DPD messages if acknowledgement is received (which are informational messages in IKEv2).
- Define the interval for DPD messages if the last DPD message was not acknowledged.
- Define the mode: periodic or on-demand (the same as in IKEv1).
- Use a non-configurable value of 6 retransmits after the last non-acknowledged DPD message.

Configuration

```
R1:
no crypto ikev2 policy default
no crypto ikev2 proposal default
no crypto ipsec transform-set default
!
!
crypto ikev2 proposal R1_R2_PROPOSAL
 encryption 3des
 integrity md5
 group 2
!
crypto ikev2 policy R1_R2_POLICY
 proposal R1_R2_PROPOSAL
 match address local 136.1.18.1
!
!
crypto pki trustpoint R1_CLIENT
 enrollment url http://150.1.1.1
!
!
crypto pki authenticate R1_CLIENT
crypto pki enroll R1_CLIENT
```

https://t.me/learningnets

```
!
!
crypto ikev2 profile R1_R2_PROFILE
 match identity remote fqdn R2.dvti.inelab.local
 identity local fqdn R1.dvti.inelab.local
 authentication local rsa-sig
 authentication remote rsa-sig
 pki trustpoint R1_CLIENT
 dpd 30 15 periodic
 lifetime 7200
!
!
ip access-list extended R1_R2_ACL
 permit ip host 150.1.1.1 host 150.1.2.2
!
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto map VPN 100 ipsec-isakmp
 set peer 136.1.28.2
 match address R1_R2_ACL
 set transform-set 3DES_MD5
 set ikev2-profile R1_R2_PROFILE
 set security-association lifetime seconds 7200
!
interface GigabitEthernet0/0.18
 crypto map VPN
```

R2:

```
no crypto ikev2 policy default
no crypto ikev2 proposal default
no crypto ipsec transform-set default
no crypto ipsec profile default
!
!
crypto ikev2 proposal R1_R2_PROPOSAL
 encryption 3des
 integrity md5
 group 2
!
!
crypto ikev2 policy R1_R2_POLICY
 proposal R1_R2_PROPOSAL
```

https://t.me/learningnets

```
 match address local 136.1.28.2
!
!
crypto pki trustpoint R1_CLIENT
 enrollment url http://150.1.1.1
!
!
crypto pki authenticate R1_CLIENT
crypto pki enroll R1_CLIENT
!
!
crypto ikev2 profile R1_R2_PROFILE
 match identity remote fqdn R1.dvti.inelab.local
 identity local fqdn R2.dvti.inelab.local
 authentication local rsa-sig
 authentication remote rsa-sig
 pki trustpoint R1_CLIENT
 dpd 30 15 periodic
 lifetime 7200
 virtual-template 1
!
crypto ipsec transform-set 3DES_MD5 esp-3des esp-md5-hmac
!
crypto ipsec profile R1_R2_IPSEC
 set ikev2-profile R1_R2_PROFILE
 set transform-set 3DES_MD5
 set security-association lifetime seconds 7200
!
interface Virtual-Template1 type tunnel
 ip unnumbered GigabitEthernet0/0.28
 tunnel source GigabitEthernet0/0.28
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile R1_R2_IPSEC
```

Verification

Verify all the configured policies that are used by our configuration; default IPsec profile is not disabled on R1, because R1 has a crypto-map based configuration, so it doesn't use IPsec profiles:

```
R2#show crypto ikev2 proposal
 IKEv2 proposal: R1_R2_PROPOSAL
      Encryption : 3DES
      Integrity  : MD596
```

https://t.me/learningnets

```
     PRF        : MD5
     DH Group   : DH_GROUP_1024_MODP/Group 2 IKEv2 proposal: default Disabled
!
!R2#show crypto ikev2 policy


 IKEv2 policy : R1_R2_POLICY
      Match fvrf  : global
      Match address local : 136.1.28.2
      Proposal    : R1_R2_PROPOSAL
 IKEv2 policy : default Disabled
!
!R1#show crypto ipsec transform-set
 Transform set default: disabled
Transform set 3DES_MD5: { esp-3des esp-md5-hmac  }
   will negotiate = { Tunnel,  },
!
!R2#show crypto ipsec profile
IPSEC profile R1_R2_IPSEC
        IKEV2 profile R1_R2_PROFILE
        Security association lifetime: 4608000 kilobytes/7200 seconds
        Responder-Only (Y/N): N
        PFS (Y/N): N
        Transform sets={
                3DES_MD5:  { esp-3des esp-md5-hmac  } ,
        }
 IPSec profile default: disabled
```

Verify the configured IKEv2 profile and its settings:

```
R1#show crypto ikev2 profile


IKEv2 profile: R1_R2_PROFILE
 Ref Count: 2
 Match criteria:
  Fvrf: global
  Local address/interface: none
  Identities:  fqdn R2.dvti.inelab.local
  Certificate maps: none Local identity: fqdn R1.dvti.inelab.local
 Remote identity: none Local authentication method: rsa-sig
 Remote authentication method(s): rsa-sig
 EAP options: none
 Keyring: none
```

https://t.me/learningnets

```
Trustpoint(s):
 R1_CLIENT
Lifetime: 7200 seconds
DPD: interval 30, retry-interval 15, periodic

 NAT-keepalive: disabled
 Ivrf: none
 Virtual-template: none
 AAA EAP authentication mlist: none
 AAA Accounting: none
 AAA group authorization: none
 AAA user authorization: none
```

Generate interesting traffic so that IKEv2 negotiation is triggered:

```
R1#ping 150.1.2.2 source loopback0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.2.2, timeout is 2 seconds:
Packet sent with a source address of 150.1.1.1 .!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Verify the IKEv2 tunnel, and note negotiated algorithms and that DPD is active:

```
R1#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session
Session-id:30, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local                 Remote                fvrf/ivrf            Status
1         136.1.18.1/500        136.1.28.2/500        none/none            READY
   Encr: 3DES, Hash: MD596, DH Grp:2, Auth sign: RSA, Auth verify: RSA
Life/Active Time: 7200/25 sec
      CE id: 1044, Session-id: 30
      Status Description: Negotiation done
      Local spi: 15E2B28D142E3495       Remote spi: B8A33809888B9C1A Local id: R1.dvti.inelab.local
Remote id: R2.dvti.inelab.local
      Local req msg id:  2              Remote req msg id:  0
      Local next msg id: 2              Remote next msg id: 0
      Local req queued:  2              Remote req queued:  0
      Local window:      5              Remote window:      5
DPD configured for 30 seconds, retry 15
      NAT-T is not detected
      Cisco Trust Security SGT is disabled
      Initiator of SA : YesChild sa: local selector  150.1.1.1/0 - 150.1.1.1/65535
```

https://t.me/learningnets

```
remote selector 150.1.2.2/0 - 150.1.2.2/65535

          ESP spi in/out: 0x2E7F4A45/0x1797F2B1
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0   Encr: 3DES, esp_hmac: MD596

          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel

 IPv6 Crypto IKEv2 Session
```

Verify the IPsec tunnel (child-SA) lifetime:

```
R1#show crypto ipsec sa


interface: GigabitEthernet0/0.18
    Crypto map tag: VPN, local addr 136.1.18.1

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (150.1.1.1/255.255.255.255/0/0)
   remote ident (addr/mask/prot/port): (150.1.2.2/255.255.255.255/0/0)
   current_peer 136.1.28.2 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 4, #pkts encrypt: 4, #pkts digest: 4
 #pkts decaps: 4, #pkts decrypt: 4, #pkts verify: 4
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.1, remote crypto endpt.: 136.1.28.2
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.18
     current outbound spi: 0x1797F2B1(395834033)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0x2E7F4A45(780094021)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2118, flow_id: Onboard VPN:118, sibling_flags 80000040, crypto map: VPN
sa timing: remaining key lifetime (k/sec): (4254891/7032)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
```

https://t.me/learningnets

```
 outbound esp sas:
      spi: 0x1797F2B1(395834033)
        transform: esp-3des esp-md5-hmac ,
        in use settings ={Tunnel, }
        conn id: 2117, flow_id: Onboard VPN:117, sibling_flags 80000040, crypto map: VPN
sa timing: remaining key lifetime (k/sec): (4254891/7032)
        IV size: 8 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

Verify that the session was successfully attached to the DVTI:

```
R2#show crypto  sockets

 Number of Crypto Socket connections 1
 Vi1 Peers (local/remote): 136.1.28.2/136.1.18.1
 Local Ident  (addr/mask/port/prot): (150.1.2.2/255.255.255.255/0/0)
 Remote Ident (addr/mask/port/prot): (150.1.1.1/255.255.255.255/0/0)
 IPSec Profile: "R1_R2_IPSEC"
       Socket State: Open Client: "TUNNEL SEC" (Client State: Active)

Crypto Sockets in Listen state:
Client: "TUNNEL SEC" Profile: "R1_R2_IPSEC" Map-name: "Virtual-Template1-head-0"
!
!R2#show interfaces virtual-access 1
Virtual-Access1 is up, line protocol is up
  Hardware is Virtual Access interface
  Interface is unnumbered. Using address of GigabitEthernet0/0.28 (136.1.28.2)
  MTU 17886 bytes, BW 100000 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL Tunnel vaccess, cloned from Virtual-Template1
  Vaccess status 0x0, loopback not set
  Keepalive not set Tunnel source 136.1.28.2 (GigabitEthernet0/0.28), destination 136.1.18.1
   Tunnel Subblocks:
      src-track: Virtual-Access1 source tracking subblock associated with GigabitEthernet0/0.28
          Set of tunnels with source GigabitEthernet0/0.28, 2 members (includes iterators), on interface <OK>
 Tunnel protocol/transport IPSEC/IP
  Tunnel TTL 255
  Tunnel transport MTU 1446 bytes
  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps)
```

https://t.me/learningnets

```
Tunnel protection via IPSec (profile "R1_R2_IPSEC")

  Last input never, output never, output hang never
  Last clearing of "show interface" counters 00:06:15
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 86
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     4 packets input, 400 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     90 packets output, 5624 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv2 FlexVPN Server and Client with PSK

Task

- Configure R1 as the FlexVPN client and R2 as the FlexVPN server:
  - Make use of the default IKEv2/IPsec configurations.
  - Enable PSK authentication with the string **flexkey**.
  - The flexvpn server should assign an IPv4 address from subnet 136.1.12.0/24 to the flexvpn client.
  - Ensure that R1 installs a default route and that R2 installs a route for VLAN 11 subnet through the IPsec tunnel, both with a tag value of 12.
- Make sure to match the following output by using an IKEv2 ID of FQDN:

```
R1#show crypto ikev2 session detailed | i inelab.local

      Local id: flexclient.inelab.local
      Remote id: flexserver.inelab.local
```

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IKEv1 has multiple options for configuring IPsec VPNs on IOS (crypto map, DVTI,

https://t.me/learningnets

SVTI, EzVPN, DMVPN, GET VPN), each with its own set of advantages and limitations, but most importantly there is poor or no interoperability possible between different VPN types. For example, a DMVPN configuration is not compatible with EzVPN clients, or crypto map VPNs do not support dynamic routing and per-peer QoS. With the addition of IKEv2 on IOS, Cisco created a new IPsec configuration framework named FlexVPN, which unifies all types of VPNs into a common command and configuration block. As we've seen from previous IKEv2 tasks, all IKEv1 IPsec VPNs have also been migrated to IKEv2 IPsec VPNs (except GET VPN, which is in development, and EzVPN client/server, which now uses flexvpn client/server functionality); so FlexVPN and previous VPN configurations can coexist on the same IOS device. While with FlexVPN there is a common command block for all types of IKEv2 IPsec VPNs (site-to-site, remote-access hub-and-spoke), there are additional specific configurations for each type of VPN.

The core of FlexVPN is based on DVTI, which offers a unique connection point and common configuration for all VPNs through the use of virtual-templates. However, after the IPsec tunnel is successfully negotiated, a virtual-access interface is automatically created that allows per-peer specific configurations to be applied in the form of authorization attributes (QoS, ACL, firewall), using either the local AAA functionality or a remote RADIUS server.

Just like in EzVPN, the FlexVPN server is an IOS device, whereas the FlexVPN client can be hardware (IOS device) or software (compatible IKEv2 clients such as AnyConnect or the one built in to Windows 7 and Windows 8); ASA code does not support any FlexVPN functionality. From a configuration perspective, FlexVPN server uses a regular DVTI setup, whereas FlexVPN client use a special SVTI setup (crypto map is not supported for FlexVPN server/client). The SVTI has a **`tunnel destination dynamic`** configuration, which allows the same tunnel interface on the FlexVPN client to be used with multiple FlexVPN servers; the tunnel destination is automatically detected from the FlexVPN client configuration. For example:

```
crypto ikev2 client flexvpn R1_R2_CLIENT
 peer 1 136.1.28.2
 peer 2 136.1.38.3
 client connect Tunnel12
```

All mandatory and optional settings pushed from the FlexVPN server to the FlexVPN client (such as IP addressing, DNS servers, and domain-name) can be locally configured on the FlexVPN server or received from the AAA server through RADIUS. If the local option is used, the configuration occurs under an IKEv2 authorization policy, and you can use the default one or create your own:

```
R1#show crypto ikev2 authorization policy default
```

https://t.me/learningnets

```
IKEv2 Authorization Policy : default
 route set interface
 route accept any tag : 1 distance : 1
```

The IKEv2 authorization policy must to be applied at the IKEv2 profile level using the command **`aaa authorization [group|user] [cert|eap|psk] list <AAA_LIST_NAME> <IKEv2_AUTHORIZATION_NAME>`**, and it requires the following:

- AAA new-model to be enabled
- AAA network authorization list name
- IKEv2 remote authentication method: RSA-SIG, EAP, PSK

With FlexVPN, RRI is no longer used. Instead, IKEv2 routing is implemented, which allows each VPN peer to install one or more routes in the routing table of the remote VPN peer; this configuration is done at the IKEv2 authorization policy level as follows:

- The command **`route set access-list <ACL_NAME>`** pushes the prefixes from the ACL as static routes in the remote VPN peer routing table.
- The command **`route set interface`** is configured on the FlexVPN client and pushes its dynamically assigned IP address as static routes in the FlexVPN server routing table.
- The command **`route accept any [tag|distance] <value>`** accepts all routes from the remote VPN peer; optionally, you can assign a tag value or change the default administrative distance of 1.

All the authorization attributes are being negotiated in the IKE_AUTH phase through informational type of messages; Cisco calls this the **Configuration Exchange**, or **Configuration Payload** functionality, and it consists of four messages:

- Initiator sends a CP(CFG_REQUEST) in the IKE_AUTH_i message, asking for configuration (like IP address).
- Responder sends a CP(CFG_REPLY) in the IKE_AUTH_r message, replying with the requested configuration.
- Initiator sends a CP(CFG_SET) in the IKE_AUTH_r message, pushing configuration like IKEv2 routes.
- Responder sends a CP(CFG_ACK) in the ACK message, which is just the acknowledgement.

https://t.me/learningnets

FlexVPN framework supports only IKEv2.

Configuration

R1:

```
default crypto ikev2 policy
default crypto ikev2 proposal
default crypto ipsec transform-set
default crypto ipsec profile
default crypto ikev2 authorization policy
!
!
aaa new-model
aaa authorization network IKE_LIST local
!
!
interface Tunnel12
 ip address negotiated
 tunnel source GigabitEthernet0/0.18
 tunnel destination dynamic
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile default
!
!
crypto ikev2 client flexvpn R1_R2_CLIENT
 peer 1 136.1.28.2
 connect auto
 client connect Tunnel12
!
!
ip access-list standard PROTECTED_ACL
 permit 136.1.11.0 0.0.0.255
!
!
crypto ikev2 authorization policy default
 route accept any tag 12
 route set interface
 route set access-list PROTECTED_ACL
!
!
crypto ikev2 keyring R1_R2_KEYRING
 peer R2
  address 136.1.28.2
  pre-shared-key local flexkey
```

https://t.me/learningnets

```
  pre-shared-key remote flexkey
!
!
crypto ikev2 profile R1_R2_PROFILE
 match identity remote fqdn flexserver.inelab.local
 identity local fqdn flexclient.inelab.local
 authentication local pre-share
 authentication remote pre-share
 keyring local R1_R2_KEYRING
 aaa authorization group psk list IKE_LIST default
!
!
crypto ipsec profile default
 set ikev2-profile R1_R2_PROFILE
```

R2:

```
default crypto ikev2 policy
default crypto ikev2 proposal
default crypto ipsec transform-set
default crypto ipsec profile
default crypto ikev2 authorization policy
!
!
aaa new-model
aaa authorization network IKE_LIST local
!
!
ip local pool IKE_POOL 136.1.12.1 136.1.12.254
!
ip access-list standard PROTECTED_ACL
 permit any
!
!
crypto ikev2 authorization policy default
 pool IKE_POOL
 route accept any tag 12
 no route set interface
 route set access-list PROTECTED_ACL
!
!
crypto ikev2 keyring R1_R2_KEYRING
 peer ALL
  address 0.0.0.0 0.0.0.0
  pre-shared-key local flexkey
```

https://t.me/learningnets

```
  pre-shared-key remote flexkey
!
!
crypto ikev2 profile R1_R2_PROFILE
 match identity remote fqdn flexclient.inelab.local
 identity local fqdn flexserver.inelab.local
 authentication local pre-share
 authentication remote pre-share
 keyring local R1_R2_KEYRING
 aaa authorization group psk list IKE_LIST default
 virtual-template 1
!
!
crypto ipsec profile default
 set ikev2-profile R1_R2_PROFILE
!
!
interface Virtual-Template1 type tunnel
 ip unnumbered GigabitEthernet0/0.28
 tunnel source GigabitEthernet0/0.28
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile default
```

Verification

Verify the default authorization policy; settings have been modified:

```
R1#show crypto ikev2 authorization policy default
 IKEv2 Authorization Policy : default route set interface
route set acl: PROTECTED_ACL
route accept any tag : 12 distance : 1
!
!R2#show crypto ikev2 authorization policy default
 IKEv2 Authorization Policy : default IPV4 Address Pool : IKE_POOL
route set acl: PROTECTED_ACL
route accept any tag : 12 distance : 1
```

Verify the flexvpn client connection status:

```
R1#show crypto ikev2 client flexvpn


  Profile : R1_R2_CLIENT
```

https://t.me/learningnets

```
Current state:ACTIVE

  Peer : 136.1.28.2
  Source : GigabitEthernet0/0.18
  ivrf : IP DEFAULT
  fvrf : IP DEFAULT
  Backup group: Default
  Tunnel interface : Tunnel12 Assigned IP address: 136.1.12.1
```

Verify the flexvpn client tunnel interface:

https://t.me/learningnets

```
R1#show interfaces tunnel12
Tunnel12 is up, line protocol is up

  Hardware is Tunnel Internet address is 136.1.12.1/32

  MTU 17878 bytes, BW 100 Kbit/sec, DLY 50000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 136.1.18.1 (GigabitEthernet0/0.18), destination 136.1.28.2
   Tunnel Subblocks:
      src-track:
         Tunnel12 source tracking subblock associated with GigabitEthernet0/0.18
          Set of tunnels with source GigabitEthernet0/0.18, 1 member (includes iterators), on interface <OK>
  Tunnel protocol/transport IPSEC/IP
  Tunnel TTL 255
  Tunnel transport MTU 1438 bytes
  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps) Tunnel protection via IPSec (profile "default")

  Last input never, output never, output hang never
  Last clearing of "show interface" counters 01:31:20
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     10 packets input, 1000 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     17 packets output, 3176 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
```

Verify the IKEv2 session status:

```
R1#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session
 Session-id:1, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local                 Remote                fvrf/ivrf            Status
```

https://t.me/learningnets

```
1         136.1.18.1/500        136.1.28.2/500        none/none            READY
      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/3287 sec
      CE id: 1001, Session-id: 1
      Status Description: Negotiation done
      Local spi: A293B893EB2B86BC       Remote spi: 9DF37FD25EF5F0F0 Local id: flexclient.inelab.local
Remote id: flexserver.inelab.local
      Local req msg id:  3              Remote req msg id:  0
      Local next msg id: 3              Remote next msg id: 0
      Local req queued:  3              Remote req queued:  0
      Local window:      5              Remote window:      5
      DPD configured for 0 seconds, retry 0
      NAT-T is not detected
      Cisco Trust Security SGT is disabled Initiator of SA : Yes
Pushed IP address: 136.1.12.1
Remote subnets:
0.0.0.0 0.0.0.0
Child sa: local selector  0.0.0.0/0 - 255.255.255.255/65535
          remote selector 0.0.0.0/0 - 255.255.255.255/65535
          ESP spi in/out: 0xAAB54941/0xD1DC0EF3
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0
          Encr: AES-CBC, keysize: 128, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel

 IPv6 Crypto IKEv2 Session
!
!R2#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session
Session-id:1, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local                 Remote                fvrf/ivrf            Status
1         136.1.28.2/500        136.1.18.1/500        none/none            READY
      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/3565 sec
      CE id: 1001, Session-id: 1
      Status Description: Negotiation done
      Local spi: 9DF37FD25EF5F0F0       Remote spi: A293B893EB2B86BC
      Local id: flexserver.inelab.local
      Remote id: flexclient.inelab.local
      Local req msg id:  0              Remote req msg id:  5
      Local next msg id: 0              Remote next msg id: 5
      Local req queued:  0              Remote req queued:  5
      Local window:      5              Remote window:      5
      DPD configured for 0 seconds, retry 0
```

https://t.me/learningnets

```
      NAT-T is not detected
      Cisco Trust Security SGT is disabled Assigned host addr: 136.1.12.1
 Initiator of SA : No
 Remote subnets:
 136.1.12.1 255.255.255.255
 136.1.11.0 255.255.255.0

Child sa: local selector  0.0.0.0/0 - 255.255.255.255/65535
          remote selector 0.0.0.0/0 - 255.255.255.255/65535
          ESP spi in/out: 0x1ABC0A0F/0x2E4C3B5B
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0
          Encr: AES-CBC, keysize: 128, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode tunnel

 IPv6 Crypto IKEv2 Session
```

Verify the installed IKEv2 routes:

```
R1#show ip route static | b Gateway
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
 S*    0.0.0.0/0 is directly connected, Tunnel12
!
!R1#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  Known via "static", distance 1, metric 0 (connected), candidate default path Tag 12
  Routing Descriptor Blocks:
  * directly connected, via Tunnel12, permanent
      Route metric is 0, traffic share count is 1
      Route tag 12
!
!R2#show ip route static | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 9 subnets, 2 masks
 S        136.1.11.0/24 is directly connected, Virtual-Access1
 S        136.1.12.1/32 is directly connected, Virtual-Access1
!
!R2#show ip route 136.1.11.0
Routing entry for 136.1.11.0/24
  Known via "static", distance 1, metric 0 (connected) Tag 12

  Routing Descriptor Blocks:
  * directly connected, via Virtual-Access1, permanent
      Route metric is 0, traffic share count is 1
```

https://t.me/learningnets

```
      Route tag 12
```

Verify that traffic is encrypted through the IPsec tunnel:

```
R2#ping 136.1.11.1 source gigabitEthernet0/0.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.11.1, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R2#ping 136.1.12.1 source gigabitEthernet0/0.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.12.1, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R2#show crypto ipsec sa


interface: Virtual-Access1
    Crypto map tag: Virtual-Access1-head-0, local addr 136.1.28.2

   protected vrf: (none)
   local  ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
   remote ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
   current_peer 136.1.18.1 port 500
     PERMIT, flags={origin_is_acl,} #pkts encaps: 10, #pkts encrypt: 10, #pkts digest: 10
#pkts decaps: 10, #pkts decrypt: 10, #pkts verify: 10
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.28.2, remote crypto endpt.: 136.1.18.1
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.28
     current outbound spi: 0x2E4C3B5B(776747867)
     PFS (Y/N): N, DH group: none
 inbound esp sas:
      spi: 0x1ABC0A0F(448530959)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2003, flow_id: Onboard VPN:3, sibling_flags 80000040, crypto map: Virtual-Access1-head-0
        sa timing: remaining key lifetime (k/sec): (4180326/2211)
        IV size: 16 bytes
        replay detection support: Y
```

https://t.me/learningnets

```
Status: ACTIVE(ACTIVE)



     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0x2E4C3B5B(776747867)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Tunnel, }
        conn id: 2004, flow_id: Onboard VPN:4, sibling_flags 80000040, crypto map: Virtual-Access1-head-0
        sa timing: remaining key lifetime (k/sec): (4180327/2211)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)



     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 7: Confidentiality and Secure Access

IKEv2 FlexVPN Spoke to Spoke with PSK

Task

- Using FlexVPN framework, configure R2 as DMVPN hub and R1 as DMVPN spoke:
  - Make use of the default IKEv2/IPsec configurations.
  - Enable PSK authentication with the string **dmvpnkey**.
  - Using EIGRP AS 12, allow traffic between VLAN 11 and VLAN 22 to be protected by IPsec.
  - Using the e-mail address IKev2 ID, ensure that ICMP traffic between R1 and R2 is logged on the hub; this should not apply to any other spokes.
- Make sure to match the following output:

```
R1#show ip eigrp 12 neighbors

EIGRP-IPv4 Neighbors for AS(12)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   2.2.2.2                 Tu12                     10 00:23:38   13  1434  0  3
!
!R2#show ip eigrp 12 neighbors

EIGRP-IPv4 Neighbors for AS(12)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   1.1.1.1                 Vi1                      12 00:23:59    2  1434  0  3
```

https://t.me/learningnets

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the **DMVPN & GETVPN & IKEv2** module of this section, which can be found in the **Section 7 Introduction** by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

DMVPN can be implemented with IKEv2 using the same configuration as in IKEv1, except that IKEv1 is replaced with IKev2, so from a functionality perspective nothing changes on the DMVPN side (remember that encryption/IPsec is optional on DMVPN); this is basically the now legacy DMVPN, which works with both IKEv1 and IKEv2. However, within the FlexVPN framework, Cisco added DMVPN functionality and changed both the configuration (to map it to the FlexVPN command set) and the functionality (to further optimize it).

Here is the configuration required on the hub side, using the IKEv2 smart defaults:

- Configure IKEv2 keyring (if PSK authentication is used) or enroll with a PKI (if certificate authentication is used).
- Configure IKEv2 profile:
  - Specify local/remote authentication.
  - Apply the IKEv2 keyring and/or the trustpoint.
  - Match the remote IKEv2 ID and optionally configure the local IKEv2 ID.
  - Attach the virtual-template.
- Apply the IKEv2 profile to an IPsec profile.
- Configure the DVTI (virtual-template), used to clone a virtual-access interface for each spoke IPsec tunnel:
  - Use GRE mode so that non-IP traffic (NHRP) is allowed through the tunnel.
  - Specify the NHRP ID (identifying the DMVPN cloud and enabling NHRP).
  - Configure NHRP redirect (to trigger spoke-to-spoke tunnel negotiation).
  - Apply the IPsec profile.

Here is the configuration required on the spoke side, using the IKEv2 smart defaults:

- Configure IKEv2 keyring (if PSK authentication is used) or enroll with a PKI (if

https://t.me/learningnets

certificate authentication is used).
- Configure IKEv2 profile:
  - Specify local/remote authentication.
  - Apply the IKEv2 keyring and/or the truspoint.
  - Match the remote IKEv2 ID and optionally configure the local IKEv2 ID.
  - Attach the virtual-template.
- Apply the IKEv2 profile to an IPsec profile.
- Configure an SVTI, used for spoke-to-hub IPsec tunnel:
  - Use GRE mode so that non-IP traffic (NHRP) is allowed through the tunnel.
  - Specify the tunnel destination as being the hub.
  - Specify the NHRP ID (identifying the DMVPN cloud and enabling NHRP).
  - Configure NHRP shortcut (to allow cloning a virtual-access interface for each spoke-to-spoke tunnel).
  - Apply the IPsec profile.
- Configure the DVTI (virtual-template), used for spoke-to-spoke IPsec tunnels:
  - Use GRE mode so that non-IP traffic (NHRP) is allowed through the tunnel.
  - Specify the NHRP ID (identifying the DMVPN cloud and enabling NHRP).
  - Configure NHRP shortcut (to allow remote spoke routes to be installed by NHO or NHRP through the cloned virtual-access interface).
  - Apply the IPsec profile.

From a functionality perspective, taking one hub and two spokes (A and B), the following occurs:

- First the spoke-to-hub IPsec tunnel is established by both spokes; however, there is no NHRP traffic at this point, so the spokes no longer use NHRP Registration Request.
  - This was somehow visible from the configuration, as NHS is no longer configured on the spoke.
  - This is an enhancement from the control-plane perspective, because on DMVPN Phase3 there was no real advantage of the hub having dynamic NHRP entries for all spokes.
  - Basically NHRP is now used only for spoke-to-spoke IPsec tunnel.
- Spoke A sends traffic to spoke B, and the hub receives it inbound on its virtual-access interface for spoke A and routes it over the virtual-access interface of spoke B.
- Spoke B replies to spoke A, and the hub receives it inbound on its virtual-access interface for spoke B and routes it over the virtual-access interface of spoke A.

https://t.me/learningnets

- Because the hub has NHRP redirect configured and ingress/egress virtual-access interfaces for the spoke-to-spoke traffic belongs to same DMVPN cloud (NHRP network ID), the hub sends NHRP redirect to both spokes.
    - This triggers the NHRP Resolution request on both spokes.
- Both spokes send NHRP Resolution requests through the hub IPsec SVTI tunnel interface and the hub relays it between spokes (through virtual-access interfaces).
- Both spokes receive the NHRP Resolution request and trigger the spoke-to-spoke IPsec tunnel negotiation using the virtual-template attached to the hub IPsec SVTI tunnel interface.
- The spoke-to-spoke IPsec tunnel is negotiated, virtual-access interfaces are created on spokes, and NHRP Resolution reply is sent over it.
- The NHRP Resolution reply, based on the DMVPN cloud routing setup, triggers the NHO (Next-Hop-Override) or NHRP routes to be installed for remote spoke subnet.

Optionally, because using DVTI templates clones a virtual-access interface for each spoke, the FlexVPN server (hub) can apply specific authorizations to each spoke—for example, interface-level configurations: policing or shaping, access-list or firewall, MTU, tunnel mode, dynamic routing. This is similar to IKEv2 routing configuration, and thus can be locally configured or received from the RADIUS server. Some authorization attributes, such as ACL attachment or CBAC inspection (which basically require dynamic application of some commands to the virtual-access interface after the session is established), cannot be directly configured in the IKEv2 authorization policy and are defined through AAA attribute-list.

In the previous task, we applied the IKEv2 authorization list under the IKEv2 profile, which will force the same attributes to all IPsec VPN sessions matching the IKEv2 profile. If you want to apply different attributes to each remote VPN peer, you must:

- Configure a name-mangler and match on the complete IKEv2 ID of the remote-peer or on specific parts of it (not supported when remote IKEv2 ID is IP address).
- Configure IKEv2 authorization policies with names identical to the matched IKEv2 ID from the name-mangler.
- Apply the name-mangler under the IKEv2 profile.

Based on the IKEv2 ID, within a name-mangler you can match on the complete IKEv2 ID or parts of it:

- If the IKEv2 ID is e-mail: match on username or domain-name, where for **instructor@ine.com** , **instructor** is the username and **ine.com** is the domain-name.
- If the IKEv2 ID is FQDN: match on hostname or domain-name, where for **router.ine.com**

https://t.me/learningnets

, **router** is the hostname and **ine.com** is the domain-name.

- If the IKEv2 ID is DN: match different certificate attributes such as: common-name, country, domain, locality, organization, organization-unit, state.

If, for example, you are using IKEv2 ID of FQDN and you want to apply different authorization attributes to remote VPN peers based on the hostname, assuming your spokes are named **ROUTER-A**, **ROUTER-B** and **ROUTER-C**, you must configure IKEv2 authorization profiles with these identical names. For this process to work, you no longer apply the IKEv2 authorization policy at the IKEv2 profile level; instead you apply the name-mangler using the command `aaa authorization [group|user] [cert|eap|psk] list <AAA_LIST_NAME> name-mangler <MANGLER_NAME> .`

The order of processing is:

1. IKE_SA_INIT is successfully finished and IKEv2 ID exchange is performed.
2. The initiator (spoke) sends its IKEv2 ID in the first IKE_AUTH message.
3. The responder (hub) will match the received IKEv2 ID against its IKEv2 profile, which also has the name-mangler applied.
4. The responder queries the name-mangler and receives the part to be retrieved from the received IKEv2 ID of the remote-peer.
5. The responder will perform authorization based on the applied network authorization list from the IKEv2 profile.
6. For local authorization, the responder will look in the configuration for a IKEv2 authorization policy name matching the IKEv2 ID part identified by the name-mangler.
7. Authorization attributes from the IKEv2 authorization policy are applied to the session.

Configuration

R1:

```
crypto ikev2 keyring DMVPN_KEYRING
 peer ALL
  address 0.0.0.0 0.0.0.0
  pre-shared-key local dmvpnkey
  pre-shared-key remote dmvpnkey
!
!
crypto ikev2 profile DMVPN_PROFILE
 match identity remote email domain inelab.local
 identity local email SPOKE1@inelab.local
 authentication local pre-share
 authentication remote pre-share
 keyring local DMVPN_KEYRING
```

https://t.me/learningnets

```
 virtual-template 1
!
!
crypto ipsec profile default
 set ikev2-profile DMVPN_PROFILE
!
!
interface Loopback1
 ip address 1.1.1.1 255.255.255.255
!
interface Tunnel12
 ip unnumbered Loopback1
 tunnel source GigabitEthernet0/0.18
 tunnel destination 136.1.28.2
 tunnel mode gre ip
 ip nhrp network-id 12
 ip nhrp shortcut virtual-template 1
 tunnel protection ipsec profile default
!
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback1
 tunnel source GigabitEthernet0/0.18
 tunnel mode gre ip
 ip nhrp network-id 12
 ip nhrp shortcut virtual-template 1
 tunnel protection ipsec profile default
!
!
router eigrp 12
 no auto-summary
 network 1.1.1.1 0.0.0.0
 network 136.1.11.0 0.0.0.255
 passive-interface default
 no passive-interface Tunnel12
```

R2:

```
aaa new-model
aaa authorization network DMVPN_LIST local
!
aaa attribute list SPOKE1_LIST
 attribute type interface-config “ip access-group DMVPN_ACL in” protocol ip
 attribute type interface-config “ip access-group DMVPN_ACL out” protocol ip
```

https://t.me/learningnets

```
!
!
ip access-list extended DMVPN_ACL
 permit icmp any any log
 permit ip any any
!
!
crypto ikev2 authorization policy SPOKE1
 no route set interface
 no route accept any
 aaa attribute list SPOKE1_LIST
!
!
crypto ikev2 name-mangler DMVPN_MANGLER
 email username
!
!
crypto ikev2 keyring DMVPN_KEYRING
 peer ALL
  address 0.0.0.0 0.0.0.0
  pre-shared-key local dmvpnkey
  pre-shared-key remote dmvpnkey
!
!
crypto ikev2 profile DMVPN_PROFILE
 match identity remote email domain inelab.local
 identity local email HUB@inelab.local
 authentication local pre-share
 authentication remote pre-share
 keyring local DMVPN_KEYRING
 aaa authorization group psk list DMVPN_LIST name-mangler DMVPN_MANGLER
 virtual-template 1
!
!
crypto ipsec profile default
 set ikev2-profile DMVPN_PROFILE
!
!
interface Loopback1
 ip address 2.2.2.2 255.255.255.255
!
interface Virtual-Template1 type tunnel
 ip unnumbered Loopback1
 tunnel source GigabitEthernet0/0.28
 tunnel mode gre ip
 ip nhrp network-id 12
```

https://t.me/learningnets

```
 ip nhrp redirect
 tunnel protection ipsec profile default
!
!
router eigrp 12
 no auto-summary
 network 2.2.2.2 0.0.0.0
 network 136.1.22.0 0.0.0.255
 passive-interface default
 no passive-interface Virtual-Template1
```

Verification

Verify the IKEv2 session status:

```
R1#show crypto session detail

Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
X - IKE Extended Authentication, F - IKE Fragmentation

Interface: Tunnel12
Uptime: 00:58:02 Session status: UP-ACTIVE
     Peer: 136.1.28.2 port 500 fvrf: (none) ivrf: (none)
      Phase1_id: HUB@inelab.local
      Desc: (none)
  IKEv2 SA: local 136.1.18.1/500 remote 136.1.28.2/500 Active
         Capabilities:(none) connid:1 lifetime:23:01:58
 IPSEC FLOW: permit 47 host 136.1.18.1 host 136.1.28.2
        Active SAs: 2, origin: crypto map
        Inbound:  #pkts dec'ed 797 drop 0 life (KB/Sec) 4250430/3592
        Outbound: #pkts enc'ed 830 drop 0 life (KB/Sec) 4250430/3592
!
!R1#show crypto ikev2 session detailed

 IPv4 Crypto IKEv2 Session
 Session-id:1, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local                 Remote                fvrf/ivrf            Status
1         136.1.18.1/500        136.1.28.2/500        none/none            READY
      Encr: AES-CBC, keysize: 256, Hash: SHA512, DH Grp:5, Auth sign: PSK, Auth verify: PSK
      Life/Active Time: 86400/3574 sec
      CE id: 1001, Session-id: 1
```

https://t.me/learningnets

```
      Status Description: Negotiation done
      Local spi: 8FD083CF06691B30       Remote spi: E70D9FE9A275A0C4 Local id: SPOKE1@inelab.local
Remote id: HUB@inelab.local
      Local req msg id:  2              Remote req msg id:  2
      Local next msg id: 2              Remote next msg id: 2
      Local req queued:  2              Remote req queued:  2
      Local window:      5              Remote window:      5
      DPD configured for 0 seconds, retry 0
      NAT-T is not detected
      Cisco Trust Security SGT is disabled
      Initiator of SA : YesChild sa: local selector  136.1.18.1/0 - 136.1.18.1/65535
remote selector 136.1.28.2/0 - 136.1.28.2/65535

          ESP spi in/out: 0x5B06DFFC/0x809A188A
          AH spi in/out: 0x0/0x0
          CPI in/out: 0x0/0x0
          Encr: AES-CBC, keysize: 128, esp_hmac: SHA96
          ah_hmac: None, comp: IPCOMP_NONE, mode transport

 IPv6 Crypto IKEv2 Session
```

Because there is only one spoke, and therefore no spoke-to-spoke session, no NHRP entries are created:

```
R1#show ip nhrp
 R1#
!
!R2#show ip nhrp
 R2#
```

Verify that the EIGRP adjacency in AS 12:

https://t.me/learningnets

```
R1#show ip eigrp 12 neighbors

EIGRP-IPv4 Neighbors for AS(12)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   2.2.2.2                 Tu12                     10 00:23:38   13  1434  0  3

!
!R2#show ip eigrp 12 neighbors

EIGRP-IPv4 Neighbors for AS(12)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   1.1.1.1                 Vi1                      12 00:23:59    2  1434  0  3
```

Verify learned routes through EIGRP AS 12 and IP connectivity:

```
R1#show ip route eigrp 12 | b Gateway
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/27008000] via 2.2.2.2, 00:45:15, Tunnel12
      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D        136.1.22.0/24 [90/26882560] via 2.2.2.2, 00:45:15, Tunnel12
!
!R2#show ip route eigrp 12 | b Gateway
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
D        1.1.1.1 [90/27008000] via 1.1.1.1, 00:46:07, Virtual-Access1
      136.1.0.0/16 is variably subnetted, 8 subnets, 2 masks
D        136.1.11.0/24 [90/26882560] via 1.1.1.1, 00:46:07, Virtual-Access1
!
!R2#ping 136.1.11.1 source gigabitEthernet0/0.22
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.11.1, timeout is 2 seconds:
Packet sent with a source address of 136.1.22.2 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
!
!R2#ping 1.1.1.1 source loopback1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2 !!!!!
```

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Because of the authorization applied on the hub for the cloned virtual-access interface, ICMP traffic will be logged and the following messages will be displayed:

```
%SEC-6-IPACCESSLOGDP: list DMVPN_ACL permitted icmp 1.1.1.1 -> 2.2.2.2 (0/0), 5 packet
%SEC-6-IPACCESSLOGDP: list DMVPN_ACL permitted icmp 136.1.11.1 -> 136.1.22.2 (0/0), 5 packets
```

Verify that the ACL has been dynamically applied on the virtual-access interface:

```
R2#show ip interface virtual-template1 | i access list
 Outgoing access list is not set
 Inbound  access list is not set
!
!R2#show ip interface virtual-access 1 | i access list
 Outgoing access list is DMVPN_ACL, default is not set
 Inbound  access list is DMVPN_ACL, default is not set
```

Verify that on the spoke the IPsec tunnel with the hub is attached to the SVTI interface, not to the virtual-template interface:

```
R1#show crypto ipsec sa

 interface: Tunnel12
 Crypto map tag: Tunnel12-head-0, local addr 136.1.18.1

   protected vrf: (none) local  ident (addr/mask/prot/port): (136.1.18.1/255.255.255.255/47/0)
 remote ident (addr/mask/prot/port): (136.1.28.2/255.255.255.255/47/0)
   current_peer 136.1.28.2 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 2761, #pkts encrypt: 2761, #pkts digest: 2761
    #pkts decaps: 2731, #pkts decrypt: 2731, #pkts verify: 2731
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0

     local crypto endpt.: 136.1.18.1, remote crypto endpt.: 136.1.28.2
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.18
     current outbound spi: 0xD006F589(3490117001)
     PFS (Y/N): N, DH group: none
```

https://t.me/learningnets

```
inbound esp sas:

      spi: 0xFB82AD4C(4219645260)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Transport, }
        conn id: 2007, flow_id: Onboard VPN:7, sibling_flags 80000000, crypto map: Tunnel12-head-0
        sa timing: remaining key lifetime (k/sec): (4166278/1541)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:
      spi: 0xD006F589(3490117001)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Transport, }
        conn id: 2008, flow_id: Onboard VPN:8, sibling_flags 80000000, crypto map: Tunnel12-head-0
        sa timing: remaining key lifetime (k/sec): (4166278/1541)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

Verify the IPsec tunnel on the hub, which is attached to the cloned virtual-access interface:

```
R2#show crypto ipsec sa

interface: Virtual-Access1
Crypto map tag: Virtual-Access1-head-0, local addr 136.1.28.2

   protected vrf: (none) local  ident (addr/mask/prot/port): (136.1.28.2/255.255.255.255/47/0)
remote ident (addr/mask/prot/port): (136.1.18.1/255.255.255.255/47/0)
   current_peer 136.1.18.1 port 500
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 2743, #pkts encrypt: 2743, #pkts digest: 2743
    #pkts decaps: 2774, #pkts decrypt: 2774, #pkts verify: 2774
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0
```

https://t.me/learningnets

```
     local crypto endpt.: 136.1.28.2, remote crypto endpt.: 136.1.18.1
     path mtu 1500, ip mtu 1500, ip mtu idb GigabitEthernet0/0.28
     current outbound spi: 0xFB82AD4C(4219645260)
     PFS (Y/N): N, DH group: none
inbound esp sas:
      spi: 0xD006F589(3490117001)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Transport, }
        conn id: 2008, flow_id: Onboard VPN:8, sibling_flags 80000000, crypto map: Virtual-Access1-head-0
        sa timing: remaining key lifetime (k/sec): (4285287/1483)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)

     inbound ah sas:

     inbound pcp sas:
outbound esp sas:)
      spi: 0xFB82AD4C(4219645260)
        transform: esp-aes esp-sha-hmac ,
        in use settings ={Transport, }
        conn id: 2007, flow_id: Onboard VPN:7, sibling_flags 80000000, crypto map: Virtual-Access1-head-0
        sa timing: remaining key lifetime (k/sec): (4285287/1483)
        IV size: 16 bytes
        replay detection support: Y Status: ACTIVE(ACTIVE)


     outbound ah sas:

     outbound pcp sas:
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

Section 6 Introduction

This section utilizes both Firewall Services and Security Services.

Lab Topics

- Cisco IOS Zone-Based Firewall
  - Network, Secure Group, and User Based Policy
  - Performance Tuning
  - Network, Protocol, and Application Inspection
- Perimeter Security Services
  - Cisco IOS QoS and Packet marking techniques
  - Traffic Filtering using Access-Lists
  - Cisco IOS NAT
  - uRPF
  - PAM - Port to Application Mapping
  - Policy Routing and Route Maps

These topics are further expanded in the following checklist.

Technology	Mastered?
Cisco IOS Firewalls	
Zone-Based Policy Firewall Using Multiple-Zone Scenarios	

https://t.me/learningnets

Technology	Mastered?
User-Based Firewall	
Secure-Group Firewall	
Transparent Cisco IOS Firewall (Layer 2)	
Context-Based Access Control (CBAC)	
Proxy Authentication (Auth Proxy)	
Port-to-Application Mapping (PAM) Usage with ACLs	
Use of PAM to Change System Default Ports	
PAM Custom Ports for Specific Applications	
Mapping Nonstandard Ports to Standard Applications	
Performance Tuning	
Tuning Half-Open Connections	
Understanding and Interpreting the show ip port-map Commands	
Understanding and Interpreting the show ip inspect Commands	
Understanding and Interpreting the debug ip inspect Commands	

https://t.me/learningnets

Technology	Mastered?
Understanding and Interpreting the show zone zone-pair Commands	
Understanding and Interpreting the debug zone Commands	
Cisco IOS Services	
Marking Packets Using DSCP and IP Precedence and Other Values	
Unicast RPF (uRPF) With or Without an ACL (Strict and Loose Mode)	
RTBH Filtering (Remote Triggered Black Hole)	
Basic Traffic Filtering Using Access Lists: SYN Flags, Established, etc. (Named vs. Numbered ACLs)	
Managing Time-Based Access Lists	
Enabling NAT and PAT on a Router	
Conditional NAT on a Router	
Multihome NAT on a Router	
CAR Rate Limiting with Traffic Classification Using ACLs	
PBR (Policy-Based Routing) and Use of Route Maps	

https://t.me/learningnets

Technology	Mastered?
Traffic Policing on a Router	
Traffic Characterization	
Packet Classification	
Packet-Marking Techniques	

Configurations

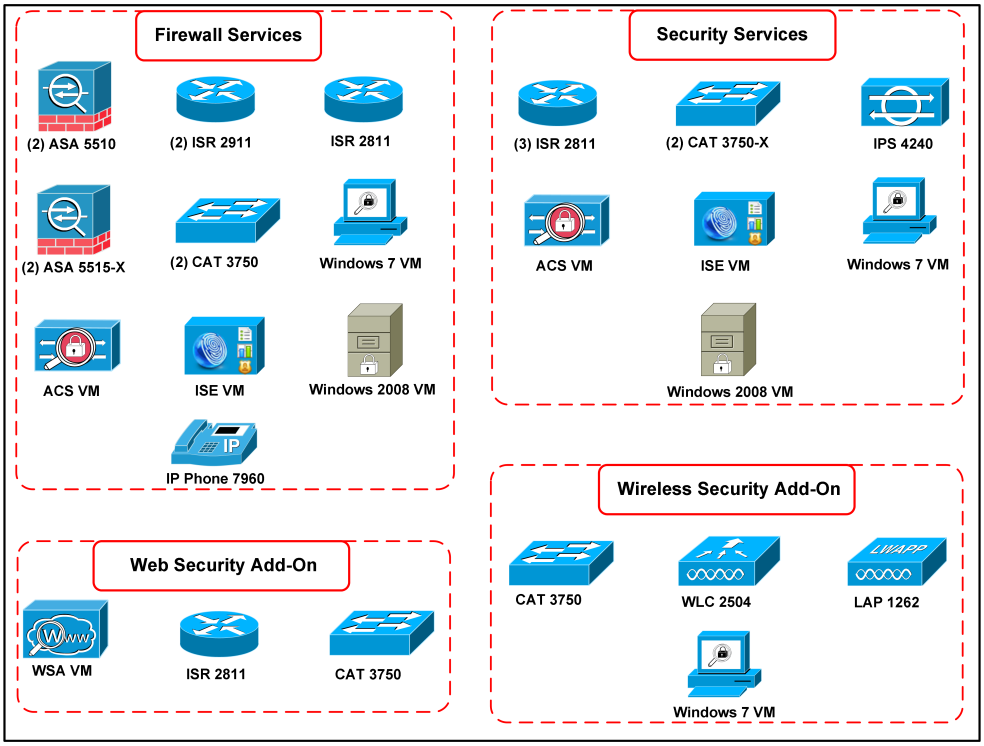
Initial configuration files are provided at the top of this page in the **Resources** panel. Each lab has an initial configuration that is assumed. The initial configuration in some cases is the configuration that the previous lab ended with. If there is only one configuration file, the other devices did not change in the previous lab.

Load the Section 6 Initial Configuration Files to initialize your rack. Pay attention to the instructions in each task, because you may be required to load a different set of initial configuration files.

Section 6 General Lab Topology

The INE CCIE Security Racks are divided into modules as seen in the figure below. This section, Section 6, utilizes both Firewall Services and Security Services.

https://t.me/learningnets



In this section, you will focus on how to provide various security services using Cisco IOS routers as a perimeter device. As you work on each task, focus on the small elements of what's happening with the technology. Gaining an in-depth knowledge of the workings of small tasks will lead to better troubleshooting of large scenarios in the CCIE Lab exam as well as in production environments. Pay close attention to the topology, which may change.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124

https://t.me/learningnets

- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Basic Access-Lists

Task

- Configure IPv4 ingress and egress ACL filtering on R1's VLAN 123 interface.
- Egress filtering should restrict traffic from VLAN 128 as follows:
  - Permit HTTP/HTTPS traffic only during business hours (Monday-Friday 08:00--17:00).
  - Permit remote access via standard virtual terminal protocols.
  - Permit file transfer via active FTP.
  - Permit ping and unix-style traceroute.
- Ingress filtering should restrict traffic destined to VLAN 128 as follows:
  - Permit DNS query and DNS zone transfer.
  - Permit email access for both secure and non-secure versions of the commonly used three protocols using the minimum number of ACL entries.
- Path MTU Discovery should work bidirectionally.
- R1 should restrict IPv6 traffic inbound on its VLAN 123 interface:
  - Deny all telnet traffic.
  - Deny all packets with the hop-by-hop extension header.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.

Overview

The main problem with IPv4/IPv6 basic access-lists is that they are not stateful and have no idea of session tracking. Thus, as in our example where we have both ingress and egress filtering for an interface, if a packet is allowed by the egress access-list, we must ensure that return traffic is allowed by the ingress access-list

https://t.me/learningnets

and vice-versa. Although it is a good idea to memorize the common protocol port numbers, a good list of them can be found in the ASA Command Line Configuration Guide, in the **Reference** section, and the sub-section **Addresses, Protocols and Ports**. A quicker method to find the ports used by most protocols is to check the IOS PAM (Port to Application Mapping) database with the command `show ip port-map .`

FTP protocol has two functional modes, Active and Passive. FTP Active mode is when the client connects to the server on port 21 (control connection), and the server opens a secondary (data connection) session sourced from port 20 back to the client. FTP Passive is when the client connects to the server on port 21 (control connection), the server tells the client the port number for the secondary (data connection) session, and the client initiates data connection on that port.

Common UNIX and IOS traceroute implementation sends out UDP packets destined toward a well-known port range of **33434-33464** by default, and expects two types of ICMP messages in reply: **Time-Exceeded** and **Port-Unreachable**. Microsoft implementation of traceroute uses ICMP Echo-Request packets, so the reply can be one of two ICMP messages: **Time-Exceeded** or **Echo-Reply**. Additionally, PMTU discovery works by sending packets at their maximum MTU size, which defaults to 1500 bytes, with the DF (Don't Fragment) bit set to 1. Whenever a transit device would need to fragment the packet due to adding additional overhead (for example, when doing encryption), it will not be allowed to and will send back to the source with an ICMP message type 3 with code 4 **Packet too Big**.

DNS query (for domain name resolution requests) uses UDP port 53, whereas DNS server zone transfer (when there is a database transfer between two DNS servers responsible for a domain) uses TCP port 53.

Remember that router-generated traffic is not subject to checks by egress access-lists; it is simply permitted. However, router-destined traffic is subject to ingress access-list, so make sure to allow control-plane traffic, such as EIGRP in our case.

By default, the IP protocol allows packet fragmenting. This feature has been a long-known weakness, exploited by various types of attacks (such as ping of death). Overlapping fragments, fragments exceeding the assembly buffer, and fragments arriving out of order are just a few examples of traffic that can severely degrade end system performance when an attacker sends them at high rates. Additionally, fragmented packets are often used to bypass IDS or firewall systems. The attacker splits a packets in such a way that the firewall or IDS system is not able to extract information about something like the port numbers. Advanced firewalls and IDS systems support packet stream reassembly, but this procedure may degrade overall security system performance. For these reasons, it is a good practice to avoid traffic fragmentation in your network and protect your servers against fragmented packets. One solution is to configure matching MTU values and enable the PMTU Discovery

https://t.me/learningnets

process. The IOS router can classify IP packets as one of the following:

- Non-fragmented packets or initial fragments. These packets have a fragment offset of zero and commonly contain some upper-level protocol payload (TCP) to allow the extraction of information about application ports.
- Non-initial fragments have a non-zero fragment offset and are the remaining parts of a fragmented packet. These packets do not have upper-level protocol port information, and the IOS router cannot match them against ACL entries configured with TCP/UDP port numbers. Instead of matching the port numbers, the firewall uses only Layer 3 information in an ACL entry (such as source/destination IP addresses) to match against the source/destination IP addresses in the packet.

You can configure an access-list entry with the log keyword, making it log the matching packets via syslog. Using the `log-input` keyword, you can also log the input interface and the source MAC address or the DLCI if frame-relay is used. There are two global configuration options to be used with access-list-based logging. The first one is the logging update-threshold. It specifies how many times a packet must hit the access-list entry to generate a logging message. By default, this value is zero, which means that the router would generate an entry based on a periodic timeout of 5 minutes. This command serves the purpose of aggregating the hits on the entry, and the format is `ip access-list log-update threshold <HIT-COUNT>` for IPv4 or `ipv6 access-list log-update threshold <HIT-COUNT>` for IPv6. When this value is non-zero, the router generates a log entry for every number of hits and also produces the periodic 5-minute logging message.

The second option relates to the behavior of packet logging. You should remember that every logged packet is process-switched, so a large packet flow may easily consume all of your CPU resources. For this reason, you may want to rate-limit the amount of process-switched packets by using the command `ip access-list logging interval <TIME-INTERVAL>` . This command allows only one packet per interval to be process-switched. All exceeding packets are not process-switched and therefore are not accounted for logging purposes. By default, this feature is off and all packets are process-switched. Note that this interval does not apply to packets destined to the router itself, because these are process-switched by default. Using this command limits the effect of packet logging on the CPU, but it may result in an unpredictable number of packets being logged. In reality, it is useful when you just need a general hint of packet matches, not the detailed statistics. In general, adding an explicit `deny ip any any log` at the end of your access-lists may assist in troubleshooting during your lab exam.

Object-groups can be used in IOS routers to optimize the number of IPv4 ACL entries, but only network and service type object-groups are supported. Additionally, support is only for extended (named or numbered) access-lists. Compared with the

https://t.me/learningnets

ASA, the disadvantage of using it is that you can't have a view of the expanded access-list. So in ASA, issuing `show running-config access-list` would show the ACL content with object-groups, whereas issuing `show access-list` would show the expanded ACL content. When using the **service** object-group, command syntax is different than on ASA, because the service object is referenced instead of the protocol field in the ACL, and if TCP/UDP port numbers are also specified, they will be used as the destination ports, not as the source ports. Basically, you cannot match on the TCP/UDP source port of the traffic when using service object-groups inside an ACL entry.

There are several IPv6 extension headers used by the protocol, each with its own specific role. One that can be matched within an IPv6 access-list is the **Hop-by-Hop Options**, which is used to signal each router in the transit path to process the packet additionally to forwarding it, so the packet will be process-switched. For example, IPv6 multicast uses this extension header in its MLD messages to force each transit router process MLD messages sent to multicast addresses for which the router is not a group member.

Configuration

R1:

```
time-range HTTP
 periodic weekdays 08:00 to 17:00
!
object-group network VLAN128
 136.1.128.0 /24
!
object-group service TCP_VLAN128
 tcp eq 25
 tcp eq 110
 tcp eq 995
 tcp eq 143
 tcp eq 220
 tcp eq 993
!
ip access-list extended EGRESS_VLAN123
 permit tcp 136.1.128.0 0.0.0.255 any eq 80 time-range HTTP
 permit tcp 136.1.128.0 0.0.0.255 any eq 443 time-range HTTP
 permit tcp 136.1.128.0 0.0.0.255 any range 22 23
 permit tcp 136.1.128.0 0.0.0.255 any range 20 21
 permit icmp 136.1.128.0 0.0.0.255 any echo
 permit udp 136.1.128.0 0.0.0.255 any range 33434 33464
!
```

https://t.me/learningnets

```
 permit udp 136.1.128.0 0.0.0.255 eq 53 any
 permit tcp 136.1.128.0 0.0.0.255 eq 25 any
 permit tcp 136.1.128.0 0.0.0.255 eq 110 any
 permit tcp 136.1.128.0 0.0.0.255 eq 995 any
 permit tcp 136.1.128.0 0.0.0.255 eq 143 any
 permit tcp 136.1.128.0 0.0.0.255 eq 220 any
 permit tcp 136.1.128.0 0.0.0.255 eq 993 any
!
 permit icmp any any packet-too-big
 deny ip 136.1.128.0 0.0.0.255 any
 permit ip any any
!
!
ip access-list extended INGRESS_VLAN123
 permit eigrp any any
 permit udp any 136.1.128.0 0.0.0.255 eq 53
 permit tcp any 136.1.128.0 0.0.0.255 eq 53
 permit object-group TCP_VLAN128 any object-group VLAN128
!
 permit tcp any eq 80 136.1.128.0 0.0.0.255 time-range HTTP
 permit tcp any eq 443 136.1.128.0 0.0.0.255 time-range HTTP
 permit tcp any range 22 23 136.1.128.0 0.0.0.255
 permit tcp any range 20 21 136.1.128.0 0.0.0.255
 permit icmp any 136.1.128.0 0.0.0.255 echo-reply
 permit icmp any 136.1.128.0 0.0.0.255 time-exceeded
 permit icmp any 136.1.128.0 0.0.0.255 port-unreachable
!
 permit icmp any any packet-too-big
 deny ip any 136.1.128.0 0.0.0.255
 permit ip any any
!
!
ipv6 access-list DENY_TELNET
 deny tcp ::/0 ::/0 eq 23
 deny tcp ::/0 eq 23 ::/0
 deny ipv6 any any hbh
 permit ipv6 ::/0 ::/0
!
!
interface GigabitEthernet0/1.123
 ip access-group INGRESS_VLAN123 in
 ip access-group EGRESS_VLAN123 out
 ipv6 traffic-filter DENY_TELNET in
```

https://t.me/learningnets

Verification

Verify the configured object-groups and that filters are correctly applied:

```
R1#show object-group
Service object group TCP_VLAN128
 tcp eq smtp
 tcp eq pop3
 tcp eq 995
 tcp eq 143
 tcp eq 220
 tcp eq 993

Network object group VLAN128
 136.1.128.0 255.255.255.0
!
!R1#show ip interface gigabitEthernet0/1.123 | i VLAN
 Outgoing access list is EGRESS_VLAN123
 Inbound  access list is INGRESS_VLAN123
!
!R1#show ipv6 interface gigabitEthernet 0/1.123 | i TELNET
 Inbound access list DENY_TELNET
```

Identify ports for the most common email protocols, IMAP/POP3/SMTP, using PAM database:

```
R2#show ip port-map | i smtp|imap|pop3

Default mapping:  imap                tcp port 143                    system defined
Default mapping:  smtp                tcp port 25                     system defined
Default mapping:  imap3               tcp port 220                    system defined
Default mapping:  pop3s               tcp port 995                    system defined
Default mapping:  pop3                tcp port 110                    system defined
Default mapping:  imaps               tcp port 993                    system defined
```

Verify PING, telnet, and traceroute services from VLAN 128. For example:

```
SW1#ping 150.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:!!!!!
```

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
!
!SW1#telnet 150.1.3.3
Trying 150.1.3.3 ... Open



User Access Verification

Password: R3>
!
!SW1#traceroute 150.1.3.3

Type escape sequence to abort.
Tracing the route to 150.1.3.3

  1 136.1.128.1 0 msec 0 msec 0 msec 2 136.1.99.3 9 msec *  0 msec
```

Check the system clock. If outside of the configured time-range, HTTP/HTTPS traffic from VLAN 128 should be dropped. If sourced from Loopback0, subnet should be allowed:

https://t.me/learningnets

```
R1#show time-range
time-range entry: HTTP (inactive)

   periodic weekdays 8:00 to 17:00
   used in: IP ACL entry
   used in: IP ACL entry
   used in: IP ACL entry
   used in: IP ACL entry
!
!SW1#telnet 150.1.3.3 80

Trying 150.1.3.3, 80 ... % Destination unreachable; gateway or host down

!
!SW1#telnet 150.1.3.3 80 /source-interface loopback0

Trying 150.1.3.3, 80 ... Open
?HTTP/1.1 400 Bad Request

Date: Sat, 12 Oct 2013 08:00:19 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.3.3 closed by foreign host]
```

Verify that ICMP traffic toward VLAN 128 is dropped but traffic toward SW1's Loopback0 is allowed:

https://t.me/learningnets

```
R4#ping 136.1.128.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.128.9, timeout is 2 seconds:U.U.U

Success rate is 0 percent (0/5)
!
!R4#ping 150.1.9.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.9.9, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
```

Verify that email traffic toward VLAN 128 is allowed (for example, SMTP and POP3), but that traffic for other protocols not allowed by the ACL is dropped. The **% Connection refused by remote host** message means that the target host refused the connection because it does not have the requested service active; the **% Destination unreachable; gateway or host down** message means that a transit device dropped the packet and sent back an ICMP message; the most common reasons for this are an ACL filter or lack of routing information:

```
R4#telnet 136.1.128.9 25
Trying 136.1.128.9, 25 ... % Connection refused by remote host
!
!R4#telnet 136.1.128.9 110
Trying 136.1.128.9, 110 ... % Connection refused by remote host
!
!R4#telnet 136.1.128.9 21
Trying 136.1.128.9, 21 ... % Destination unreachable; gateway or host down
```

Verify that IPv6 telnet traffic is bidirectionally dropped:

https://t.me/learningnets

```
R4#telnet 150::1:1:1

Trying 150::1:1:1 ... % Destination unreachable; gateway or host down

!
!R1#telnet 150::1:4:4

Trying 150::1:4:4 ... % Connection timed out; remote host not responding
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Dynamic Access-Lists

Task

- Configure R1 for dynamic access-lists inbound on its VLAN 128 interface:
  - TCP traffic should be allowed only after successful authentication.
  - Each IP address should be authenticated for TCP traffic.
  - Authenticate with a local username/password pair of **CISCO/CISCO**.
  - Timeout the dynamic entries after 1 minute of inactivity and allow it for a maximum of 5 minutes.
  - Allow the user to extend the session before the 5-minute window expires.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Dynamic access-lists is a type of access policy that is activated by a user logging in to the router. The important factor here is the `access-enable` exec command, which activates all dynamic access entries in all access-lists. This command may be assigned to a particular user's profile, or it could be attached to a virtual terminal line, affecting all authenticated users. The user must be authorized for **exec** access after successful authentication for the `access-enable` command to be automatically executed. Because of the autocommand option, upon authentication, the user will not be allowed access to manage the router and be disconnected, and the dynamic

https://t.me/learningnets

access-list entries will be activated; however, if the user's authorization cannot execute the `access-enable` command, the user will be granted access to the router's VTY lines. You may use two optional arguments along with the `access-enable` command:

- The `host` keyword, which creates dynamic ACL entry ONLY for a host that triggered the authentication session.
- The `timeout` keyword, which specifies inactivity timeout to remove the dynamic ACL entry.

While creating a dynamic template entry in an access-list, you may also specify absolute timeout with the `timeout` keyword; the idle timeout should be less than the absolute timeout. The global command `access-list dynamic-extend` extends the absolute timeout by 6 minutes each time the user authenticates again via telnet on the router. So for example, if you've authenticated once, the dynamic entry will be available for the absolute timeout value, but if you reauthenticate before the absolute timeout expires, the absolute timeout value is extended by 6 minutes. Both idle and absolute timeout values are configured in minutes.

Note that dynamic or lock-and-key access-lists are not a form of stateful firewall, but represent a basic form of authenticating users before allowing access to specific resources. It is similar to authentication-proxy functionality; however:

- Dynamic ACL can only be triggered by a telnet connection to the router; with authentication proxy, the router can transparently intercept telnet/HTTP/FTP traffic and request for authentication.
- Both can use the local database or an external RADIUS/TACACS server for authentication and authorization.
- Authentication proxy can handle the addition of multiple dynamic access-lists, whereas lock and key can only add one.
- Authentication proxy has an absolute timeout but no idle timeout; lock and key has both.

Dynamic access-lists have no support for IPv6 traffic.

Configuration

R1:

```
aaa new-model
aaa authentication login VTY local
aaa authorization exec VTY local
```

https://t.me/learningnets

```
!
username CISCO password CISCO
username CISCO autocommand access-enable host timeout 1
!
ip access-list extended INBOUND_VLAN128
 dynamic DACL timeout 5 permit tcp any any
 permit tcp any host 150.1.1.1 eq 23
 deny tcp any any
 permit ip any any
!
interface GigabitEthernet0/1.128
 ip access-group INBOUND_VLAN128 in
!
line vty 0 15
 login authentication VTY
 authorization exec VTY
!
access-list dynamic-extended
```

Verification

Verify that TCP traffic is restricted before authentication:

```
SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... % Destination unreachable; gateway or host down
!
!SW1#ping 150.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
```

Authenticate to R1 via telnet, and test again to see if TCP traffic is now allowed:

```
SW1#telnet 150.1.1.1
Trying 150.1.1.1 ... Open



User Access Verification

Username: CISCO
Password:
[Connection to 150.1.1.1 closed by foreign host]
```

https://t.me/learningnets

```
!
!SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open



User Access Verification

Password: R4>
!
!R1#show ip access-lists
Extended IP access list INBOUND_VLAN128
    10 Dynamic DACL permit tcp any any permit tcp host 136.1.128.9 any (17 matches) (time left 50)

    20 permit tcp any host 150.1.1.1 eq telnet (79 matches)
    30 deny tcp any any (3 matches)
    40 permit ip any any (218 matches)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Reflexive Access-Lists

Task

- Using reflexive access-lists:
  - Configure R1 to permit IPv4 TCP, UDP, and ICMP traffic from VLAN 128, outbound on its VLAN 123 interface.
  - Configure R1 to permit IPv6 TCP and ICMP traffic outbound on its VLAN 123 interface.
  - Dynamic ACL entries should be cleared after 10 minutes of inactivity, except for ICMP, which should allow a 10-second window of no traffic.
  - Drop all other traffic.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Reflexive access-lists represent a basic form of stateful firewall, in that a router can track initiated TCP/UDP/ICMP session flows and allow return traffic automatically by inserting dynamic entries in the access-list applied in the opposite direction. To implement traffic reflection, a named access-list is required. This feature only works for basic protocols such as SSH, telnet, and HTTP, because session tracking is purely based on layer3/layer4 information. It does not support application inspection, and for this reason, protocols negotiating secondary sessions such as FTP/TFTP

https://t.me/learningnets

are not supported.

A dynamically created entry will be expired after an idle timeout, if no traffic is seen for a certain interval. The idle timeout can be globally configured using the command `ip reflexive-list timeout <seconds>` or specific per access-list entry using the keyword `timeout <seconds>` . Additionally, for TCP traffic, if the session is closed gracefully by both the client and the server, the router will understand that the session is over and clear the dynamic entry without waiting for the idle timeout. For UDP and ICMP traffic, if there is no real layer4 information for session state, it waits for the idle timeout to close the dynamic entry.

Reflexive ACLs have poor scalability, because the router opens the pinholes temporarily and needs to age them out, constantly keeping track of every new session. Remember that router-originated traffic is not subject to reflection unless you use local policy routing to divert it into a loopback interface, therefore making the router interpret the traffic as ingress.

Reflexive ACLs are supported for both IPv4 and IPv6 traffic, with the same options and functionality. IPv6 no longer uses ARP for IP-to-MAC address resolution; instead, it uses ICMPv6 with two dedicated types: 135 for neighbor solicitation and 136 for neighbor advertisement. Neighbor solicitation replaces ARP Request and uses the IPv6 solicited-node multicast address, whereas Neighbor advertisement replaces ARP Reply. For this reason, make sure to allow required ICMPv6 types in the inbound access-list; otherwise, when the IPv6 neighbor table becomes stale, IP-to-MAC resolution will no longer work and link-layer packet encapsulation will fail.

Configuration Option 1

R1:

```
ip access-list extended IPV4_OUTBOUND_VLAN123
 permit tcp 136.1.128.0 0.0.0.255 any reflect IPV4_MIRROR
 permit udp 136.1.128.0 0.0.0.255 any reflect IPV4_MIRROR
 permit icmp 136.1.128.0 0.0.0.255 any reflect IPV4_MIRROR timeout 10
!
ipv6 access-list IPV6_OUTBOUND_VLAN123
 permit tcp ::/0 ::/0 reflect IPV6_MIRROR
 permit icmp ::/0 ::/0 reflect IPV6_MIRROR timeout 10
!
ip access-list extended IPV4_INBOUND_VLAN123
 permit eigrp any any
 evaluate IPV4_MIRROR
!
ipv6 access-list IPV6_INBOUND_VLAN123
```

https://t.me/learningnets

```
 permit 88 any any
 permit icmp any any nd-ns
 permit icmp any any nd-na
 evaluate IPV6_MIRROR
!
interface GigabitEthernet0/1.123
 ip access-group IPV4_OUTBOUND_VLAN123 out
 ip access-group IPV4_INBOUND_VLAN123 in
 ipv6 traffic-filter IPV6_OUTBOUND_VLAN123 out
 ipv6 traffic-filter IPV6_INBOUND_VLAN123 in
!
ip reflexive-list timeout 600
```

Configuration Option 2

If the task had not required you to do the filtering outbound on R1's VLAN123 interface, the following alternative configuration could have been used:

R1:

```
ip access-list extended IPV4_INBOUND_VLAN128
 permit eigrp any any
 permit tcp 136.1.128.0 0.0.0.255 any reflect IPV4_MIRROR
 permit udp 136.1.128.0 0.0.0.255 any reflect IPV4_MIRROR
 permit icmp 136.1.128.0 0.0.0.255 any reflect IPV4_MIRROR timeout 10
!
ip access-list extended IPV4_INBOUND_VLAN123
 permit eigrp any any
 evaluate IPV4_MIRROR
!
ipv6 access-list IPV6_INBOUND_VLAN128
 permit 88 any any
 permit icmp any any nd-ns
 permit icmp any any nd-na
 permit tcp ::/0 ::/0 reflect IPV6_MIRROR
 permit icmp ::/0 ::/0 reflect IPV6_MIRROR timeout 10
!
ipv6 access-list IPV6_INBOUND_VLAN123
 permit 88 any any
 permit icmp any any nd-ns
 permit icmp any any nd-na
 evaluate IPV6_MIRROR
!
interface GigabitEthernet0/1.123
```

https://t.me/learningnets

```
 ip access-group IPV4_INBOUND_VLAN123 in
 ipv6 traffic-filter IPV6_INBOUND_VLAN123 in
!
interface GigabitEthernet0/1.128
 ip access-group IPV4_INBOUND_VLAN128 in
 ipv6 traffic-filter IPV6_INBOUND_VLAN128 in
!
ip reflexive-list timeout 600
```

Verification

Verify that IPv4 ICMP and TCP traffic from VLAN 128 works, and that if sourced from SW1's Loopback0 interface it does not:

```
SW1#ping 150.1.4.4 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 U.U.U
Success rate is 0 percent (0/5)
!
!SW1#ping 150.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
!
!SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

Password: R4>
!
!R1#show ip access-lists IPV4_MIRROR
Reflexive IP access list IPV4_MIRROR
 permit icmp host 150.1.4.4 host 136.1.128.9  (10 matches) (time left 8)
 permit tcp host 150.1.4.4 eq telnet host 136.1.128.9 eq 11625 (30 matches) (time left 592)
```

https://t.me/learningnets

For IPv6 testing, SW1's default SDM must be changed using the command `sdm prefer dual-ipv4-and-ipv6 routing` , and IPv6 EIGRP must be configured on SW1.

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

CBAC Routed Firewall

Task

- Configure CBAC on R1 according to the following requirements:
  - Inspect IPv4/IPv6 TCP, UDP, and ICMP for both transit and router-generated traffic.
  - Ensure that FTP sessions can be successfully established.
  - Apply the policy outbound on R1's VLAN 123 interface.
  - Drop any traffic received inbound on VLAN 123 interface if it doesn't exist in the state table.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

CBAC (context-based access-control) introduces true stateful application inspection to IOS code, being the successor to reflexive access-lists. Compared with the reflexive ACL feature, it is an advanced inspection technology. The general idea is to inspect protocol-specific information in traffic flows going across the router and dynamically open holes in access-lists for returning traffic. It supports both IPv4 and IPv6, but true application inspection capabilities for IPv6 is limited and basically supports only FTP. CBAC configurations commonly consist of two parts:

https://t.me/learningnets

- The CBAC inspection rule (defines traffic classes to inspect) applied to an interface in the direction matching the initial flow of traffic (usually from the trusted network to the untrusted network).
- An access-list applied to an interface in the direction matching the returning traffic flow (usually from the untrusted network to the trusted network).

The inspection rule is globally defined using these commands:

```
ip inspect name <RULE-NAME> <protocol> <options>
ipv6 inspect name <RULE-NAME> <protocol> <options>
```

You apply it to an interface using the following interface-level commands:

```
ip inspect <RULE-NAME> [in|out]
ipv6 inspect <RULE-NAME> [in|out]
```

The rule inspects traffic flowing in the direction specified and installs a special ACL bypass entry in the packet-switching path to allow returning packets to pass through the router, even if there is an ACL denying them (in older IOS releases, CBAC simply added new ACL entries dynamically, but now it's a bypass rule that performs more efficiently). Therefore, even if there is an inbound ACL in the path of returning packets, it will be bypassed by packets identified as part of the returning protocol traffic flow. For example, with the FTP protocol in active mode, CBAC dynamically permits the data transfer connection (which is in the reverse direction), even without knowing the communication port in advance.

The inspection rule and the access-list can be assigned to different interfaces; as long as they apply to the initial connections and returning traffic flows respectively, it's fine. For example, you can apply the inspection rule as outgoing on the outside interface and an access-list in the incoming direction. You can also put the inspection rule in the incoming direction on the inside interface and the access-list in the incoming direction on the outside interface.

CBAC protocol inspection can provide you with alarms, notifying you of protocol violations and potential attacks. By default, alerts are enabled for all inspected protocols. You can disable alarms globally and configure them selectively, per protocol. In addition to providing alarm messages, CBAC can perform basic traffic accounting (duration, bytes transferred) by logging audit trails in syslog. This feature is disabled by default, and it could be enabled globally for all inspected protocols or tuned selectively per protocol in an inspection rule.

CBAC may inspect generic TCP/UDP connections (without looking for specific

https://t.me/learningnets

upper-level protocol information) just to check their integrity and open a hole for returning traffic. If you enable TCP inspection along with HTTP inspection at the same time, HTTP inspection will be used for connections made on TCP port 80. By default, CBAC (as well as outgoing ACLs) does not apply to router-generated traffic. You must add manual ACL entries to permit returning traffic for router-generated packets. Alternatively, you can use local policy routing or a special CBAC feature for router-generated traffic inspection that only works for IPv4 packets.

CBAC has several configurable inspection timeouts. For every inspected protocol, you can set up an inactivity timeout and CBAC will close the inspection entry after the timeout expires. There is a special command, **`ip inspect dns-timeout <TIMEOUT>`**, which applies when you inspect generic UDP sessions. Because DNS represents a special case for the UDP protocol, it has a separate timeout. DNS sessions are quick request/response, so by default when CBAC sees a UDP packet on port 53, it holds the inspection entry for 5 seconds, not the default interval specified for the UDP protocol. In current releases of IOS, you can define separate DNS inspection rules with their own timeouts.

Configuration

R1:

```
ip inspect name IPV4_CBAC udp router-traffic
ip inspect name IPV4_CBAC tcp router-traffic
ip inspect name IPV4_CBAC icmp router-traffic
ip inspect name IPV4_CBAC ftp
!
ipv6 inspect name IPV6_CBAC udp
ipv6 inspect name IPV6_CBAC tcp
ipv6 inspect name IPV6_CBAC icmp
ipv6 inspect name IPV6_CBAC ftp
!
ip access-list extended IPV4_INBOUND
 permit eigrp any any
 deny ip any any
!
ipv6 access-list IPV6_INBOUND
 permit 88 any any
 permit icmp any any nd-ns
 permit icmp any any nd-na
 deny ipv6 any any
!
interface GigabitEthernet0/1.123
 ip inspect IPV4_CBAC out
```

https://t.me/learningnets

```
 ipv6 inspect IPV6_CBAC out
 ip access-group IPV4_INBOUND in
 ipv6 traffic-filter IPV6_INBOUND in
```

Verification

Test IPv4 TCP/ICMP traffic from behind R1:

```
SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open



User Access Verification

Password: R4>
!
!SW1#ping 150.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
!
!R1#show ip inspect sessions  detail
Established Sessions Session 2C86649C (136.1.128.9:8)=>(150.1.4.4:0) icmp SIS_OPEN
  Created 00:00:08, Last heard 00:00:08
   ECHO request
  Bytes sent (initiator:responder) [360:360]
  In  SID 150.1.4.4[0:0]=>136.1.128.9[0:0] on ACL IPV4_INBOUND  (5 matches)
  In  SID 0.0.0.0[0:0]=>136.1.128.9[3:3] on ACL IPV4_INBOUND
  In  SID 0.0.0.0[0:0]=>136.1.128.9[11:11] on ACL IPV4_INBOUND
Session 2C866274 (136.1.128.9:54646)=>(150.1.4.4:23) tcp SIS_OPEN

  Created 00:00:28, Last heard 00:00:12
  Bytes sent (initiator:responder) [40:75]
  In  SID 150.1.4.4[23:23]=>136.1.128.9[54646:54646] on ACL IPV4_INBOUND  (15 matches)
```

Verify that IPv4 router-initiated traffic is subject to CBAC inspection engine:

```
R1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open



User Access Verification
```

https://t.me/learningnets

```
Password: R4>

!
!R1#ping 150.1.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#show ip inspect sessions  detail
Established Sessions Session 2C86649C (136.1.99.1:37446)=>(150.1.4.4:23) tcp SIS_OPEN
  Created 00:00:15, Last heard 00:00:11
  Bytes sent (initiator:responder) [31:71]
  In  SID 150.1.4.4[23:23]=>136.1.99.1[37446:37446] on ACL IPV4_INBOUND  (24 matches)
Session 2C866274 (136.1.99.1:8)=>(150.1.4.4:0) icmp SIS_OPEN

  Created 00:00:03, Last heard 00:00:03
   ECHO request
  Bytes sent (initiator:responder) [360:360]
  In  SID 150.1.4.4[0:0]=>136.1.99.1[0:0] on ACL IPV4_INBOUND  (10 matches)
  In  SID 0.0.0.0[0:0]=>136.1.99.1[3:3] on ACL IPV4_INBOUND
  In  SID 0.0.0.0[0:0]=>136.1.99.1[11:11] on ACL IPV4_INBOUND
```

Change SW1's default SDM to use dual stack with the command **`sdm prefer dual-ipv4-and-ipv6 routing`**. Save the configuration and reload, and then enable IPv6 routing and IPv6 EIGRP on SW1:

```
SW1:

ipv6 unicast-routing
!
interface Vlan128
 ipv6 enable
 ipv6 eigrp 111
!
interface Loopback0
 ipv6 address 150::1:9:9/128
 ipv6 eigrp 111
!
ipv6 router eigrp 111
 no shutdown
```

Test IPv6 traffic from behind R1:

```
SW1#telnet 150::1:4:4
```

https://t.me/learningnets

```
Trying 150::1:4:4 ... Open


User Access Verification

Password: R4>
SW1#ping 150::1:4:4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150::1:4:4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
!
!R1#show ipv6 inspect sessions detail
Established Sessions Session 30058214 (150::1:9:9:21872)=>(150::1:4:4:23) tcp SIS_OPEN
   Created 00:00:27, Last heard 00:00:21
  Bytes sent (initiator:responder) [45:80]
  In  SID 150::1:4:4[23:23]=>150::1:9:9[21872:21872] on ACL IPV6_INBOUND  (18 matches)
Session 30058594 (150::1:9:9:0)=>(150::1:4:4:0) icmp SIS_OPEN

   Created 00:00:02, Last heard 00:00:02
   Destinations: 1
       Dest addr [150::1:4:4]
  Bytes sent (initiator:responder) [260:260]
  In  SID 150::1:4:4[0:0]=>150::1:9:9[129:129] on ACL IPV6_INBOUND  (5 matches)
  In  SID 150::1:4:4[0:0]=>150::1:9:9[2:2] on ACL IPV6_INBOUND
  In  SID 150::1:4:4[0:0]=>150::1:9:9[1:1] on ACL IPV6_INBOUND
  In  SID 150::1:4:4[0:0]=>150::1:9:9[3:3] on ACL IPV6_INBOUND
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

CBAC Transparent Firewall

Task

- Configure CBAC on R3 according to the following requirements:
  - Inspect IPv4 TCP and ICMP traffic bridged between VLANs 123 and 34.
  - Inspect IPv4 TCP and ICMP traffic routed between VLANs 34 and 27.
  - Drop all other traffic.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

CBAC transparent firewall was introduced in IOS version 12.3(7)T along with support for IPv6. It allows for traffic filtering and stateful inspection when the router acts as a layer 2 bridge. The benefit of using this operational mode is the ease of deployment; adding the firewall does not require any IP addressing change in the network. It acts as a bump on the wire. In transparent mode, the router does not attempt any routing and simply learns Ethernet MAC addresses on all interfaces, performing Ethernet frame switching based on destination MAC addresses. To configure the IOS firewall as a bridge, you should first define a type of bridging configured in the firewall. This could be either CRB (Concurrent Routing and Bridging) or IRB (Integrated Routing and Bridging). Both use the concept of a **bridge-group**, which is essentially a virtual bridge inside the router, with its own

https://t.me/learningnets

MAC address table:

- CRB mode is enabled using the command `bridge crb` . It allows the router to bridge frames between interfaces configured as members of a single bridge group. Interfaces not in a bridge group with IP addresses assigned may accept and forward IP packets. Essentially, the router acts as a bridge for one set of interfaces and layer3 router for the rest of the interfaces.
- IRB mode is enabled using the command `bridge irb` . It allows configuring a special BVI (Bridge Virtual Interface) for every bridge group, which represents the router as a layer3 device within the respective bridge group. All devices in the group aware of this BVI's MAC address may send IP packets to the router and communicate with other layer2 domains. Essentially, the router acts as both a layer2 bridge and layer3 router for the bridging interfaces, and as a layer3 router for the rest of the interfaces.

You may compare the BVI with an SVI (Switch Virtual Interface) from a layer3 Catalyst switch. In addition to the `bridge irb` command, you must add the command `bridge <number> route ip` ; otherwise, the BVI interface will not route IP packets. To configure an interface/sub-interface as member of a bridge group, use the command `bridge-group <number>` in interface configuration mode. By default, the new group starts with Spanning Tree Protocol disabled, so you need the command `bridge <number> protocol ieee` to enable the IEEE-compatible spanning tree for the bridge-group. Notice that for 802.1Q sub-interfaces, the firewall will run the PVST+ variant of STP. You may apply the classic firewall inspect rules and access-lists to the interface configured with the bridge-group numbers after this. For example:

```
bridge irb
bridge-group 1 protocol ieee
bridge 1 route ip
!
ip inspect name CBAC tcp
ip inspect name CBAC udp
ip inspect name CBAC icmp
!
ip acess-list extended OUTSIDE_IN
 deny ip any any log
!
interface FastEthernet0/0
 bridge-group 1
 ip access-group OUTSIDE_IN in
!
interface FastEthernet0/1
```

https://t.me/learningnets

```
bridge-group 1
ip inspect CBAC in
```

Only TCP/UDP/ICMP flows are supported by layer2 transparent firewall, and non-IP packets are permitted by default across the firewall. To block a non-IP protocol, you must apply a protocol-type access-list to the bridged interface. Those lists have their numbers in the range 201-299 and permit traffic based on Ethertype of SNAP PID value. For example, the following list will permit IP and ARP but block IPX:

```
access-list 201 permit 0x800
access-list 201 permit 0x806
access-list 201 deny 0x8137
access-list 201 permit 0x0 0xFFFF
```

The Ethertype value becoming popular for IPv6 is 0x86DD, and you can use this value to filter it. To associate a protocol-type list with an interface, use the interface-level command **`bridge-group <number> input-type-list <ACL-NUMBER>`** . CBAC inspection is not supported for IPv6 when running transparent mode.

Transparent firewall does not inspect or process multicast Ethernet frames; they are permitted regardless of any input access-lists. However, broadcast frames (such as NetBIOS over TCP/UDP or DHCP) are subject to transparent ACL checks. Transparent firewall provides special handling for DHCP packets because they are important for host auto-configuration. The global command to permit DHCP packets across the firewall regardless of the IP access-lists configuration is **`ip inspect l2-transparent dhcp-passthrough .`**

The IP address assigned to the BVI interface must be on the same subnet as the protected segments. You may treat is as equivalent to the ASA firewall's default management IP address. The traffic originated from the router across the bridged interface is not subject to the checks by the ACL applied on the bridged interfaces. Instead, to control the management access, you may need a separate access-list attached to the BVI.

Although EIGRP multicast hello packets will be allowed by default regardless of any applied access-lists, EIGRP updates and acknowledgments are unicast, so EIGRP traffic must still be allowed by the inbound access-list. For traffic between VLANs 34 and 27, as it will be routed through the BVI interface, inspection is applied at the BVI level, whereas for bridged traffic, inspection is applied at the layer2 interface level. IRB has been pre-configured from the initial configurations, so there is no need to add it.

https://t.me/learningnets

Configuration

```
R3:

ip inspect name CBAC_VLAN123 tcp
ip inspect name CBAC_VLAN123 icmp
!
ip access-list extended INBOUND_VLAN34
 permit eigrp any any
 deny ip any any
!
interface FastEthernet0/1.123
 ip inspect CBAC_VLAN123 in
!
interface FastEthernet0/1.34
 ip access-group INBOUND_VLAN34 in
!
!
ip inspect name CBAC_VLAN34 tcp
ip inspect name CBAC_VLAN34 icmp
!
ip access-list extended INBOUND_VLAN27
 deny ip any any
!
interface BVI99
 ip inspect CBAC_VLAN34 in
!
interface FastEthernet0/0
 ip access-group INBOUND_VLAN27 in
```

Verification

Verify IPv4 traffic received inbound on R3's VLAN 123 interface:

```
R1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

Password: R4>
```

https://t.me/learningnets

```
!
!R1#ping 150.1.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R3#show ip inspect sessions
Established Sessions Session 4A0FFA10 (136.1.99.1:29347)=>(150.1.4.4:23) tcp SIS_OPEN
 Session 4A0FF7F0 (136.1.99.1:8)=>(150.1.4.4:0) icmp SIS_OPEN
```

Verify traffic received inbound on R3's VLAN 34 interface:

```
R3#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.27.54         0148.f8b3.2e24.23       Oct 13 2013 10:37 PM    Automatic
!
!R4#ping 136.1.27.54

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.27.54, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/12 ms
!
!R3#show ip inspect sessions
Established Sessions Session 4A0FF7F0 (136.1.99.4:8)=>(136.1.27.54:0) icmp SIS_OPEN
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

CBAC URL Filtering

Task

- Configure CBAC on R1 as follows:
  - Inspect HTTP traffic inbound on its VLAN 128 interface, and do not allow any sessions to be initiated from the opposite directions.
  - Filter all Java applets except from HTTP servers 150.1.4.4 and 150.1.10.10.
  - HTTP servers on 150.1.4.4 and 150.1.10.10 can be reachable on both TCP 80 and 8080.
  - Configure URL Filtering through a Websense server at 136.1.27.100, and allow all HTTP requests if the Websense server is down.
  - The Websense server allows connection only from IP addresses in the 150.1.0.0/16 range.
  - Whitelist the **cisco.com** domain.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

CBAC supports the following advanced functionalities for HTTP traffic inspection:

- Inspecting HTTP headers for IPv4 traffic to allow or drop Java applets based on a standard access-lists, where permit entries from the ACL are HTTP servers from

https://t.me/learningnets

which Java is allowed.
- Integration with two URL Filtering vendors for IPv4 traffic, Websense and N2H2, and also local URL whitelisting; URL Filtering can be configured for fail-open or fail-closed mode when the URL Filtering server is unreachable.

When HTTP traffic inspection is enabled, java scanning engine is automatically enabled to drop all java applets, which affects the CPU performance. However, if an access-list with the `permit any` statement is used, basically allowing java, the scanning engine is disabled.

CBAC knows which protocols/applications use which ports based on the PAM database, visible with the command `show ip port-map` . You can add custom ports for all PAM known protocols if, for example, you have a service running on a non-standard port, such as FTP running on TCP port 2121 . When defining custom ports for a protocol, an optional standard/extended access-list can be specified to identify the traffic for which this mapping is relevant; with standard access-lists, you define the destination, whereas with extended access-lists, you define the source/destination pattern. For HTTP traffic to be inspected, there is no need to do TCP inspection, so when you configure a layer7/application inspection the transport protocol is automatically included, but only for the respective application. Because of the limited application inspection capabilities of CBAC with IPv6 traffic (basically only FTP being supported), the URL filtering and Java inspections are not available. For IPv6, PAM custom ports can be configured as well.

Configuration

R1:

```
access-list 80 permit host 150.1.4.4
access-list 80 permit host 150.1.10.10
!
access-list 99 permit host 150.1.4.4
access-list 99 permit host 150.1.10.10
access-list 99 deny any
!
ip urlfilter server vendor websense 136.1.27.100
ip urlfilter exclusive-domain permit cisco.com
ip urlfilter source-interface Loopback0
ip urlfilter allow-mode on
!
ip port-map http port tcp 8080 list 80
ip inspect name CBAC http java-list 99 urlfilter
!
```

https://t.me/learningnets

```
ip access-list extended INBOUND_VLAN19
 permit eigrp any any
 deny ip any any
!
ip access-list extended INBOUND_VLAN123
 permit eigrp any any
 permit tcp host 136.1.27.100 eq 15868 host 150.1.1.1
 deny ip any any
!
interface GigabitEthernet0/1.128
 ip inspect CBAC in
!
interface GigabitEthernet0/0
 ip access-group INBOUND_VLAN19 in
!
interface GigabitEthernet0/1.123
 ip access-group INBOUND_VLAN123 in
```

Verification

Activate the HTTP server on R4 to listen on port 8080 by using the command **`ip http port 8080`**. Simulate an HTTP session toward R4's Loopback on port 8080 from SW1:

https://t.me/learningnets

```
SW1#ping 150.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)
!
!SW1#telnet 150.1.4.4 8080

Trying 150.1.4.4, 8080 ... Open
?HTTP/1.1 400 Bad Request

Date: Sat, 12 Oct 2013 23:28:01 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.4.4 closed by foreign host]

!
!R1#show ip inspect sessions

Established Sessions Session 2C86649C (136.1.128.9:60458)=>(150.1.4.4:8080) http SIS_OPEN
```

Verify PAM HTTP custom ports and general CBAC configuration:

```
R1#show ip port-map | i http
Default mapping:  http                tcp port 80                        system defined
Host specific:    http                tcp port 8080          in list 80  user defined
Default mapping:  https               tcp port 443                       system defined
!
!R1#show ip inspect config
Session audit trail is disabled
Session alert is enabled
one-minute (sampling period) thresholds are [unlimited : unlimited] connections
max-incomplete sessions thresholds are [unlimited : unlimited]
max-incomplete tcp connections per host is unlimited. Block-time 0 minute.
tcp synwait-time is 30 sec -- tcp finwait-time is 5 sec
tcp idle-time is 3600 sec -- udp idle-time is 30 sec
tcp reassembly queue length 16; timeout 5 sec; memory-limit 1024 kilo bytes
dns-timeout is 5 sec
Inspection Rule Configuration
```

https://t.me/learningnets

```
Inspection name CBAC
 http java-list 99 url-filter is on alert is on audit-trail is off timeout 3600
!
!R1#show ip inspect interfaces
Interface Configuration Interface GigabitEthernet0/1.128
 Inbound inspection rule is CBAC
 http java-list 99 url-filter is on alert is on audit-trail is off timeout 3600

  Outgoing inspection rule is not set
  Inbound access list is not set
  Outgoing access list is not set
```

Verify the Websense/Urlfiltering configuration; it appears as disabled because there is no Websense server at the IP address of 136.1.27.100, so the router cannot establish the TCP connection with it:

```
R1#show ip urlfilter config
 Websense URL Filtering is DISABLED

 Primary Websense server configurations
 ======================================Websense server IP address Or Host Name: 136.1.27.100
 Websense server port: 15868
 Websense retransmission time out: 6 (in seconds)
 Websense number of retransmission: 2

 Secondary Websense servers configurations
 ==========================================
 Other configurations
 ====================Allow Mode: ON

 System Alert: ENABLED
 Audit Trail: DISABLED
 Log message on Websense server: DISABLED
 Truncate Hostname: DISABLED
 Truncate Script-parameters: DISABLED
 Maximum number of cache entries: 5000
 Maximum number of packet buffers: 200
 Maximum outstanding requests: 1000
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

CBAC Tuning

Task

- Configure R1 to allow IPv4 TCP and ICMP traffic inbound on its VLAN 128 using stateful inspection:
  - Delete half-open sessions when the total number reaches 1000, and stop when it falls below 900.
  - Delete half-open sessions when the one-minute rate reaches 100, and stop when it falls below 90.
  - Enable a per-host limit of half-open sessions to 50 and block for 5 minutes when the threshold is reached.
  - Allow a TCP session to be established within 15 seconds.
- To improve performance, disable CBAC alerts, but retain alerting for ICMP sessions.
- Ensure that CBAC audit is enabled only for ICMP sessions.
- Change the hashtable size to accommodate 4096 sessions for intensive traffic flow.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

By default, with every inspection rule, CBAC enables DoS prevention features very similar to the TCP Intercept feature, and you cannot disable this behavior. CBAC also applies DoS prevention to UDP sessions. A UDP session is considered half-

https://t.me/learningnets

open if no packets have been sent in response to the initial packet. Both UDP and TCP half-open sessions count against the maximum connection thresholds at the same time. Therefore, if you set up a limit of 100 half-open sessions, the router will count both numbers of UDP and TCP sessions against this threshold.

CBAC supports maximum thresholds and per-minute rates for half-open sessions. When the current number of connections exceeds any of the thresholds, the router starts deleting half-open connection entries and sending RST packets to TCP session endpoints. The router does not delete any sessions that have already been established. For TCP sessions, the only supported mode of operations is watch mode; the router does not intervene with TCP sessions, it just watches them as they get established. Offending sessions (sessions that did not complete in time or exceeded thresholds) receive RST messages as well when exceeding configured timeouts. The amount of time CBAC waits for TCP connection establishment is configured using the command `ip inspect tcp synwait-time <seconds>` . Note that this is the equivalent of watch-timeout for the TCP Intercept feature.

One feature special to CBAC is per-host TCP connection limits. CBAC bounds the number of half-open connections per host, in addition to having the global thresholds. As the number of half-open sessions destined to the same IP address exceeds the per-host threshold, the router can block any further connections from establishing for the amount of time specified. The command syntax is `ip inspect tcp max-incomplete host <number> block-time <minutes>` .

The core of the CBAC algorithm is protocol inspection logic and session state table. To make CBAC more effective under heavy load conditions, you should consider the following:

- State table is hash structured, with a configurable number of entries. Try to make the hashtable size close to the average number of concurrent sessions passing through the firewall using command `ip inspect hashtable-size` .
- By default, protocol inspection logic generates alerts when it finds inconsistence in protocol tracking. This may cause additional CPU load under intensive traffic. Consider disabling the alerts globally or per protocol to improve performance, using the command `ip inspect alert-off` .
- You can configure the router to log whenever packets within a session are dropped because of protocol nonconformance by using the global command `ip inspect log drop-pkt` .
- There’s an additional CBAC feature called *session audit*. It allows for logging every session's statistics for accounting or audit purposes. It can be enabled globally or per protocol.

https://t.me/learningnets

The default hash table size is 1024 buckets. It may vary in size from 2048, 4096, or 8192 buckets. When the number of concurrent connections across the firewall exceeds 4000, you may want to adjust the default size to match approximately half of the maximum connections count; for example, with 4500 sessions, the optimal hash table size would be 2048 buckets.

Configuration

R1:

```
ip inspect hashtable-size 2048
no ip inspect audit-trail
ip inspect alert-off
!
ip inspect name CBAC tcp
ip inspect name CBAC icmp audit-trail on
ip inspect name CBAC icmp alert on
!
ip inspect max-incomplete low 900
ip inspect max-incomplete high 1000
ip inspect one-minute low 90
ip inspect one-minute high 100
ip inspect tcp synwait-time 15
ip inspect tcp max-incomplete host 50 block-time 5
!
ip access-list extended INBOUND
 permit eigrp any any
 deny ip any any
!
interface GigabitEthernet0/1.128
 ip inspect CBAC in
!
interface range GigabitEthernet0/0 , GigabitEthernet0/1.123
 ip access-group INBOUND in
```

The task does not require it, but here is an identical configuration for IPv6 traffic; as long as CBAC is used in layer3 mode, the same settings are available for both IPv4 and IPv6 traffic:

R1:

```
ipv6 inspect hashtable-size 2048
```

https://t.me/learningnets

```
no ip inspect audit-trail
ipv6 inspect alert-off
!
ipv6 inspect name CBAC tcp
ipv6 inspect name CBAC icmp audit-trail on
ipv6 inspect name CBAC icmp alert on
!
ipv6 inspect max-incomplete low 900
ipv6 inspect max-incomplete high 1000
ipv6 inspect one-minute low 90
ipv6 inspect one-minute high 100
ipv6 inspect tcp synwait-time 15
ipv6 inspect tcp max-incomplete host 50 block-time 5
!
ipv6 access-list INBOUND
 permit 88 any any
 permit icmp any any nd-ns
 permit icmp any any nd-na
 deny ipv6 any any
!
interface GigabitEthernet0/1.128
 ipv6 inspect CBAC in
!
interface range GigabitEthernet0/0 , GigabitEthernet0/1.123
 ipv6 traffic-filter INBOUND in
```

Verification

Verify the CBAC configuration:

```
R1#show ip inspect all
Session audit trail is disabled
Session alert is disabled
one-minute (sampling period) thresholds are [90 : 100] connections
max-incomplete sessions thresholds are [900 : 1000]
max-incomplete tcp connections per host is 50. Block-time 5 minutes.
tcp synwait-time is 15 sec
 -- tcp finwait-time is 5 sec
tcp idle-time is 3600 sec -- udp idle-time is 30 sec
tcp reassembly queue length 16; timeout 5 sec; memory-limit 1024 kilo bytes
dns-timeout is 5 sec
Inspection Rule Configuration
 Inspection name CBAC tcp alert is off audit-trail is off timeout 3600
 icmp alert is on audit-trail is on timeout 10
```

https://t.me/learningnets

```
Interface Configuration Interface GigabitEthernet0/1.128
 Inbound inspection rule is CBAC

    tcp alert is off audit-trail is off timeout 3600
    icmp alert is on audit-trail is on timeout 10
  Outgoing inspection rule is not set
  Inbound access list is not set
  Outgoing access list is not set
```

Verify that session audit is on only for ICMP traffic:

```
SW1#ping 150.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds: !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
!
!SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

Password:  R4>
!
!R1:
%FW-6-SESS_AUDIT_TRAIL_START: Start icmp session: initiator (136.1.128.9:8) -- responder (150.1.4.4:0)
 %FW-6-SESS_AUDIT_TRAIL: Stop icmp session: initiator (136.1.128.9:8) sent 360 bytes -- responder (150.1.4.4:0) sent
```

Verify the TCP SYN timeout settings:

https://t.me/learningnets

```
R1#debug ip inspect events
 INSPECT special events debugging is on

!
!SW1#telnet 136.1.27.100

Trying 136.1.27.100 ... % Connection timed out; remote host not responding

!
!R1:

FIREWALL sent a TCP pkt (136.1.27.100:23) tcp flag:0x4 -> 136.1.128.9:33590 seq 0 ack 0 wnd 4128, GigabitEthernet0/1
 FIREWALL sent a TCP pkt (136.1.128.9:33590) tcp flag:0x4 -> 136.1.27.100:23 seq 2541624063 ack 0 wnd 0, GigabitEther
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

CBAC High Availability

Task

- Configure R1 and R2 for CBAC stateful traffic inspection and failover inbound on VLAN 128 interface.
- R1 and R2 should synchronize the session states, with R1 as the preferred forwarder at all times.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

CBAC supports a high-availability (HA)configuration involving two routers to provide redundancy. Only one router can be the active forwarder (so traffic load-balancing is not supported) and there is no support for asymmetric routing; traffic needs to flow bidirectionally through the active forwarder. In general, it is your responsibility to properly configure routing to satisfy these requirements. For this scenario, it has been assured from the loaded initial configurations. Like other technologies supported by IOS in HA mode, such as Zone-Based Policy Firewall (ZBPF), IPsec, or PKI, HA uses the Stateful Switch-Over (SSO) redundancy feature. This inter-device redundancy function relies on two protocols:

- **HSRP** determines the router roles: ACTIVE and STANDBY.

https://t.me/learningnets

- **SCTP** ensures automatic synchronization between ACTIVE and STANDBY; for CBAC, this is related to session state replication but only for TCP/UDP sessions; ICMP and layer7/application session states are not synchronized. There is no option for configuration synchronization; the two devices must be manually configured.

Because the order of operations may cause your configuration to be non-functional, the following steps are recommended, in order, for configuring SSO with CBAC:

1. Configure and verify HSRP.
2. Configure and verify SSO.
3. Configure CBAC.
4. Activate and verify CBAC for SSO.

> CBAC HA is only supported for IPv4 traffic, not for IPv6, and only in routed mode.

Configuration

R1:

```
interface GigabitEthernet0/1.128
 standby ip 136.1.128.12
 standby priority 150
 standby preempt
 standby name CBAC
!
ipc zone default
 association 1
  no shutdown
  protocol sctp
  local-port 5000
  local-ip 136.1.128.1
  remote-port 5000
  remote-ip 136.1.128.2
!
redundancy inter-device
 scheme standby CBAC
```

R2:

```
interface GigabitEthernet0/1.128
 standby ip 136.1.128.12
 standby preempt
```

https://t.me/learningnets

```
 standby name CBAC
!
ipc zone default
 association 1
  no shutdown
  protocol sctp
  local-port 5000
  local-ip 136.1.128.2
  remote-port 5000
  remote-ip 136.1.128.1
!
redundancy inter-device
 scheme standby CBAC
```

At this point, you must save the configuration and reload both devices to activate the redundancy. When R2 detects itself as standby, it may force itself another reload. The following output shows the initial required reload:

```
R1#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_INIT
Pending Scheme: Standby (Will not take effect until next reload)
      Pending Groupname: CBAC
  Scheme: <NOT CONFIGURED>
  Peer present: UNKNOWN
  Security: Not configured
!
!R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_INIT
Pending Scheme: Standby (Will not take effect until next reload)

      Pending Groupname: CBAC
  Scheme: <NOT CONFIGURED>
  Peer present: UNKNOWN
  Security: Not configured
```

After SSO is functional, configure CBAC identically on both routers:

```
R1: and R2:

ip inspect name CBAC tcp
ip inspect name CBAC udp
!
ip access-list extended INBOUND
 permit eigrp any any
```

https://t.me/learningnets

```
 deny ip any any
!
interface GigabitEthernet0/1.128
 ip inspect CBAC in redundancy stateful CBAC
!
interface GigabitEthernet0/1.123
 ip access-group INBOUND in
```

Verification

First, verify correct inter-device redundancy functionality:

```
R1#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_ACT
  Scheme: Standby Groupname: CBAC Group State: Active
  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
!
!R1#show redundancy states
my state = 13 -ACTIVE
peer state = 8  -STANDBY HOT
           Mode = Duplex
        Unit ID = 0

     Maintenance Mode = Disabled
    Manual Swact = enabled Communications = Up

   client count = 13
 client_notification_TMR = 60000 milliseconds
           RF debug mask = 0x0
!
!R2#show redundancy inter-device
Redundancy inter-device state: RF_INTERDEV_STATE_STDBY
  Scheme: Standby Groupname: CBAC Group State: Standby
  Peer present: RF_INTERDEV_PEER_COMM
  Security: Not configured
!
!R2#show redundancy states
my state = 8  -STANDBY HOT
peer state = 13 -ACTIVE
           Mode = Duplex
        Unit ID = 0

     Maintenance Mode = Disabled
    Manual Swact = cannot be initiated from this the standby unit
```

https://t.me/learningnets

```
Communications = Up


   client count = 13
 client_notification_TMR = 60000 milliseconds
           RF debug mask = 0x0
```

Initiate a TCP session from behind R1, and verify that session state is replicated from ACTIVE to STANDBY:

```
SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

Password: R4>
!
!R1#show ip inspect sessions detail
Established Sessions Session 31B1E238 (136.1.128.9:27243)=>(150.1.4.4:23) tcp SIS_OPEN
  Created 00:00:16, Last heard 00:00:14
  Bytes sent (initiator:responder) [37:74]
In  SID 150.1.4.4[23:23]=>136.1.128.9[27243:27243] on ACL INBOUND  (14 matches)
 HA state: HA_ACTIVE
!
!R2#show ip inspect sessions detail
Established Sessions Session 2D125CF4 (136.1.128.9:27243)=>(150.1.4.4:23) tcp SIS_OPEN
  Created 00:00:19, Last heard never
  Bytes sent (initiator:responder) [0:0]
In  SID 150.1.4.4[23:23]=>136.1.128.9[27243:27243] on ACL INBOUND
 HA state: HA_STANDBY
```

Enable CBAC application inspection on both R1 and R2 (for example, HTTP), and verify that session state is not replicated between ACTIVE and STANDBY, and that there is no HA state for the session on ACTIVE device:

```
R1: and R2:
ip inspect name CBAC http
!
!SW1#telnet 150.1.4.4 80
Trying 150.1.4.4, 80 ... Open
?
HTTP/1.1 400 Bad Request
Date: Sun, 13 Oct 2013 07:16:45 GMT
```

https://t.me/learningnets

```
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.4.4 closed by foreign host]
!
!R1#show ip inspect sessions detail
Established Sessions Session 31B1E238 (136.1.128.9:36430)=>(150.1.4.4:80) http SIS_OPEN
  Created 00:00:08, Last heard 00:00:08
  Bytes sent (initiator:responder) [0:0]
  In  SID 150.1.4.4[80:80]=>136.1.128.9[36430:36430] on ACL INBOUND  (2 matches)
!
!R2#show ip inspect sessions detail
R2#
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

CBAC Application Inspection

Task

- Configure CBAC inspection for messenger traffic inbound on R1's VLAN 128 interface.
- For Yahoo Messenger:
  - Allow users access only to chat.
  - Inspect traffic toward server 136.1.27.101 on random TCP ports.
  - Log information about session traffic.
- For MSN Messenger:
  - Drop and log traffic toward server 136.1.27.102.
  - Log information about session traffic.
- Messenger traffic tunneled over HTTP should not be allowed.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

CBAC supports advanced layer7 inspection for HTTP and Instant Messenger (IM, MSN, AOL) traffic. By default, all IM client protocols have native TCP ports used for connecting to IM servers. If, however, connection is not successful, it will fail over to HTTP tunneling. HTTP tunneling means that an application uses the default HTTP/HTTPS TCP ports to function. IM protocols can even be manually configured

https://t.me/learningnets

to natively use HTTP tunneling, instead of the default ports. HTTP layer7 inspection can detect if TCP port 80 is used for purely HTTP traffic or by application tunneling and take actions. Additionally, CBAC IM inspection can be configured to allow or drop per application service; for example, it can allow text/chat and block voice/video calling or file transfer.

> CBAC application firewall is not supported for IPv6.

Configuration

```
R1:

appfw policy-name IM_POLICY
 application im yahoo
  server permit ip 136.1.27.101
  service text-chat action allow
  service default action reset
  audit-trail on
!
 application im msn
  server deny ip 136.1.27.102
  audit-trail on
!
 application http
  port-misuse im action reset
!
ip inspect name CBAC appfw IM_POLICY
ip inspect log drop-pkt
!
interface GigabitEthernet0/1.128
 ip inspect CBAC in
```

V Verification

Identify the native ports used by MSN and YMSGR and check CBAC configuration:

```
R1#show ip port-map | i msgr
Default mapping:  ymsgr                tcp port 5050                  system defined
Default mapping:  msnmsgr              tcp port 1863                  system defined
!
!R1#show ip inspect interfaces
```

https://t.me/learningnets

```
Interface Configuration Interface GigabitEthernet0/1.128
Inbound inspection rule is CBAC
    Application Policy name IM_POLICY Application: im yahoo
       service default action: reset
       service text-chat action: allow
       server: permit ip 136.1.27.101
       audit-trail: on Application: im msn
       server: deny ip 136.1.27.102
       audit-trail: on Application http

        port-misuse im action reset
  Outgoing inspection rule is not set
  Inbound access list is not set
  Outgoing access list is not set
```

Simulate Yahoo Messenger traffic on standard port toward any destination; check the session state in CBAC and verify audit logs:

```
SW1#telnet 136.1.27.11 5050
Trying 136.1.27.11, 5050 ... % Connection timed out; remote host not responding
!
!R1:
%FW-6-SESS_AUDIT_TRAIL_START: Start im-yahoo session: initiator (136.1.128.9:47145) -- responder (136.1.27.11:5050)
!
!R1#show ip inspect sessions
Half-open Sessions
 Session 315F71E0 (136.1.128.9:47145)=>(136.1.27.11:5050) im-yahoo SIS_OPENING
!
!
%FW-6-SESS_AUDIT_TRAIL: Stop im-yahoo session: initiator (136.1.128.9:47145) sent 0 bytes -- responder (136.1.27.11:
```

Simulate Yahoo Messenger traffic toward server 136.1.27.101 on non-standard port; check the session state in CBAC and verify audit logs:

```
SW1#telnet 136.1.27.101 57171
Trying 136.1.27.101, 57171 ... % Connection timed out; remote host not responding
!
!R1:
%FW-6-SESS_AUDIT_TRAIL_START: Start im-yahoo session: initiator (136.1.128.9:33326) -- responder (136.1.27.101:57171
!
!R1#show ip inspect sessions
Half-open Sessions
 Session 315F71E0 (136.1.128.9:33326)=>(136.1.27.101:57171) im-yahoo SIS_OPENING
!
```

https://t.me/learningnets

```
!
%FW-6-SESS_AUDIT_TRAIL: Stop im-yahoo session: initiator (136.1.128.9:33326) sent 0 bytes -- responder (136.1.27.101
```

Simulate MSN Messenger traffic toward server 136.1.27.102 on standard port; verify that traffic is dropped and logged:

```
SW1#telnet 136.1.27.102 1863
Trying 136.1.27.102, 1863 ... % Connection timed out; remote host not responding
!
!R1:
%FW-6-DROP_PKT: Dropping im-msn session 136.1.128.9:57749 136.1.27.102:1863    with ip ident 7786 tcpflags 0x6002 se
```

Simulate MSN Messenger traffic toward any server on standard port; check the session state in CBAC and verify audit logs:

```
SW1#telnet 136.1.27.55 1863
Trying 136.1.27.55, 1863 ... % Connection timed out; remote host not responding
!
!R1:
%FW-6-SESS_AUDIT_TRAIL_START: Start im-msn session: initiator (136.1.128.9:13058) -- responder (136.1.27.55:1863)
!
!R1#show ip inspect sessions
Half-open Sessions
 Session 315F71E0 (136.1.128.9:13058)=>(136.1.27.55:1863) im-msn SIS_OPENING
!
!
%FW-6-SESS_AUDIT_TRAIL: Stop im-msn session: initiator (136.1.128.9:13058) sent 0 bytes -- responder (136.1.27.55:18
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Routed Mode

Task

- Configure two security zones on R1, VLAN 19 member of **INSIDE**, VLAN 123 and VLAN 128 member of **OUTSIDE**.
- R1 should be managed only from subnet 136.1.19.0/24 of **INSIDE** zone via telnet and HTTP.
- Users logged in to the router should be allowed to ping any destination.
- Inspect FTP/TFTP/HTTP/TELNET traffic from **INSIDE** to **OUTSIDE**, and ensure that telnet traffic is allowed only when the source and destination match the 150.1.0.0/16 subnet.
- Inspect HTTP/ICMP/SSH traffic within the **OUTSIDE** zone, and ensure that ICMP is not allowed if the source and destination match the 150.1.0.0/16 subnet.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Zone-Based Policy Firewall (ZBPF) is the successor of CBAC for configuring access control in the IOS firewall. Before this feature was introduced, traffic filtering was accomplished using access-lists and stateful traffic inspection rules (CBAC). Both the access-lists and inspection rules apply directly to the physical interfaces, which may poorly reflect the organization's security policy, because the policy deals with

https://t.me/learningnets

more high-level objects than interfaces. A new core concept of ZBPF is zone, which groups different interfaces sharing the same security attributes or the same level of trust. Permissions for traffic forwarding is made between the zones or within a zone, not between physical interfaces. By default, traffic is permitted between the interfaces within the same security zone and blocked between different zones, whereas traffic between an interface configured in a security zone and an interface not in any zone is blocked. CBAC inspection cannot be applied to an interface that is member of a security zone, whereas access-lists can still be applied but it is not recommended; all firewall policies should be done solely at the ZBPF level.

There is one default zone in every router, known as **self**, which encompasses the router's own IP addresses. Traffic to and from this zone to any other zone is permitted by default, to allow for control plane and management plane traffic. However, you may apply an explicit bidirectional policy between the self zone and any other configured zone to control the router-originated or router-destined traffic.

To permit certain traffic flows between two different security zones, you must define a zone pair and apply an inspection policy to this pair. The first zone in the pair is the source zone, and the second is the destination zone. When you assign the inspection policy to a zone pair, it applies to the traffic flowing from source to destination zone. The inspection policy is configured as a set of traffic-classes along with associated actions, such as inspect, pass, or drop. Traffic classes are generally defined using access-lists and/or protocol matching.

Zone-based firewall implements the same set of features as the classic IOS firewall. It allows defining advanced inspection options, such as deep protocol inspection and parameter tuning. To accomplish all listed features, ZBPF uses a new configuration framework called Cisco Policy Language (CPL). The syntax for the new configuration language was heavily borrowed from the well-known MQC (Modular QoS CLI) and resembles it in many ways, by using policy-maps, class-maps, and actions association. Here is a brief overview of ZBPF configuration steps:

1. Define security zones. After initial planning, you define security regions and configure the zone names. The command to accomplish this is `zone security <NAME>` , and every zone may have a meaningful description assigned to it.

2. Define security zone-pairs. Zone-pairs are created only for source/destination zones for which traffic policies will be applied. As mentioned before, zone pairs specify the direction of traffic flow. The command to define a zone-pair is `zone-pair security <NAME> source <NAME> destination <NAME> .`

3.

https://t.me/learningnets

Define class-maps matching traffic that must have policy applied as it crosses a zone-pair. There is a new special type of class-maps used at this step, class-map type inspect, which specifies inspection scope.

4. Define policy-maps to apply action for your class-maps traffic. A new special type of policy-map called policy-map type inspect has been introduced by CPL to specify ZBPF policies. You assign the previously defined class-maps to the policy-map and specify associated actions.
5. Optionally define other advanced policy attributes. Advanced options include inspection parameter tuning, TCP/UDP session limiting, and application inspection rules. This will be detailed in a separate task.
6. Apply policy-maps to zone-pairs. This will apply a set of configured rules for traffic between source and destination zone. The command to apply the security policy to a zone pair is **`service-policy type inspect <POLICY-NAME> .`** Notice that policy applies just in one direction, and in some cases you may need to apply a policy in the opposite direction as well, depending on the actions defined in the policy-map from previous steps.
7. Assign interfaces to zones. This will basically activate the ZBPF rules, accomplished using the interface-level command **`zone-member security <ZONE> .`**

Defining Traffic Classes

Inspection type class-maps can be of type **match-all** or **match-any**, just like the regular MQC class-maps. In the first case, all match conditions must be met for the class-map to match the traffic. You can match an access-list using the command **`match access-group [number|name <NAME>]`** , which is the only method to identify traffic source or destination. In addition to matching access-lists, you can match protocols supported by the inspection engine. The list of the protocols is the same as supported by CBAC. Unlike the classic class-map, when you configure the match protocol command you do not activate NBAR engine; the protocol will be selected for inspection when the containing policy-map is assigned to a zone pair. It is common to combine access-list matching with protocol matching. For example,

```
class-map type inspect match-all CMAP_HTTP
 match access-group name ACL_INSIDE
 match protocol http
```

https://t.me/learningnets

would match HTTP traffic only for packets matching the ACL. You can use the command `show ip port-map` to list all supported protocols available for matching.

You may apply class-map nesting using the command `match class-map` . This may be very helpful in situations where you need to apply complex AND/OR logic. For example, if you want to select a set of protocols (such as HTTP, DNS, FTP, ICMP) used by a group of users, you cannot combine multiple protocols and ACL matching within a single class-map, but you may use class-map nesting as follows:

```
class-map match-any CMAP_PROTOCOLS
 match protocol http
 match protocol ftp
 match protocol dns
 match protocol icmp

!
ip access-list extended ACL_INSIDE
 perm ip 192.168.0.0 0.0.0.255 any
!
class-map match-all CMAP_INSIDE
 match class-map CMAP_PROTOCOLS
 match access-group name ACL_INSIDE
```

Applying Policy Actions

There are three general policy actions applicable under inspect-type policy-maps: **inspect**, **drop**, and **pass**, the last two supporting the additional **log** option. The first action is similar to the CBAC inspection rule. It allows for stateful inspection of traffic flowing from source to destination zone, and automatically permits returning traffic flows even for complex protocols, such as H.323. The drop action simply discards matching traffic, and the pass action permits packet flow without stateful inspection, similar to the permit action in access-lists. Here is an example:

policy-map PMAP_INSIDE class-map CMAP_INSIDE inspect

If you are using the pass action, make sure that the returning traffic is also permitted, because the pinholes will not be opened dynamically. Also, the class-map order within a policy-map is important; for example, traffic matching a pass action one way must match a pass action when returning; if an inspect action is matched it will be dropped. Every inspection policy-map has an implicit class-default configured with action drop, which can optionally be modified to pass but not to inspect. This simulates the implicit deny any action found in traditional access-lists.

https://t.me/learningnets

Self-Zone Considerations

Self-zone policy has limited functionality compared with the policies available for transit-traffic zone-pairs. First, for the stateful inspection, router-generated traffic is limited to TCP, UDP, ICMP, and complex-protocol inspection for H.323. Application Inspection (HTTP, FTP, Telnet) is not available for self-zone policies; that is, you cannot apply any advanced protocol filtering for traffic destined to/from the router. Additionally, traffic policing cannot be configured on self-zone policies; for this, the control-plane policing function should be used. You may want to configure stateful inspection for protocols terminated/originated at the router if you apply a policy between the zone self and any other configured zone. The policy-map by default drops all non-matched traffic, and this may cause control-plane problems, the most common being that routing protocol traffic is not taken into consideration.

A recommended order of operation for configuring ZBPF is:

1. Define access-lists for traffic matching.
2. Define class-maps for traffic matching (access-list and protocol based).
3. Define policy-maps.
4. Configure zones and zone-pairs, and apply the policies.
5. Assign interfaces to zones.

Although CBAC supports both IPv4 and IPv6 traffic, it uses a different set of commands for each of the two protocol versions. With ZBPF, there is a single unified command syntax for both IPv4 and IPv6 traffic, the only difference being ACL configuration, which, for IPv6, uses the `ipv6 access-list <NAME>` syntax. One important difference is that all traffic destined/originated for the self-zone using link-local addresses is not subject to ZBPF policies, but is implicitly allowed.

> ZBPF cannot inspect multicast or MPLS traffic.

Configuration

R1:

```
ip access-list extended INSIDE_TO_OUTSIDE
 permit ip 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
!
ip access-list extended OUTSIDE_TO_OUTSIDE
 deny icmp 150.1.0.0 0.0.255.255 150.1.0.0 0.0.255.255
 permit icmp any any
```

https://t.me/learningnets

```
!
ip access-list extended INSIDE_TO_SELF
 permit tcp 136.1.19.0 0.0.0.255 any eq 23
 permit tcp 136.1.19.0 0.0.0.255 any eq 80
!
ip access-list extended EIGRP
 permit eigrp any any
!
ip access-list extended ICMP
 permit icmp any any echo
!
!
class-map type inspect match-any INSIDE_TO_OUTSIDE_PROTOCOLS
 match protocol ftp
 match protocol tftp
 match protocol http
!
class-map type inspect match-all INSIDE_TO_OUTSIDE_TELNET
 match protocol telnet
 match access-group name INSIDE_TO_OUTSIDE
!
class-map type inspect match-any OUTSIDE_TO_OUTSIDE_PROTOCOLS
 match protocol http
 match protocol ssh
!
class-map type inspect match-all OUTSIDE_TO_OUTSIDE_ICMP
 match protocol icmp
 match access-group name OUTSIDE_TO_OUTSIDE
!
class-map type inspect match-any INSIDE_TO_SELF
 match access-group name INSIDE_TO_SELF
!
class-map type inspect match-any EIGRP
 match access-group name EIGRP
!
class-map type inspect match-any ICMP
 match access-group name ICMP
!
!
policy-map type inspect INSIDE_TO_OUTSIDE
 class INSIDE_TO_OUTSIDE_PROTOCOLS
  inspect
 class INSIDE_TO_OUTSIDE_TELNET
  inspect
!
policy-map type inspect OUTSIDE_TO_OUTSIDE
```

https://t.me/learningnets

```
 class OUTSIDE_TO_OUTSIDE_PROTOCOLS
  inspect
 class OUTSIDE_TO_OUTSIDE_ICMP
  inspect
!
policy-map type inspect INSIDE_TO_SELF
 class INSIDE_TO_SELF
  inspect
 class EIGRP
  pass
!
policy-map type inspect SELF_TO_INSIDE
 class ICMP
  inspect
 class EIGRP
  pass
!
policy-map type inspect OUTSIDE_TO_SELF
 class EIGRP
  pass
!
policy-map type inspect SELF_TO_OUTSIDE
 class ICMP
  inspect
 class EIGRP
  pass
!
!
zone security INSIDE
zone security OUTSIDE
!
!
zone-pair security INSIDE_TO_OUTSIDE source INSIDE destination OUTSIDE
 service-policy type inspect INSIDE_TO_OUTSIDE
!
zone-pair security OUTSIDE_TO_OUTSIDE source OUTSIDE destination OUTSIDE
 service-policy type inspect OUTSIDE_TO_OUTSIDE
!
zone-pair security SELF_TO_INSIDE source self destination INSIDE
 service-policy type inspect SELF_TO_INSIDE
!
zone-pair security SELF_TO_OUTSIDE source self destination OUTSIDE
 service-policy type inspect SELF_TO_OUTSIDE
!
zone-pair security INSIDE_TO_SELF source INSIDE destination self
 service-policy type inspect INSIDE_TO_SELF
```

https://t.me/learningnets

```
!
zone-pair security OUTSIDE_TO_SELF source OUTSIDE destination self
 service-policy type inspect OUTSIDE_TO_SELF
!
!
interface GigabitEthernet0/0
 zone-member security INSIDE
!
interface GigabitEthernet0/1.123
 zone-member security OUTSIDE
!
interface GigabitEthernet0/1.128
 zone-member security OUTSIDE
```

Verification

Verify that R1 still has all of its EIGRP neighbors:

```
R1#show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(111)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
5   136.1.19.10             Gi0/0                    10 00:38:14    7   100  0  59
4   136.1.99.4              Gi0/1.123                11 00:38:23    3   100  0  151
3   136.1.99.3              Gi0/1.123                11 00:38:23    5   100  0  130
2   136.1.99.2              Gi0/1.123                13 00:38:23    1   100  0  40
1   136.1.128.9             Gi0/1.128                11 00:38:23    4   100  0  56
0   136.1.128.2             Gi0/1.128                11 00:38:23    1   100  0  41
```

Verify that the router cannot be managed from the **OUTSIDE** zone, but it can be from the **INSIDE** zone:

```
SW1#telnet 150.1.1.1
Trying 150.1.1.1 ... % Connection timed out; remote host not responding
!
!SW2#telnet 150.1.1.1
Trying 150.1.1.1 ... Open



User Access Verification

Password: R1>
```

https://t.me/learningnets

```
!
!R1#show policy-firewall session zone-pair INSIDE_TO_SELF
 Zone-pair: INSIDE_TO_SELF
  Service-policy inspect : INSIDE_TO_SELF Class-map : INSIDE_TO_SELF(match-any)
 Established Sessions =  1
 Session 31733FA0 (136.1.19.10:38250)=>(150.1.1.1:23) tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:24, Last heard 00:00:22
          Bytes sent (initiator:responder) [37:74]
    Class-map : EIGRP(match-any)
    Class-map : class-
    (match-any)
```

Verify that users logged in to the router can PING but cannot intitiate other traffic, like telnet:

```
R1#telnet 150.1.9.9
Trying 150.1.9.9 ... % Connection timed out; remote host not responding
!
!R1#ping 150.1.9.9
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.9.9, timeout is 2 seconds: !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#ping 150.1.10.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.10.10, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R1#show policy-firewall session zone-pair SELF_TO_INSIDE
 Zone-pair: SELF_TO_INSIDE
  Service-policy inspect : SELF_TO_INSIDE Class-map : ICMP(match-any)
 Established Sessions =  1
 Session 31734DA0 (136.1.19.1:8)=>(150.1.10.10:0) icmp SIS_OPEN

          Created 00:00:01, Last heard 00:00:01
          ECHO request
          Bytes sent (initiator:responder) [360:360]
    Class-map : EIGRP(match-any)
    Class-map : class-default(match-any)
```

Verify that telnet traffic from **INSIDE** to **OUTSIDE** is allowed only within the 150.1.0.0/16 subnet, and not from **OUTSIDE** to **INSIDE**:

https://t.me/learningnets

```
SW2#telnet 150.1.9.9

Trying 150.1.9.9 ... % Connection timed out; remote host not responding
!
!SW2#telnet 150.1.9.9 /source-interface loopback0
Trying 150.1.9.9 ... Open


User Access Verification

Password: SW1>
!
!R1#show policy-firewall session zone-pair INSIDE_TO_OUTSIDE
 Zone-pair: INSIDE_TO_OUTSIDE
  Service-policy inspect : INSIDE_TO_OUTSIDE
    Class-map : INSIDE_TO_OUTSIDE_PROTOCOLS(match-any) Class-map : INSIDE_TO_OUTSIDE_TELNET(match-all)
 Established Sessions =  1
 Session 31735120 (150.1.10.10:53022)=>(150.1.9.9:23) telnet:tcp SIS_OPEN/TCP_ESTAB
          Created 00:00:13, Last heard 00:00:10
          Bytes sent (initiator:responder) [40:69]
    Class-map : class-default(match-any)
!
!SW1#telnet 150.1.10.10 /source-interface loopback0
Trying 150.1.10.10 ... % Connection timed out; remote host not responding
```

Verify that HTTP traffic is allowed from **INSIDE** to **OUTSIDE** outside the 150.1.0.0/16 subnet, and not from **OUTSIDE** to **INSIDE**:

```
SW1#telnet 150.1.10.10 80
Trying 150.1.10.10, 80 ... % Connection timed out; remote host not responding
!
!SW2#telnet 150.1.9.9 80
Trying 150.1.9.9, 80 ... Open
?
HTTP/1.1 400 Bad Request
Date: Sun, 13 Oct 2013 18:48:06 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.9.9 closed by foreign host]
!
!R1#show policy-firewall session zone-pair INSIDE_TO_OUTSIDE
```

https://t.me/learningnets

```
 Zone-pair: INSIDE_TO_OUTSIDE
  Service-policy inspect : INSIDE_TO_OUTSIDE Class-map : INSIDE_TO_OUTSIDE_PROTOCOLS(match-any)
Established Sessions =  1
Session 317354A0 (136.1.19.10:48839)=>(150.1.9.9:80) http:tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:03, Last heard 00:00:03
          Bytes sent (initiator:responder) [0:0]
    Class-map : INSIDE_TO_OUTSIDE_TELNET(match-all)
    Class-map : class-default(match-any)
```

Verify that intra-zone ICMP traffic from **OUTSIDE** to **OUTSIDE** is not allowed when both the source and destination match the 150.1.0.0/16 subnet:

```
SW1#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 .....
Success rate is 0 percent (0/5)
!
!SW1#ping 150.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
!
!R1#show policy-firewall session zone-pair OUTSIDE_TO_OUTSIDE
 Zone-pair: OUTSIDE_TO_OUTSIDE
  Service-policy inspect : OUTSIDE_TO_OUTSIDE
    Class-map : OUTSIDE_TO_OUTSIDE_PROTOCOLS(match-any) Class-map : OUTSIDE_TO_OUTSIDE_ICMP(match-all)
Established Sessions =  1
Session 31735820 (136.1.128.9:8)=>(150.1.3.3:0) icmp SIS_OPEN

          Created 00:00:02, Last heard 00:00:02
          ECHO request
          Bytes sent (initiator:responder) [360:360]
    Class-map : class-default(match-any)
```

Verify that, for example, telnet traffic is not allowed from **OUTSIDE** to **OUTSIDE**, and that HTTP traffic is allowed, even within the 150.1.0.0/16 subnet:

```
SW1#telnet 150.1.3.3 /source-interface loopback0
Trying 150.1.3.3 ... % Connection timed out; remote host not responding
!
!SW1#telnet 150.1.3.3 80 /source-interface loopback0
```

https://t.me/learningnets

```
Trying 150.1.3.3, 80 ... Open
?
HTTP/1.1 400 Bad Request
Date: Sun, 13 Oct 2013 18:59:06 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.3.3 closed by foreign host]
!
!R1#show policy-firewall session zone-pair OUTSIDE_TO_OUTSIDE
 Zone-pair: OUTSIDE_TO_OUTSIDE
  Service-policy inspect : OUTSIDE_TO_OUTSIDE Class-map : OUTSIDE_TO_OUTSIDE_PROTOCOLS(match-any)
 Established Sessions =  1
 Session 31735BA0 (150.1.9.9:48082)=>(150.1.3.3:80) http:tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:03, Last heard 00:00:03
          Bytes sent (initiator:responder) [0:0]
    Class-map : OUTSIDE_TO_OUTSIDE_ICMP(match-all)
    Class-map : class-default(match-any)
```

Change SW1's default SDM to use dual stack with the command **`sdm prefer dual-ipv4-and-ipv6 routing`**. Save the configuration and reload, and then enable IPv6 routing and IPv6 EIGRP on SW1:

```
SW1:

ipv6 unicast-routing
!
interface Vlan128
 ipv6 enable
 ipv6 eigrp 111
!
interface Loopback0
 ipv6 address 150::1:9:9/128
 ipv6 eigrp 111
!
ipv6 router eigrp 111
 no shutdown
```

Confirm that IPv6 traffic matches on configured policies. For example, test HTTP traffic from **OUTSIDE** to **OUTSIDE**:

```
SW1#telnet 150::1:3:3 80
```

https://t.me/learningnets

```
Trying 150::1:3:3, 80 ... Open
?
HTTP/1.1 400 Bad Request
Date: Sun, 13 Oct 2013 21:06:36 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150::1:3:3 closed by foreign host]
!
!R1#show policy-firewall session zone-pair OUTSIDE_TO_OUTSIDE
 Zone-pair: OUTSIDE_TO_OUTSIDE
  Service-policy inspect : OUTSIDE_TO_OUTSIDE Class-map : OUTSIDE_TO_OUTSIDE_PROTOCOLS(match-any)
Established Sessions =  1
Session 317362A0 [150::1:9:9]:26548=>[150::1:3:3]:80 tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:07, Last heard 00:00:07
          Bytes sent (initiator:responder) [0:0]
    Class-map : OUTSIDE_TO_OUTSIDE_ICMP(match-all)
    Class-map : class-default(match-any
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Transparent Mode

Task

- Configure two security zones on R3: **VLAN123** and **VLAN34**.
- ICMP traffic initiated from **VLAN123** zone should match the following output:

```
%FW-6-PASS_PKT: (target:class)-(VLAN123_TO_VLAN34:ICMP) Passing icmp pkt
136.1.99.1:0 => 150.1.4.4:0 with ip ident 0
```

- Inspect HTTP traffic on ports 80 and 21 destined for R4's Loopback0 address.
- Configure R3 to log for transit traffic not matching any of the configured class-maps.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

It is possible to use CPL (Cisco Policy Language) and define transparent firewall rules using ZBPF syntax. In fact, this approach has many advantages, because it operates with a better abstraction (zone) instead of interfaces and allows for easy-to-understand and concise configurations. To configure ZBPF, simply apply zone names to the interfaces configured for Ethernet bridging. It is possible, but not recommended, to configure layer3 and layer2 interfaces in the same zone. However, in most cases you may want to keep those domains separate.

Any BVI interfaces on the router automatically belong to the self zone, and the

https://t.me/learningnets

policy should be configured appropriately. If you want to treat the whole layer2 domain behind the BVI as a single layer3 security zone, you may simply assign a zone to the BVI, but not to the bridging interfaces. As for multicast and broadcast traffic types, both are not inspected by ZBPF. Moreover, in transparent mode both traffic types are permitted across the firewall configured with zones. Other than this, there are no functional differences between ZBPF in routed and transparent modes.

You may compare the BVI with an SVI (Switch Virtual Interface) from a layer3 Catalyst switch. In addition to the **`bridge irb`** command, you must add the command **`bridge <number> route ip`** ; otherwise, the BVI interface will not route IP packets. To configure an interface/sub-interface as a member of a bridge group, use the command **`bridge-group <number>`** in interface configuration mode. By default, the new group starts with Spanning Tree Protocol disabled, so you need the command **`bridge <number> protocol ieee`** to enable the IEEE-compatible spanning tree for the bridge-group. Notice that for 802.1Q sub-interfaces, the firewall will run the PVST+ variant of STP. AT the interface level, just like in routed mode, assign security zones. For example:

```
bridge irb
bridge-group 1 protocol ieee
bridge 1 route ip
!
interface FastEthernet0/0
 bridge-group 1
 zone-member security INSIDE
!
interface FastEthernet0/1
 bridge-group 1
 zone-member security OUTSIDE
```

As mentioned previously, ZBPF supports the same protocol set as CBAC. You may see the port-numbers used by various applications by issuing the commands **`show ip port-map`** or **`show ipv6 port-map`** . If you need to define a custom protocol port, you may use the command **`ip port-map <protocol>`** or **`ipv6 port-map <protocol>`** to assign new port numbers to the protocol. If you need to re-use a well-known port for a different application, such as to assign a telnet application to port 80, you must specify an access-list for which the mapping is valid, using the command **`ip port-map telnet port <PORT> list <ACL>`** or **`ipv6 port-map telnet port <PORT> list <ACL>`** . These are the same requirements as with the classic CBAC firewall. For example:

https://t.me/learningnets

```
access-list 1 permit 136.1.27.0 0.0.0.255
!
ip port-map telnet port tcp 80 list 99
```

> Just as with CBAC, IPv6 traffic is automatically bridged and cannot be inspected by ZPFW in transparent mode.

Configuration

R3:

```
ip access-list extended ICMP
 permit icmp any any
!
ip access-list extended HTTP
 permit tcp any host 150.1.4.4 eq 80
 permit tcp any host 150.1.4.4 eq 21
!
access-list 1 permit host 150.1.4.4
!
!
ip port-map http port tcp 21 list 1
!
!
class-map type inspect match-all ICMP
 match protocol icmp
 match access-group name ICMP
!
class-map type inspect match-all HTTP
 match access-group name HTTP
 match protocol http
!
!
policy-map type inspect VLAN123_TO_VLAN34
 class ICMP
  pass log
 class HTTP
  inspect
 class class-default
  drop log
!
```

https://t.me/learningnets

```
policy-map type inspect VLAN34_TO_VLAN123
 class ICMP
  pass log
 class class-default
  drop log
!
!
zone security VLAN123
zone security VLAN34
!
!
zone-pair security VLAN123_TO_VLAN34 source VLAN123 destination VLAN34
 service-policy type inspect VLAN123_TO_VLAN34
!
zone-pair security VLAN34_TO_VLAN123 source VLAN34 destination VLAN123
 service-policy type inspect VLAN34_TO_VLAN123
!
!
interface FastEthernet0/1.123
 zone-member security VLAN123
!
interface FastEthernet0/1.34
 zone-member security VLAN34
```

Verification

Verify ZBPF configuration:

```
R3#show policy-firewall config zone
Zone: self
  Description: System defined zone
Zone: VLAN123
  Member Interfaces:
    FastEthernet0/1.123
  Member of Zone-pairs:
    Zone VLAN123 is source of Zone-pair VLAN123_TO_VLAN34
    Zone VLAN123 is destination of Zone-pair VLAN34_TO_VLAN123
Zone: VLAN34
  Member Interfaces:
    FastEthernet0/1.34
  Member of Zone-pairs:
    Zone VLAN34 is destination of Zone-pair VLAN123_TO_VLAN34
    Zone VLAN34 is source of Zone-pair VLAN34_TO_VLAN123
!
!R3#show policy-firewall config policy-map
```

https://t.me/learningnets

```
  Policy Map type inspect VLAN123_TO_VLAN34 Class ICMP
      Pass log Class HTTP
      Inspect
    Class class-default
      Drop log

  Policy Map type inspect VLAN34_TO_VLAN123
    Class ICMP
      Pass log
    Class class-default
      Drop log
!
!R3#show policy-firewall config class-map
 Class Map type inspect match-all ICMP (id 1)
   Match access-group name ICMP
   Match protocol icmp
 Class Map type inspect match-all HTTP (id 2)

   Match access-group name HTTP
   Match protocol http
```

Verify that ICMP traffic initiated from VLAN123 zone matched the requested output; no session states should be created due to the **pass** action:

```
R1#ping 150.1.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.4.4, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R3:
%FW-6-PASS_PKT: (target:class)-(VLAN123_TO_VLAN34:ICMP) Passing icmp pkt 136.1.99.1:0 => 150.1.4.4:0 with ip ident 0
!
!R4#ping 150.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R3:
%FW-6-PASS_PKT: (target:class)-(VLAN34_TO_VLAN123:ICMP) Passing icmp pkt 136.1.99.4:0 => 150.1.1.1:0 with ip ident 0
!
!R3#show policy-firewall session zone-pair VLAN123_TO_VLAN34
 Zone-pair: VLAN123_TO_VLAN34
  Service-policy inspect : VLAN123_TO_VLAN34
    Class-map : ICMP(match-all)
```

https://t.me/learningnets

```
    Class-map : HTTP(match-all)
    Class-map : class-default(match-any)
!
!R3#show policy-firewall stats zone-pair VLAN123_TO_VLAN34

policy exists on zp VLAN123_TO_VLAN34
 Zone-pair: VLAN123_TO_VLAN34

  Service-policy inspect : VLAN123_TO_VLAN34
 Class-map: ICMP (match-all)
      Match: access-group name ICMP
      Match: protocol icmp Pass
 10 packets, 800 bytes


    Class-map: HTTP (match-all)
      Match: access-group name HTTP
      Match: protocol http

   Inspect
        Session creations since subsystem startup or last reset 0
        Current session counts (estab/half-open/terminating) [0:0:0]
        Maxever session counts (estab/half-open/terminating) [0:0:0]
        Last session created 00:06:21
        Last statistic reset 00:00:14
        Last session creation rate 0
        Maxever session creation rate 0
        Last half-open session total 0
        TCP reassembly statistics
        received 0 packets out-of-order; dropped 0
        peak memory usage 0 KB; current usage: 0 KB
        peak queue length 0


    Class-map: class-default (match-any)
      Match: any
      Drop
        0 packets, 0 bytes
```

Verify that HTTP traffic is inspected:

```
R1#telnet 150.1.4.4 80
Trying 150.1.4.4, 80 ... Open
?
HTTP/1.1 400 Bad Request
```

https://t.me/learningnets

```
Date: Tue, 15 Oct 2013 05:21:24 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request[Connection to 150.1.4.4 closed by foreign host]
!
!R3#show policy-firewall session zone-pair VLAN123_TO_VLAN34
 Zone-pair: VLAN123_TO_VLAN34
  Service-policy inspect : VLAN123_TO_VLAN34
    Class-map : ICMP(match-all) Class-map : HTTP(match-all)
Established Sessions =  1
Session 49E62320 (136.1.99.1:49957)=>(150.1.4.4:80) http:tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:16, Last heard 00:00:16
          Bytes sent (initiator:responder) [0:0]
    Class-map : class-default(match-any)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Asymmetric Routing

Task

- Configure two security zones on R1: VLAN 128 member of **INSIDE** and VLAN 123 member of **OUTSIDE**.
- Configure R1 and R2 for zone-based firewall stateful failover and bidirectionally inspect TELNET protocol.
- HTTP traffic should work bidirectionally, but no session states should be created.
- Ensure that R1 is preferred as the **ACTIVE** device as long as it has both interfaces in the **UP** state.
- Ensure that if there is asymmetric routing, traffic through the firewall is still functional.
- If an additional interface is required, use GigabitEthernet0/1.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

IOS ZBPF supports asymmetric routing at the redundancy group level. If compared with the ZBPF HA configuration, there are two additional settings required on both routers:

https://t.me/learningnets

- Enable asymmetric routing at the redundancy group level.
- Enable asymmetric routing on the interface not a member of the redundancy group, thus on the interface where return traffic will be received.

Essentially, we must configure the **interface interlink** interface, over which a tunnel will be automatically created between the **ACTIVE** and **STANDBY** device, so that return traffic received by the **STANDBY** device can be routed through this tunnel and injected back into the **ACTIVE device**. This is required because only the **ACTIVE** device maintains the session state and replicates the session state to the **STANDBY** device. The following restrictions apply to the asymmetric configuration and functionality:

- Traffic matching an **inspect** action will be automatically routed from the **STANDBY** device to the **ACTIVE** device.
- Traffic matching a **pass** action will be automatically routed from the **STANDBY** device to the **ACTIVE** device, but because there is no session state it will match on the **class-default** class-map; therefore, this class-map must have the **pass** action configured.
- Asymmetric traffic matching an **inspect** action cannot be initiated bidirectionally; initiating traffic must be received inbound on the interface where the redundancy group and VIP is configured.
- Asymmetric routing matching a **pass** action can be initiated bidirectionally.

This setup is useful, for example, when you have two Internet routers configured with ZBPF in HA, each having a different Internet connection, and you have no real control over how traffic is received back from the Internet, so it can be received by both **ACTIVE** or **STANDBY** devices.

Configuration

Configure R1 and R2 for High Availability and Asymmetric Routing:

```
SW2:
interface range fastEthernet 1/0/1 - 2
 switchport trunk allowed vlan add 1


R1:
interface GigabitEthernet0/1
 ip address 136.1.1.1 255.255.255.0
!
parameter-map type inspect global
 redundancy
 exit
!
```

https://t.me/learningnets

```
redundancy
 application redundancy
  group 1
   name ZBPF
   preempt
   priority 150
   control GigabitEthernet0/1 protocol 1
   data GigabitEthernet0/1
   asymmetric-routing interface GigabitEthernet0/1
!
interface GigabitEthernet0/1.123
 redundancy rii 10 decrement 60
 redundancy asymmetric-routing enable
!
interface GigabitEthernet0/1.128
 redundancy rii 20 decrement 60
 redundancy group 1 ip 136.1.128.12 exclusive
```

R2:

```
interface GigabitEthernet0/1
 ip address 136.1.1.2 255.255.255.0
!
parameter-map type inspect global
 redundancy
 exit
!
redundancy
 application redundancy
  group 1
   name ZBPF
   preempt
   priority 100
   control GigabitEthernet0/1 protocol 1
   data GigabitEthernet0/1
   asymmetric-routing interface GigabitEthernet0/1
!
interface GigabitEthernet0/1.123
 redundancy rii 10
 redundancy asymmetric-routing enable
!
interface GigabitEthernet0/1.128
 redundancy rii 20
```

https://t.me/learningnets

```
redundancy group 1 ip 136.1.128.12 exclusive
```

Identically configure R1 and R2 for ZBPF:

https://t.me/learningnets

R1: and R2:

```
class-map type inspect match-any TELNET
 match protocol telnet
!
class-map type inspect match-any HTTP
 match protocol http
!
policy-map type inspect INSIDE_TO_OUTSIDE
 class TELNET
  inspect
 class HTTP
  pass
 class class-default
  pass
!
policy-map type inspect OUTSIDE_TO_INSIDE
 class TELNET
  inspect
 class HTTP
  pass
 class class-default
  pass
!
!
zone security INSIDE
zone security OUTSIDE
!
!
zone-pair security INSIDE_TO_OUTSIDE source INSIDE destination OUTSIDE
 service-policy type inspect INSIDE_TO_OUTSIDE
!
zone-pair security OUTSIDE_TO_INSIDE source OUTSIDE destination INSIDE
 service-policy type inspect OUTSIDE_TO_INSIDE
!
!
interface GigabitEthernet0/1.128
 zone-member security INSIDE
!
interface GigabitEthernet0/1.123
 zone-member security OUTSIDE
```

https://t.me/learningnets

Verification

Force asymmetric routing by modifying EIGRP delay or bandwidth, initiating traffic needs to be routed through R1 and return traffic through R2:

```
R1:

interface GigabitEthernet0/1.128
 bandwidth 10000
```

Verify that there is asymmetric routing between SW1 and R3 Loopback0 interfaces, and that session is initiated through R1 and returned through R2:

```
SW1#show ip route 150.1.3.3
Routing entry for 150.1.3.3/32
  Known via "eigrp 111", distance 90, metric 156416, type internal
  Redistributing via eigrp 111
  Last update from 136.1.128.1 on Vlan128, 00:04:14 ago
  Routing Descriptor Blocks: * 136.1.128.1, from 136.1.128.1, 00:04:14 ago, via Vlan128
      Route metric is 156416, traffic share count is 1
      Total delay is 5110 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
!
!R3#show ip route 150.1.9.9
Routing entry for 150.1.9.9/32
  Known via "eigrp 111", distance 90, metric 307200, type internal
  Redistributing via eigrp 111
  Last update from 136.1.99.2 on BVI99, 00:02:01 ago
  Routing Descriptor Blocks: * 136.1.99.2, from 136.1.99.2, 00:02:01 ago, via BVI99

      Route metric is 307200, traffic share count is 1
      Total delay is 11000 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

Verify that asymmetric routing tunnel is established between R1 and R2, but there is no traffic sent or received:

```
R1#show redundancy application asymmetric-routing interface group 1
AR Group ID:1 interface GigabitEthernet0/1 neighbor 136.1.1.2,
  transport context:
    my ip 136.1.1.1, my port 53000
```

https://t.me/learningnets

```
    peer ip 136.1.1.2, peer port 53000
!
!R1#show redundancy application asymmetric-routing tunnel group 1
Group ID:1
  rii    10, idb GigabitEthernet0/1.123 packet sent: 0, packet received: 0
byte sent:   0, byte recv:        0
        encap: length 32
         IP :45 00 00 00 00 00 00 00 FF 11 00 00 88 01 01 01 88 01 01 02
         UDP:CF 08 CF 08 00 00 00 00
         AR :00 01 00 0A
!
!R2#show redundancy application asymmetric-routing interface group 1
AR Group ID:1 interface GigabitEthernet0/1 neighbor 136.1.1.1,
  transport context:
    my ip 136.1.1.2, my port 53000
    peer ip 136.1.1.1, peer port 53000
!
!R2#show redundancy application asymmetric-routing tunnel group 1
Group ID:1
  rii    10, idb GigabitEthernet0/1.123 packet sent: 0, packet received: 0
byte sent:   0, byte recv:        0

        encap: length 32
         IP :45 00 00 00 00 00 00 00 FF 11 00 00 88 01 01 02 88 01 01 01
         UDP:CF 08 CF 08 00 00 00 00
         AR :00 01 00 0A
```

Initiate a telnet session from SW1, and verify that the session state is replicated and traffic was sent routed by R2 through the tunnel with R1:

```
SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open



User Access Verification

Password: R4>
!
!R1#show policy-firewall session ha
Session 318AEFA0 (136.1.128.9:40390)=>(150.1.4.4:23) telnet:tcp SIS_OPEN/TCP_ESTAB
          Created 00:01:15, Last heard 00:01:13
          Bytes sent (initiator:responder) [37:74] HA State: ACTIVE, RG ID: 1
      Established Sessions = 1
!
!R1#show redundancy application asymmetric-routing tunnel group 1
```

https://t.me/learningnets

```
Group ID:1
  rii    10, idb GigabitEthernet0/1.123        packet sent: 0,packet received: 14
        byte sent:   0,byte recv:        1134
        encap: length 32
         IP :45 00 00 00 00 00 00 00 FF 11 00 00 88 01 01 01 88 01 01 02
         UDP:CF 08 CF 08 00 00 00 00
         AR :00 01 00 0A
!
!R2#show policy-firewall session ha
Session 3144E7C0 (136.1.128.9:40390)=>(150.1.4.4:23) telnet:tcp SIS_OPEN/TCP_ESTAB
          Created 00:01:27, Last heard never
          Bytes sent (initiator:responder) [0:0]HA State: STANDBY, RG ID: 1
      Established Sessions = 1
!
!R2#show redundancy application asymmetric-routing tunnel group 1
Group ID:1
  rii    10, idb GigabitEthernet0/1.123 packet sent: 14
, packet received: 0 byte sent:   1134
, byte recv:       0
        encap: length 32
         IP :45 00 00 00 00 00 00 00 FF 11 00 00 88 01 01 02 88 01 01 01
         UDP:CF 08 CF 08 00 00 00 00
         AR :00 01 00 0A
```

Note that asymmetric routing does not work in the opposite direction for traffic matching a inspect action, because interfaces configured for asymmetric routing need to receive the return traffic flow, not the initiating traffic flow:

```
R3#telnet 150.1.9.9
Trying 150.1.9.9 ... % Connection timed out; remote host not responding
```

Initiate HTTP traffic from SW1; there will be no session state because of pass action, but traffic is allowed and on the return path it will match on the **class-default** class-map of R2. Asymmetric traffic matching a pass action and sent through the tunnel will always match the **class-default** class map:

```
SW1#telnet 150.1.3.3 80
Trying 150.1.3.3, 80 ... Open
?
HTTP/1.1 400 Bad Request
Date: Tue, 15 Oct 2013 21:46:29 GMT
Server: cisco-IOS
Accept-Ranges: none
```

https://t.me/learningnets

```
400 Bad Request
[Connection to 150.1.3.3 closed by foreign host]
!
!R1#show redundancy application asymmetric-routing tunnel group 1
Group ID:1
  rii    10, idb GigabitEthernet0/1.123        packet sent: 0,packet received: 414
         byte sent:   0,byte recv:        32334
         encap: length 32
          IP :45 00 00 00 00 00 00 00 FF 11 00 00 88 01 01 01 88 01 01 02
          UDP:CF 08 CF 08 00 00 00 00
          AR :00 01 00 0A
!
!R2#show redundancy application asymmetric-routing tunnel group 1
Group ID:1
  rii    10, idb GigabitEthernet0/1.123 packet sent: 414
, packet received: 0 byte sent:   32334
, byte recv:       0
         encap: length 32
          IP :45 00 00 00 00 00 00 00 FF 11 00 00 88 01 01 02 88 01 01 01
          UDP:CF 08 CF 08 00 00 00 00
          AR :00 01 00 0A
```

Note that asymmetric routing also works in the opposite direction for traffic matching a pass action; it will match on the **HTTP** class-map on R2, sent over the tunnel to R1, and match on **class-default** class-map on R1:

```
R3#telnet 150.1.9.9 80
Trying 150.1.9.9, 80 ... Open
?
HTTP/1.1 400 Bad Request
Date: Tue, 15 Oct 2013 21:51:31 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 150.1.9.9 closed by foreign host]
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Application Inspection

Task

- Configure two security zones on R1: VLAN 19 member of **INSIDE** and VLAN 128 member of **OUTSIDE**.
- Configure the firewall so that the **INSIDE** users are not allowed to use insecure POP3/IMAP4 logins or invalid commands to the **OUTSIDE** servers.
- Block image downloads from **digg.com** for **INSIDE** users and log any attempts to use AOL Messenger for non text-chat service (such as file transfer).
- Image files can have the filename extensions **.jpg**, **.png**, or **.gif**.
- Log all violations.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Application inspection and control (AIC) applies at Layer 7 of the OSI model, allowing for policy configuration based on application protocol parameters. Applications are the most vulnerable part of your network, and thus enforcing strict per-application policy is important to increase the overall security. For example, you may want your firewall to block certain SMTP commands that open vulnerabilities in the servers. Alternatively, you may want to make sure that HTTP connections are not used to tunnel instant-messenger services or to connect to other non-HTTP

https://t.me/learningnets

services. This is only possible with application-level inspection. Although it is comparable to ASA's MPF functionality, it is more limited. Cisco IOS ZBPF supports application inspection and control for:

- HTTP
- SMTP/ESMTP
- POP3/IMAP
- P2P Application Traffic (eDonkey, Fasttrack, Kazaa2, GnUtella, etc.)
- IM Applications (AOL, MSN, Yahoo, ICQ)
- Sun RPC
- SIP/H323

AIC uses a special type of class-maps and policy-maps to configure application protocol inspection. These are application-specific inspection class-maps and policy-maps, such as **`class-map type inspect http`** , **`policy-map type inspect http`** , and so on for every application protocol supported. The application class-map defines protocol parameters and should be used inside the application policy map. The application policy-map should be later applied as a nested (child) service-policy under the generic inspection policy-map—for example, resetting HTTP traffic if there is a protocol violation identified:

```
class-map type inspect http match-any HTTP_APPLICATION
  match req-resp protocol-violation
!
policy-map type inspect http HTTP_APPLICATION
  class type inspect http HTTP_APPLICATION
   reset
!
class-map type inspect match-any HTTP_TRAFFIC
 match protocol http
!
policy-map type inspect INSIDE_TO_OUTSIDE
 class type inspect HTTP_TRAFFIC
  inspect
  service-policy http HTTP_APPLICATION
```

https://t.me/learningnets

It is important that the parent L3/L4 class-map, **HTTP_TRAFFIC** in our case, define only the protocols supported by AIC. For example, if the parent class-map defines FTP protocol via the command `match protocol ftp` , you cannot apply the nested service-policy for any AIC-supported protocol, because FTP is not supported by AIC. Application inspection and control (AIC) has a different set of features for every of the above-mentioned protocols.

- **HTTP inspection** offers granular filtering on several types of application activity, including different methods (GET/POST), HTTP header contents, and URIs. Additionally, HTTP inspection allows you to limit document transfer size, web address lengths, and the flow of browser requests/server replies to enforce compliance with application-behavior standards. Finally, you may limit types of content (such as block java apps, images, videos, or compressed content) that is transferred over HTTP connections.
- **SMTP inspection** can limit content length and enforce protocol compliance. It checks for command syntax and permits only the legal SMTP/ESMTP commands.
- **POP3 and IMAP inspection** can help ensure that users are using secure authentication mechanisms to prevent compromise of user credentials.
- **P2P inspection** can help restrict certain services used by these applications, such as file transfer, file search, or chat.
- **IM inspection** can be used to allow or restrict certain services, such as to allow text-chat but disallow voice/video calling.
- **SIP/H323 inspection** can be used to enforce protocol compliance with RFC specifications or restrict certain application-level functions and negotiations.

One special type of objects that is often used with the application inspection is a regular-expressions parameter map. This allows you to define sets of regular expressions to be later used with application inspection, such as to match URL strings or command parameters like FTP GET filename or SMTP recipients. The syntax to define a regular-expression parameter map is:

https://t.me/learningnets

```
parameter-map type regex REGEX_CISCO
  pattern  [cC][iI][sS][cC][oO]
!
class-map type inspect http CMAP_CISCO
 match request header host regex REGEX_CISCO
```

Configuration

R1:

```
parameter-map type regex REGEX_IMAGES
 pattern .*\.([jJ][pP][gG]|[pP][nN][gG]|[gG][iI][fF])
!
parameter-map type regex REGEX_DIGG
 pattern .*[dD][iI][gG][gG]\.[cC][oO][mM]
!
class-map type inspect http match-all CMAP_HTTP_AIC
 match request header host regex REGEX_DIGG
 match request uri regex REGEX_IMAGES
!
class-map type inspect pop3 match-any CMAP_POP3_AIC
 match login clear-text
 match invalid-command
!
class-map type inspect imap match-any CMAP_IMAP_AIC
 match login clear-text
 match invalid-command
!
class-map type inspect aol match-any CMAP_AOL_AIC
 match no service text-chat
!
policy-map type inspect http PMAP_HTTP_AIC
 class CMAP_HTTP_AIC
  reset
  log
!
policy-map type inspect pop3 PMAP_POP3_AIC
 class CMAP_POP3_AIC
  reset
  log
!
policy-map type inspect imap PMAP_IMAP_AIC
```

https://t.me/learningnets

```
 class CMAP_IMAP_AIC
  reset
  log
!
policy-map type inspect im CMAP_AOL_AIC
 class CMAP_AOL_AIC
  log
!
class-map type inspect match-all INSIDE_TO_OUTSIDE_HTTP
 match protocol http
!
class-map type inspect match-all INSIDE_TO_OUTSIDE_POP3
 match protocol pop3
!
class-map type inspect match-all INSIDE_TO_OUTSIDE_IMAP
 match protocol imap
!
!
class-map type inspect match-all INSIDE_TO_OUTSIDE_AOL
 match protocol aol
!
policy-map type inspect INSIDE_TO_OUTSIDE
 class INSIDE_TO_OUTSIDE_HTTP
  inspect
  service-policy http PMAP_HTTP_AIC
class INSIDE_TO_OUTSIDE_POP3
  inspect
  service-policy pop3 PMAP_POP3_AIC
class INSIDE_TO_OUTSIDE_IMAP
  inspect
  service-policy imap PMAP_IMAP_AIC
class INSIDE_TO_OUTSIDE_AOL
  inspect
  service-policy im PMAP_AOL_AIC
!
!
zone security INSIDE
zone security OUTSIDE
!
!
zone-pair security INSIDE_TO_OUTSIDE source INSIDE destination OUTSIDE
 service-policy type inspect INSIDE_TO_OUTSIDE
!
!
interface GigabitEthernet0/0
 zone-member security INSIDE
```

https://t.me/learningnets

```
!
interface GigabitEthernet0/1.128
 zone-member security OUTSIDE
```

Verification

Connect from **INSIDE** to **OUTSIDE** on HTTP service and simulate requesting some forbidden image files from **digg.com**:

```
SW2#telnet 150.1.9.9 80
Trying 150.1.9.9, 80 ... OpenGET /image.jpg HTTP/1.1
Host: digg.com
[Connection to 150.1.9.9 closed by foreign host]
!
!R1:

%APPFW-4-HTTP_HDR_FIELD_REGEX_MATCHED: Header field (^[Hh][Oo][Ss][Tt]:.*[dD][iI][gG][gG]\.[cC][oO][mM]) matched - r
%APPFW-4-HTTP_URI_REGEX_MATCHED: URI regex (.*\.([jJ][pP][gG]|[pP][nN][gG]|[gG][iI][fF])) matched - resetting sessio
```

Because there is no POP3/IMAP server, enable TCP small services on SW1 and configure R1 to inspect this traffic as POP3 and connect from SW2:

```
R1:
access-list 1 permit 150.1.9.9
ip port-map pop3 port tcp 7 list 1
!
!SW1:
service tcp-small-servers
!
!SW2#telnet 150.1.9.9 7
Trying 150.1.9.9, 7 ... OpenWWHOHO

[Connection to 150.1.9.9 closed by foreign host]
!
!R1:
%FW-5-POP3_INVALID_COMMAND: (target:class)-(INSIDE_TO_OUTSIDE:INSIDE_TO_OUTSIDE_POP3):Invalid POP3 command from init
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Tuning

Task

- Configure two security zones on R1: VLAN 19 member of **INSIDE**, and VLAN 128 member of **OUTSIDE**.
- Bidirectionally inspect ICMP and telnet traffic.
- To prevent DoS attacks, limit the **INSIDE** to **OUTSIDE** traffic flow to 256Kbps.
- Ensure that no more than 2000 half-open sessions, and no more than 100 new half-open sessions, generated per minute are allowed.
- Stop dropping the half-open session states when the absolute number of half-open connections falls below 1000 and the rate falls below 10.
- Ensure that the firewall will log dropped packets in case of misconfiguration.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Zone-Based Policy Firewall supports two types of rate-limiting:

- Limiting aggregate packet rate for intra-zone or inter-zone traffic.
- Limiting the maximum number and/or rate of the half-open connections for TCP/UDP sessions.

https://t.me/learningnets

The first option resembles QoS traffic policing, but it applies to traffic matching the inspect class maps and essentially limits the aggregated flows between interface zones, not out/in of a single interface. Thus, if you have multiple interfaces in a single zone, the configuration will apply aggregate limiting across all interfaces. The command to apply policing is **`police rate <Bps> burst <Bytes>`** configured under the respective class of a policy-map. Any traffic bursts exceeding the configured rate are dropped by the firewall; that is, you cannot apply advanced actions such as remarking. The smaller the burst, the less traffic could be sent instantly after an idle period. By setting a larger burst value, you ensure smoother traffic flows but risk getting sudden spikes of heavy traffic bursts. There is no optimal burst value; this is usually picked up based on experimental data. Here is an example of ZBPF policing configuration:

```
policy-map type inspect OUTSIDE_TO_INSIDE
 class ICMP
  inspect
  police rate 256000 burst 8000
```

> Notice that you need the inspect action to be configured before enabling the policing. Another limitation of the traffic-rate policing is that it is not supported with the zone self, so for any zone pair that includes this zone.

The second rate-limiting feature applies to the TCP/UDP-based connections. It helps preventing resource exhaustion by DoS attacks, such as SYN flooding, and behaves similarly to the TCP Intercept feature. However, the only mode available is watch mode, in which the firewall passively observes connections and tears them apart when thresholds are crossed. The UDP sessions are emulated as bidirectional UDP packet exchange that time out after some configurable timer value. To apply connection-based limits to a particular class applied under a policy map, you must define an inspect parameter-map first. This object defines various connection attributes, including TCP-intercept like parameters and various timeouts. The values defined in the parameter-map are similar to the CBAC parameter values familiar to you from the classic firewall configuration. Here is a full list of the inspect parameter-map configurable options.

```
parameter-map type inspect parameter-map-name
 alert {on|off}
 audit-trail {on|off}
 dns-timeout <seconds>
 icmp idle-timeout <seconds>
 ipv6 routing-header-enforcement
```

https://t.me/learningnets

```
 max-incomplete {low <num-of-conn> | high <num-of-conn>}
 one-minute {low <num-of-conn> | high <num-of-conn>}
 tcp max-incomplete host <threshold> [block-time <minutes>]
 sessions maximum <sessions>
 tcp finwait-time <seconds>
 tcp idle-time <seconds>
 tcp synwait-time <seconds>
 udp idle-time <seconds>
```

The alert and audit-trail options define whether to produce an alert for a particular event (security violation, improper TCP stream structure) and generate connection statistics information, such as number of bytes sent and received. The various timeout values (such as dns-timeout, tcp finwait-time, and tcp idle-time) specify the amount of time to expire a particular type of connection. The settings for the TCP-intercept feature are:

- **max-incomplete low/high**: specifies the low (stop dropping half-open sessions) and high (start dropping half-open sessions) absolute thresholds for the half-open TCP/UDP connections. A half-open UDP connection is the connection that sent but did not received any packet.
- **one-minute low/high**: specifies the per-minute thresholds for the number of sessions. This effectively controls the rate of the half-open connections.
- **tcp max-incomplete host**: sets the number of half-open sessions allowed per host and the time to block this host after the limit has been exceeded.

To apply the parameter-map to a zone pair, use the option to the inspect command configured under a particular class. The traffic flows matching this class will inherit the parameters specified by the parameter map. Other classes will use the default parameter-map that could be viewed using the command **`show parameter-map type inspect default`** . Here is an example that limits the number of session to one hundred:

```
parameter-map type inspect INSIDE
  sessions maximum 100
!
policy-map type inspect INSIDE_TO_OUTSIDE
  cass INSIDE_TRAFFIC
    inspect INSIDE
```

There are three possible parameter-maps:

- The **default** parameter-map, which is inherited by all user-configured class-maps as

https://t.me/learningnets

well as the class-default class-map.

- The **global** parameter-map, which inherits some settings from the **default** parameter-map; it globally affects the ZBPF functionality and has some additional configurable settings not available within the **default** or user-configurable parameter-maps.
- The user-configurable parameter-map, which inherits all settings from the **default** parameter-map unless modified, and can only be applied to user-configured class-maps, not to the class-default class-map.

The additional settings available in the **global** parameter-map are:

```
parameter-map type inspect global
 waas {enable|disable}
 l2-transparent dhcp-passthrough
 tcp reset-PSH
 log {dropped-packets|summary}
```

The two most important are:

- **`l2-transparent dhcp-passthrough`**, which has the same functionality found in CBAC, where DHCP packets are allowed to pass through the ZBPF regardless of the class-map settings.
- **`log dropped-packets enable`**, which activates ZBPF logging for several cases, such as when traffic would be dropped because one of the interfaces is not a member of any security-zone, or a traffic policy (zone-pair) is not defined between zones the traffic requires to transit.

When the **`log dropped-packets enable`** option is enabled under the **global** parameter map, the following messages may appear on the router console, depending on the mode (transparent or routed), if the traffic is dropped because of the described scenarios:

```
%FW-6-DROP_PKT: Dropping Unknown-l4 session 0.0.0.0:0 0.0.0.0:0  due to  policy match failure with ip ident 0
%FW-6-DROP_PKT: Dropping Unknown-l4 session 0.0.0.0:0 0.0.0.0:0  due to  One of the interfaces not being cfged for z
!
!
%FW-6-DROP_PKT: Dropping tcp session 150.1.10.10:58776 150.1.9.9:23  due to  Policy not present on zone-pair with ip
%FW-6-DROP_PKT: Dropping tcp session 150.1.10.10:45826 150.1.9.9:23  due to  One of the interfaces not being cfged f
```

Configuration

https://t.me/learningnets

R1:

```
parameter-map type inspect PARAMETER
 max-incomplete low 1000
 max-incomplete high 2000
 one-minute low 10
 one-minute high 100
!
parameter-map type inspect global
 log dropped-packets enable
!
class-map type inspect match-any PROTOCOLS
 match protocol icmp
 match protocol telnet
!
policy-map type inspect INSIDE_TO_OUTSIDE
 class PROTOCOLS
  inspect PARAMETER
  police rate 256000 burst 16000
!
policy-map type inspect OUTSIDE_TO_INSIDE
 class PROTOCOLS
  inspect
!
!
zone security INSIDE
zone security OUTSIDE
!
!
zone-pair security INSIDE_TO_OUTSIDE source INSIDE destination OUTSIDE
 service-policy type inspect INSIDE_TO_OUTSIDE
!
zone-pair security OUTSIDE_TO_INSIDE source OUTSIDE destination INSIDE
 service-policy type inspect OUTSIDE_TO_INSIDE
!
!
interface GigabitEthernet0/0
 zone-member security INSIDE
!
interface GigabitEthernet0/1.128
 zone-member security OUTSIDE
```

https://t.me/learningnets

Verification

Verify all three parameter-map settings, and note that the **global** parameter-map inherits some settings from the **default** parameter map:

```
R1#show parameter-map type inspect
 parameter-map type inspect default
   audit-trail off
   alert on max-incomplete low 2147483647
 max-incomplete high 2147483647
 one-minute low 2147483647
 one-minute high 2147483647
   udp idle-time 30
   icmp idle-time 10
   dns-timeout 5
   tcp idle-time 3600
   tcp finwait-time 5
   tcp synwait-time 30
   tcp max-incomplete host 4294967295 block-time 0
   sessions maximum 2147483647
 parameter-map type inspect PARAMETER
   audit-trail off
   alert on max-incomplete low 1000
 max-incomplete high 2000
 one-minute low 10
 one-minute high 100
   udp idle-time 30
   icmp idle-time 10
   dns-timeout 5
   tcp idle-time 3600
   tcp finwait-time 5
   tcp synwait-time 30
   tcp max-incomplete host 4294967295 block-time 0
   sessions maximum 2147483647
 parameter-map type inspect global
   alert on
   sessions maximum 2147483647
   waas disabled
   l2-transparent dhcp-passthrough disabled log dropped-packets enabled
   log summary disabled
   max-incomplete low 18000
   max-incomplete high 20000 one-minute low 2147483647
 one-minute high 2147483647
```

https://t.me/learningnets

```
  tcp reset-PSH disabled
```

Verify the applied parameter-maps for both default and user-configured class-maps:

```
R1#show  policy-firewall config zone-pair INSIDE_TO_OUTSIDE
Zone-pair              : INSIDE_TO_OUTSIDE
Source Zone            : INSIDE
Destination Zone       : OUTSIDE
Service-policy inspect : INSIDE_TO_OUTSIDE Class-map : PROTOCOLS(match-any)
   Match protocol icmp
   Match protocol telnet
  Action : inspect Parameter-map : PARAMETER
 Action : Police rate 256000 burst 16000
 Class-map : class-default(match-any)
   Match any
  Action : drop log Parameter-map : Default
```

Generate an ICMP packet flood from the **INSIDE** zone, to exceed the policing rate:

```
SW2#ping 150.1.9.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.9.9, timeout is 2 seconds: !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!SW2#ping 150.1.9.9 size 1400 timeout 0 repeat 1000

Type escape sequence to abort.
Sending 1000, 1400-byte ICMP Echos to 150.1.9.9, timeout is 0 seconds:
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..........................................................................
..................
```

https://t.me/learningnets

```
Success rate is 0 percent (0/1000)
!
!R1:
%FW-6-DROP_PKT: Dropping icmp session 136.1.19.10:0 150.1.9.9:0 on zone-pair INSIDE_TO_OUTSIDE class PROTOCOLS due t
!
!R1#show policy-firewall stats zone-pair INSIDE_TO_OUTSIDE

policy exists on zp INSIDE_TO_OUTSIDE
  Zone-pair: INSIDE_TO_OUTSIDE

  Service-policy inspect : INSIDE_TO_OUTSIDE

    Class-map: PROTOCOLS (match-any)
      Match: protocol icmp
        2 packets, 1460 bytes
        30 second rate 0 bps
      Match: protocol telnet
        0 packets, 0 bytes
        30 second rate 0 bps

   Inspect
        Packet inspection statistics [process switch:fast switch]
        icmp packets: [0:40]

        Session creations since subsystem startup or last reset 2
        Current session counts (estab/half-open/terminating) [0:0:0]
        Maxever session counts (estab/half-open/terminating) [1:1:0]
        Last session created 00:01:43
        Last statistic reset never
        Last session creation rate 0
        Maxever session creation rate 2
        Last half-open session total 0
        TCP reassembly statistics
        received 0 packets out-of-order; dropped 0
        peak memory usage 0 KB; current usage: 0 KB
        peak queue length 0
 Police
        rate 256000 bps,16000 limit conformed 40 packets, 43592 bytes; actions: transmit
exceeded 997 packets, 1409854 bytes; actions: drop

        conformed 0 bps, exceed 20000 bps

    Class-map: class-default (match-any)
      Match: any
```

https://t.me/learningnets

```
Drop
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF High Availability

Task

- Configure two security zones on R1: VLAN 128 member of **INSIDE** and VLAN 123 member of **OUTSIDE**.
- Configure R1 and R2 for zone-based firewall stateful failover and bidirectionally inspect HTTP/TELNET protocols.
- Ensure that R1 is preferred as the **ACTIVE** device as long as it has both interfaces in the **UP** state.
- If an additional interface is required, use GigabitEthernet0/1.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Zone-Based Policy Firewall supports a High Availability configuration, just like CBAC, but there are some differences:

- It does not rely on HSRP.
- It supports asymmetric routing, where traffic can exit through the **ACTIVE** router and return through the **STANDBY** router.
- Session state is replicated only for **inspect** actions.
- Session state is replicated between **ACTIVE** and **STANDBY** only for layer 4

https://t.me/learningnets

protocols (TCP/UDP), so not for ICMP or for layer 7 inspections (applications inspection).

ZBPF HA is somewhat similar with ASA failover in that it requires some dedicated interfaces:

- **control link** is used for router role initial negotiation, **ACTIVE** or **STANDBY**, and router availability through hello packets.
- **data link** is used to synchronize/replicate the session states from **ACTIVE** to **STANDBY**.
- **interface interlink** is used for asymmetric routing scenario, for the **STANDBY** device to route traffic back to the **ACTIVE device**.

The three roles of control-plane interfaces required for ZBPF HA can all be assigned to the same interface, but it has to be an interface not assigned to any RRI (redundant interface identifier) group, so an interface not part of any zone. Additionally, interfaces used for these roles need IPv4 addresses assigned so that IPv4 packet processing is enabled.

IOS router supports a configuration maximum of two redundancy groups, which are similar to the failover groups from the ASA. Identically, one router can be **ACTIVE** for first redundancy group and the second router can be **ACTIVE** for second redundancy group, thus achieving an **ACTIVE/ACTIVE** scenario in which each router is actually forwarding traffic. The **ACTIVE** router is selected per redundancy group in a two-step process, just like with HSRP:

- The device with highest priority takes the **ACTIVE** role; if priority is equal, proceed to the next step.
- The device with highest IP address for the control link takes the **ACTIVE** role.
- Preemption works only if **ACTIVE** device was selected based on highest priority, not based on highest IP address.
- The default priority is 100 if not configured.

From a configuration perspective, the following are necessary to be configured on both routers:

- Enable redundancy within the **global** parameter-map.
- For **ACTIVE/STANDBY** mode, configure one redundancy group; for **ACTIVE/ACTIVE** mode, configure two redundancy groups.
- Within each redundancy group, configure the name and device priority and specify control/data link interfaces.
- Map interfaces that will be assigned to zones such as RRI (redundant interface

https://t.me/learningnets

identifier); it is mandatory use the same RRI number for common subnet interfaces between the two routers.

- Configure a VIP (also a VMAC will be automatically assigned) for the redundancy group to at least one of the zone interfaces.
- Configure ZBPF settings and policies.

> ZBPF HA is supported for both IPv4 and IPv6 traffic, but only in routed mode.

Configuration

Configure R1 and R2 for High Availability:

SW2:

```
interface range fastEthernet1/0/1 - 2
 switchport trunk allowed vlan add 1
```

R1:

```
interface GigabitEthernet0/1
 ip address 136.1.1.1 255.255.255.0
!
parameter-map type inspect global
 redundancy
 exit
!
redundancy
 application redundancy
  group 1
   name ZBPF
   preempt
   priority 150
   control GigabitEthernet0/1 protocol 1
   data GigabitEthernet0/1
!
interface GigabitEthernet0/1.123
 redundancy rii 10 decrement 60
!
interface GigabitEthernet0/1.128
 redundancy rii 20 decrement 60
 redundancy group 1 ip 136.1.128.12 exclusive
```

https://t.me/learningnets

R2:

```
interface GigabitEthernet0/1
 ip address 136.1.1.2 255.255.255.0
!
parameter-map type inspect global
 redundancy
 exit
!
redundancy
 application redundancy
  group 1
   name ZBPF
   preempt
   priority 100
   control GigabitEthernet0/1 protocol 1
   data GigabitEthernet0/1
!
interface GigabitEthernet0/1.123
 redundancy rii 10
!
interface GigabitEthernet0/1.128
 redundancy rii 20
 redundancy group 1 ip 136.1.128.12 exclusive
```

Identically configure R1 and R2 for ZBPF:

R1: and R2:

```
class-map type inspect match-any ZBPF
 match protocol telnet
 match protocol http
!
policy-map type inspect INSIDE_TO_OUTSIDE
 class ZBPF
  inspect
!
policy-map type inspect OUTSIDE_TO_INSIDE
 class ZBPF
  inspect
!
!
zone security INSIDE
zone security OUTSIDE
!
```

https://t.me/learningnets

```
!
zone-pair security INSIDE_TO_OUTSIDE source INSIDE destination OUTSIDE
 service-policy type inspect INSIDE_TO_OUTSIDE
!
zone-pair security OUTSIDE_TO_INSIDE source OUTSIDE destination INSIDE
 service-policy type inspect OUTSIDE_TO_INSIDE
!
!
interface GigabitEthernet0/1.128
 zone-member security INSIDE
!
interface GigabitEthernet0/1.123
 zone-member security OUTSIDE
```

Verification

Verify the operational state of the redundancy application group as well as router roles:

```
R1#show redundancy application group 1
Group ID:1
Group Name:ZBPF

Administrative State: No Shutdown
Aggregate operational state : UpMy Role: ACTIVE
Peer Role: STANDBY
Peer Presence: Yes
Peer Comm: Yes
Peer Progression Started: Yes

RF Domain: btob-one
         RF state: ACTIVE
         Peer RF state: STANDBY HOT
!
!R2#show redundancy application group 1
Group ID:1
Group Name:ZBPF

Administrative State: No Shutdown
Aggregate operational state : UpMy Role: STANDBY
Peer Role: ACTIVE
Peer Presence: Yes
Peer Comm: Yes
```

https://t.me/learningnets

```
Peer Progression Started: Yes

RF Domain: btob-one
         RF state: STANDBY HOT
         Peer RF state: ACTIVE
```

Verify the control link interfaces:

```
R1#show redundancy application control-interface group 1
The control interface for rg[1] is GigabitEthernet0/1
Interface is Control interface associated with the following protocols: 1
BFD Enabled
Interface Neighbors:Peer: 136.1.1.2 Standby RGs: 1 BFD handle: 0
!
!R2#show redundancy application control-interface group 1
The control interface for rg[1] is GigabitEthernet0/1
Interface is Control interface associated with the following protocols: 1
BFD Enabled
Interface Neighbors:Peer: 136.1.1.1 Active RGs: 1 BFD handle: 0
```

Verify the complete settings for the application redundancy, and notice that the default hello and hold timers match the default HSRP timers as well:

```
R1#show redundancy application protocol group 1

RG Protocol RG 1
------------------ Role: Active
        Negotiation: Enabled Priority: 150
 Protocol state: Active
 Ctrl Intf(s) state: Up
 Active Peer: Local
 Standby Peer: address 136.1.1.2, priority 100, intf Gi0/1
        Log counters:
                role change to active: 4
                role change to standby: 5
                disable events: rg down state 1, rg shut 0
                ctrl intf events: up 5, down 4, admin_down 1
                reload events: local request 0, peer request 3

RG Media Context for RG 1
-------------------------- Ctx State: Active
        Protocol ID: 1
        Media type: Default Control Interface: GigabitEthernet0/1
```

https://t.me/learningnets

```
        Current Hello timer: 3000 Configured Hello timer: 3000, Hold timer: 10000
Peer Hello timer: 3000, Peer Hold timer: 10000

        Stats:
                Pkts 453, Bytes 28086, HA Seq 0, Seq Number 453, Pkt Loss 0
                Authentication not configured
                Authentication Failure: 0
                Reload Peer: TX 0, RX 0
                Resign: TX 1, RX 1
        Standby Peer: Present. Hold Timer: 10000
                Pkts 208, Bytes 7072, HA Seq 0, Seq Number 552, Pkt Loss 0
```

Verify that common subnet interfaces on R1 and R2 are assigned to the same RRI group, and that R1 will lower its priority if any of its interfaces is down so that R2 can become **ACTIVE**:

```
R1#show redundancy rii
 No. of RIIs in database: 2
 Interface                        RII Id     decrement   GigabitEthernet0/1.123        :
10              60
  GigabitEthernet0/1.128        :20              60
!
!R2#show redundancy rii
 No. of RIIs in database: 2
 Interface                        RII Id     decrement   GigabitEthernet0/1.123        :10              0
  GigabitEthernet0/1.128        :20              0
```

Initiate a telnet session from **INSIDE** to **OUTSIDE** and verify session that state is replicated from **ACTIVE** to **STANDBY**:

```
SW1#telnet 150.1.4.4
Trying 150.1.4.4 ... Open



User Access Verification

Password: R4>
!
!R1#show policy-firewall session zone-pair INSIDE_TO_OUTSIDE
 Zone-pair: INSIDE_TO_OUTSIDE
  Service-policy inspect : INSIDE_TO_OUTSIDE
    Class-map : ZBPF(match-any) Established Sessions =  1
 Session 2CF740A0 (136.1.128.9:18976)=>(150.1.4.4:23) telnet:tcp SIS_OPEN/TCP_ESTAB
```

https://t.me/learningnets

```
          Created 00:00:14, Last heard 00:00:12
          Bytes sent (initiator:responder) [37:74] HA State: ACTIVE, RG ID: 1
    Class-map : class-default(match-any)
!
!R2#show policy-firewall session zone-pair INSIDE_TO_OUTSIDE
 Zone-pair: INSIDE_TO_OUTSIDE
  Service-policy inspect : INSIDE_TO_OUTSIDE
    Class-map : ZBPF(match-any) Established Sessions =  1
 Session 312B9DC0 (136.1.128.9:18976)=>(150.1.4.4:23) telnet:tcp SIS_OPEN/TCP_ESTAB
          Created 00:00:31, Last heard never
          Bytes sent (initiator:responder) [0:0] HA State: STANDBY, RG ID: 1

    Class-map : class-default(match-any)
```

Initiate a telnet session from **OUTSIDE** to **INSIDE** and verify that session state is replicated from **ACTIVE** to **STANDBY**, even though there is no VIP configured in the reverse direction:

https://t.me/learningnets

```
R4#telnet 150.1.9.9

Trying 150.1.9.9 ... Open


User Access Verification

Password: SW1>

!
!R1#show policy-firewall session Zone-pair OUTSIDE_TO_INSIDE

 Zone-pair: OUTSIDE_TO_INSIDE
  Service-policy inspect : OUTSIDE_TO_INSIDE
    Class-map : ZBPF(match-any) Established Sessions =  1
 Session 2CF74420 (136.1.99.4:58486)=>(150.1.9.9:23) telnet:tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:18, Last heard 00:00:17
          Bytes sent (initiator:responder) [40:69] HA State: ACTIVE, RG ID: 1

    Class-map : class-default(match-any)
!
!R2#show policy-firewall session zone-pair OUTSIDE_TO_INSIDE

 Zone-pair: OUTSIDE_TO_INSIDE
  Service-policy inspect : OUTSIDE_TO_INSIDE
    Class-map : ZBPF(match-any) Established Sessions =  1
 Session 312BA140 (136.1.99.4:58486)=>(150.1.9.9:23) telnet:tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:32, Last heard never
          Bytes sent (initiator:responder) [0:0] HA State: STANDBY, RG ID: 1

    Class-map : class-default(match-any)
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF User Based Firewall Using ACS

Task

- Configure a zone-based firewall policy on R3:
  - Implicitly allow ICMP/TCP/UDP traffic between VLAN 123 and VLAN 34, and inspect ICMP/TCP/UDP traffic toward VLAN 27.
  - ICMP/TCP/UDP traffic from VLAN 27 should be authenticated first through the ACS server.
  - R3 authenticates users from VLAN 27 via telnet.
  - Do not use any ACLs for this task.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The user-based firewall is a combination of zone-based firewall with authentication-proxy services. The final goal is to allow traffic flow between zones according to ZBPF configuration, but only for IP addresses (users) that were previously authenticated. Authentication can be forced only for traffic between certain zones and can be unidirectional as well (for example, if initiated from zone A to zone B it is allowed, but if initiated from zone B to zone A it requires authentication before being allowed). Because the firewall configuration dictates what traffic is allowed and how (pass or inspect), the authentication proxy configuration no longer requires an initial

https://t.me/learningnets

ACL to be applied inbound on the interface receiving traffic (which dictates what traffic is allowed before authentication) or an ACL to be received from the RADIUS server (which dictates what traffic was allowed after authentication). Because of this, firewall configuration becomes more scalable and manageable when dealing with a huge number of users, because instead of downloading an ACL for each user, if users can be divided into groups, you'll configure a ZBPF configuration per group at router level and use RADIUS only for authentication purposes.

This is how it works:

1. The client initiates telnet/http/ftp traffic toward any destination.
2. The router intercepts telnet/http/ftp traffic (based on the auth-proxy configuration) and prompts the user to authenticate.
3. Authentication is done by a remote RADIUS server that returns in the Access-Accept message two Cisco AV Pair values, the privilege level of 15 and the group name used in the firewall configuration.
4. The router uses the received group name in its ZBPF configuration, basically activating firewall rules for the respective group; however, traffic is allowed only for the authenticated IP address.

The group name received from the RADIUS server (which can be ISE or ACS) is in the format of Cisco AV Pair, with a value of `supplicant-group=<group_name>`. Except for the regular auth-proxy and ZBPF configuration, the command that ties these two together is configured at the class-map level and matches on the group name received from the RADIUS server; it is `match user-group <group_name>`. This is the new and recommended method for configuring user-based firewall, because it is easier to understand and configure.

Configuration

R3:

```
aaa new-model
aaa authentication login CONSOLE none
aaa authentication login default group ACS
aaa authorization auth-proxy default group ACS
!
line console 0
 login authentication CONSOLE
!
aaa group server radius ACS
 server-private 172.16.1.100 key cisco
 ip radius source-interface Loopback0
```

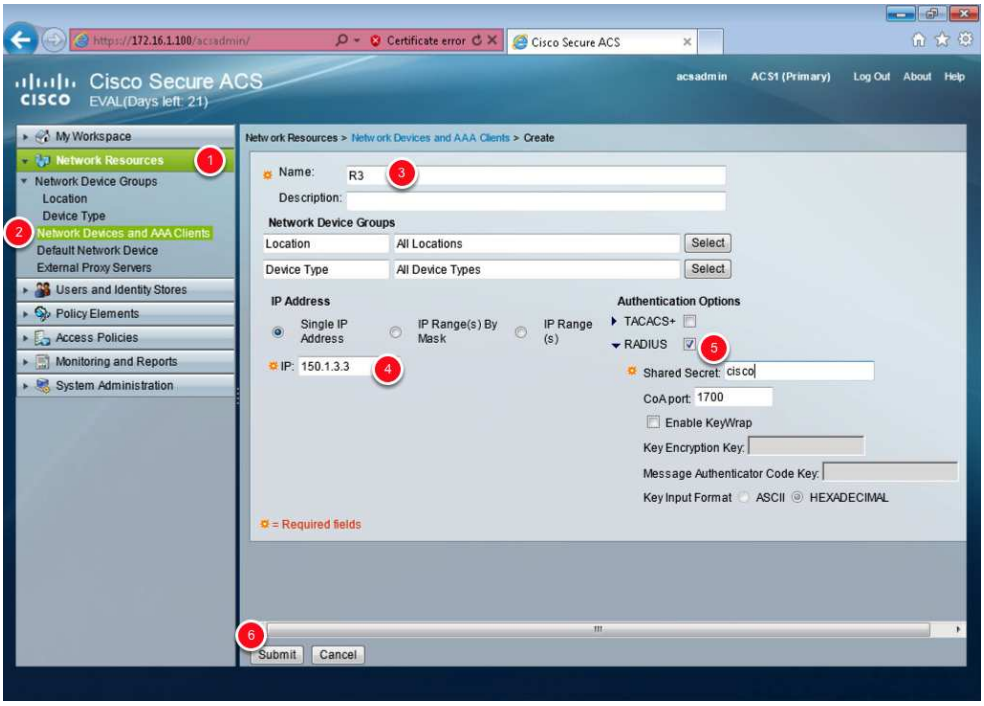
https://t.me/learningnets

```
!
ip admission name RADIUS_ACS proxy telnet
!
zone security TRUSTED
zone security UNTRUSTED
!
class-map type inspect match-any PROTOCOLS
 match protocol tcp
 match protocol udp
 match protocol icmp
!
class-map type inspect match-all USER_GROUP
 match user-group CCIE_GROUP
!
class-map type inspect match-all PROTOCOLS_AUTHENTICATED
 match class-map PROTOCOLS
 match class-map USER_GROUP
!
policy-map type inspect TRUSTED_TO_UNTRUSTED
 class type inspect PROTOCOLS
  inspect
!
policy-map type inspect UNTRUSTED_TO_TRUSTED
 class type inspect PROTOCOLS_AUTHENTICATED
  inspect
!
zone-pair security TRUSTED_TO_UNTRUSTED source TRUSTED destination UNTRUSTED
 service-policy type inspect TRUSTED_TO_UNTRUSTED
!
zone-pair security UNTRUSTED_TO_TRUSTED source UNTRUSTED destination TRUSTED
 service-policy type inspect UNTRUSTED_TO_TRUSTED
!
interface FastEthernet0/0
 zone-member security UNTRUSTED
 ip admission RADIUS_ACS
!
interface FastEthernet0/1.34
 zone-member security TRUSTED
!
interface FastEthernet0/1.123
 zone-member security TRUSTED
!
interface BVI99
```

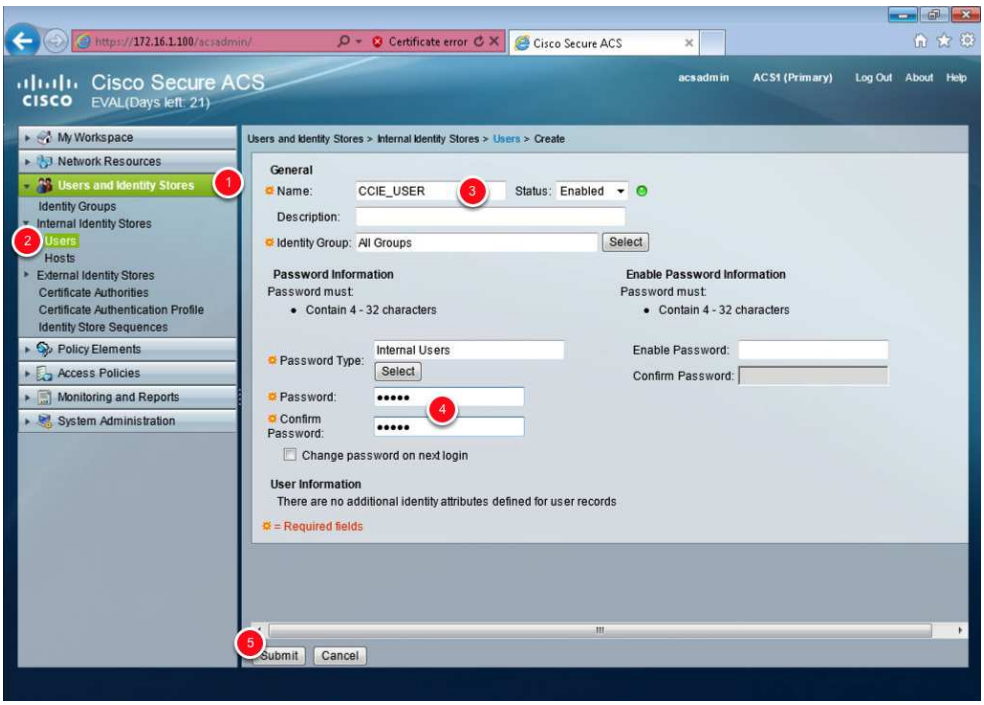
https://t.me/learningnets

```
zone-member security TRUSTED
```

Add R3 as a NAD device into ACS:

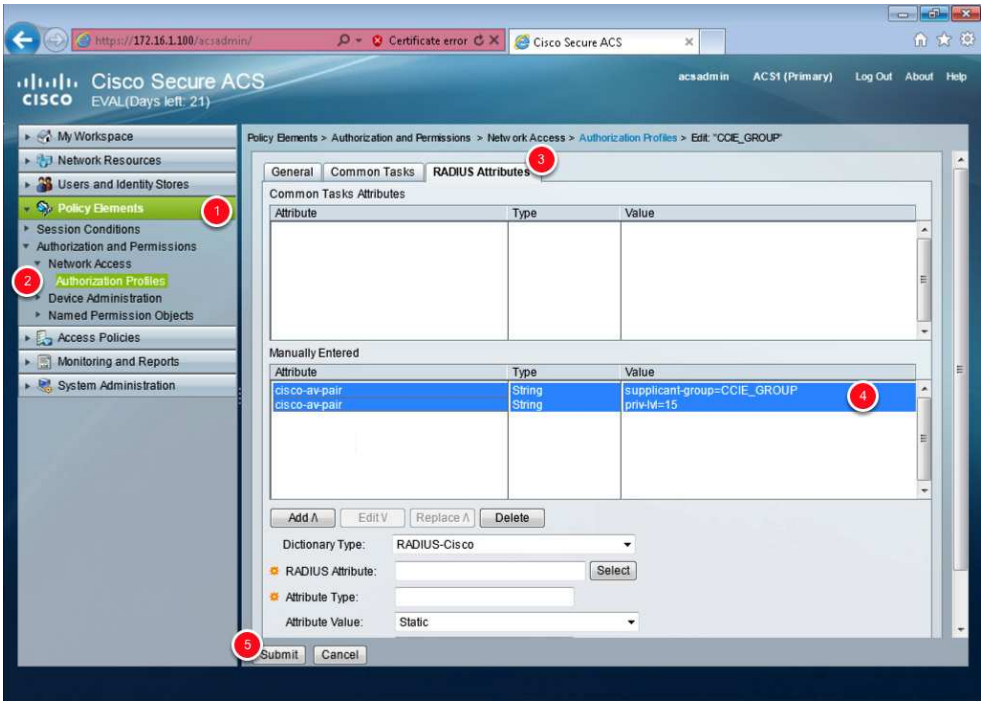


Configure a user in the ACS internal database, which will be used for authentication via auth-proxy service:

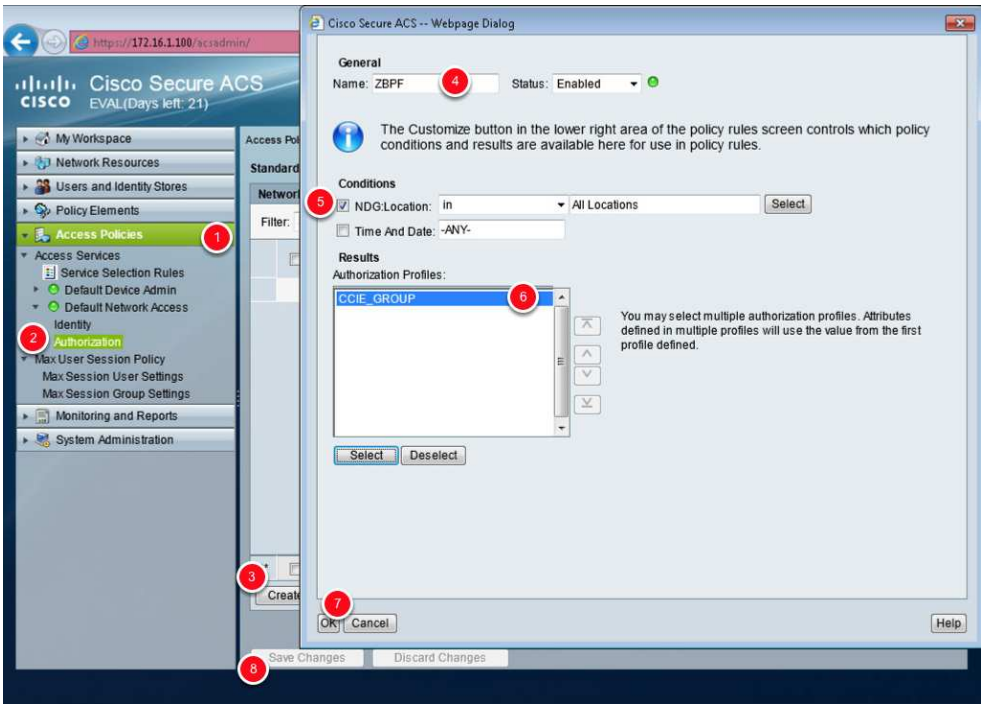


https://t.me/learningnets

Configure the RADIUS authorization profile required for the user-based firewall functionality, and assign priv-lvl of 15 and the group specified at the router level in the firewall policies:



Configure the RADIUS authorization policy, basically applying the previously configured profile:



Verification

Verify that traffic within the **TRUSTED** zone is allowed and traffic from the **TRUSTED**

https://t.me/learningnets

to the **UNTRUSTED** zone is inspected:

```
R2#telnet 150.1.4.4
Trying 150.1.4.4 ... Open

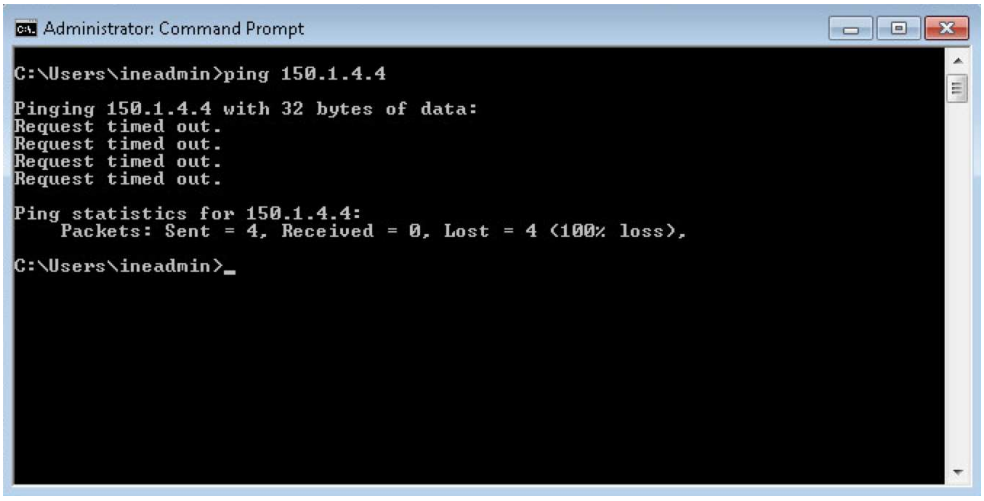


User Access Verification

Password: R4>
!
!R2#ping 136.1.27.51

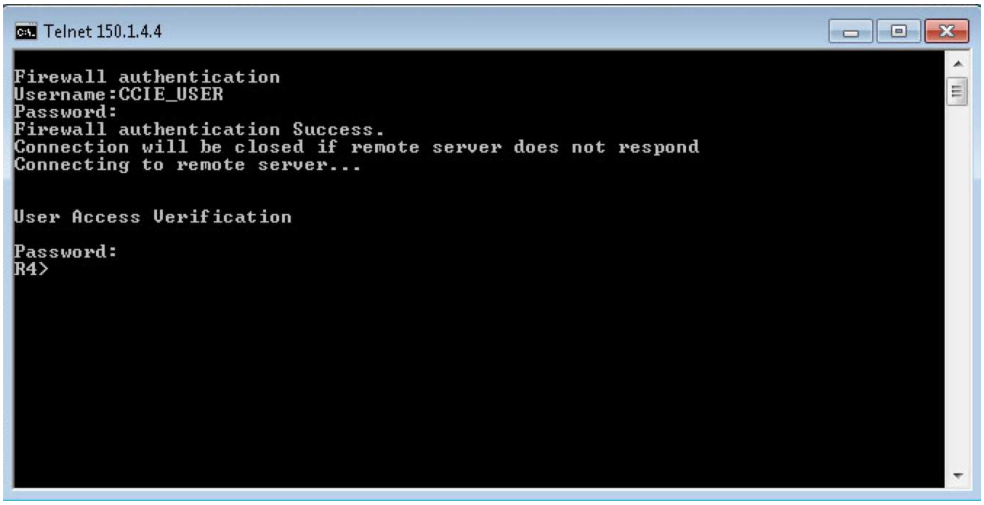
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.27.51, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/8 ms
!
!R3#show policy-firewall session
      Established Sessions = 1 Session 49F14DE0 (136.1.99.2:8)=>(136.1.27.51:0) icmp SIS_OPEN
          Created 00:00:02, Last heard 00:00:02
          ECHO request Bytes sent (initiator:responder) [360:360]
```

Verify that traffic from the **UNTRUSTED** to the **TRUSTED** zone does not work before user authentication:



Telnet from **TEST-PC-A** to any destination. For example R4, R3 will intercept the session and prompt for user authentication; if successfully authenticated, traffic from the **UNTRUSTED** zone to the **TRUSTED** zone should be allowed:

https://t.me/learningnets




```
R3#show policy-firewall session zone-pair UNTRUSTED_TO_TRUSTED
 Zone-pair: UNTRUSTED_TO_TRUSTED
  Service-policy inspect : UNTRUSTED_TO_TRUSTED
    Class-map : PROTOCOLS_AUTHENTICATED(match-all) Established Sessions =  1
 Session 49F15860 (136.1.27.51:1272)=>(150.1.4.4:23) tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:05, Last heard 00:00:03
          Bytes sent (initiator:responder) [38:65]
    Class-map : class-default(match-any)
```

Verify the authentication proxy cache:

```
R3#show ip auth-proxy cache username CCIE_USER
Authentication Proxy Cache
 User Name                : CCIE_USER
Client IP                 : 136.1.27.51
Client Port               : 1271
Timeout                   : 60
Time  Remaining           : 59              Connection state          : INTERCEPT

EPM information : Authproxy
Admission feature         : Authproxy
AAA Policies              :Supplicant-Group         : CCIE_GROUP


EOU information
--------------------------------------------------------------------------
Address         Interface              AuthType   Posture-Token Age(min)
--------------------------------------------------------------------------


EPM information : EOU
```

https://t.me/learningnets

```
Existing Firewall Sessions Information:

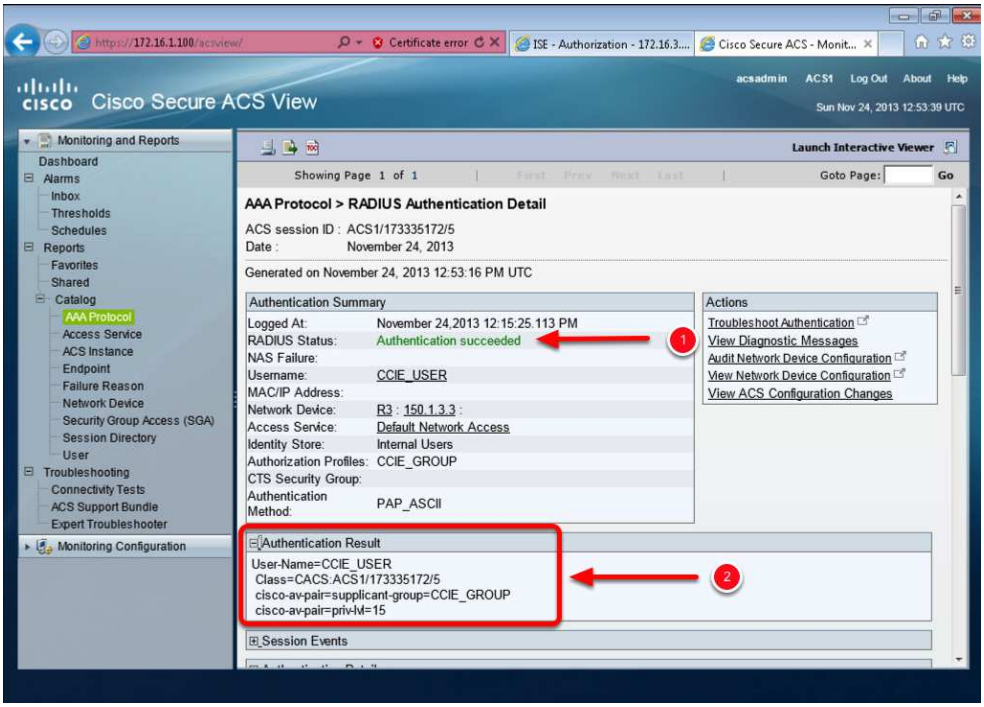
        Session 49F15160 (136.1.27.51:1272)=>(150.1.4.4:23) tcp SIS_OPEN/TCP_ESTAB
Created 00:01:15, Last heard 00:01:12
Bytes sent (initiator:responder) [32:65]
```

Verify the user group mappings for the client IP address:

```
R3#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.27.51         0148.f8b3.2e24.23       Nov 25 2013 12:00 AM    Automatic
!
!R3#show epm session ip 136.1.27.51
Admission feature       : Authproxy
AAA Policies            :Supplicant-Group         : CCIE_GROUP
!
!R3#show user-group
Usergroup : CCIE_GROUP
--------------------------------------------------------------------------------
User Name       Type     Interface              Learn         Age (min)
--------------------------------------------------------------------------------
136.1.27.51     IPv4     FastEthernet0/0        Dynamic          2
```

Verify the ACS RADIUS Authentication Logs:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF User Based Firewall Using ISE

Task

- Configure a zone-based firewall policy on R3:
  - Implicitly allow ICMP/TCP/UDP traffic between VLAN 123 and VLAN 34, and inspect ICMP/TCP/UDP traffic toward VLAN 27.
  - ICMP/TCP/UDP traffic from VLAN 27 should be authenticated first through the ISE server.
  - R3 authenticates users from VLAN 27 via http.
  - Do not use any ACLs for this task.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

This is the same configuration as the previous one, except that ISE is being used instead of ACS, and auth-proxy is achieved through HTTP, not telnet.

Configuration

R3:

```
aaa new-model
aaa authentication login CONSOLE none
```

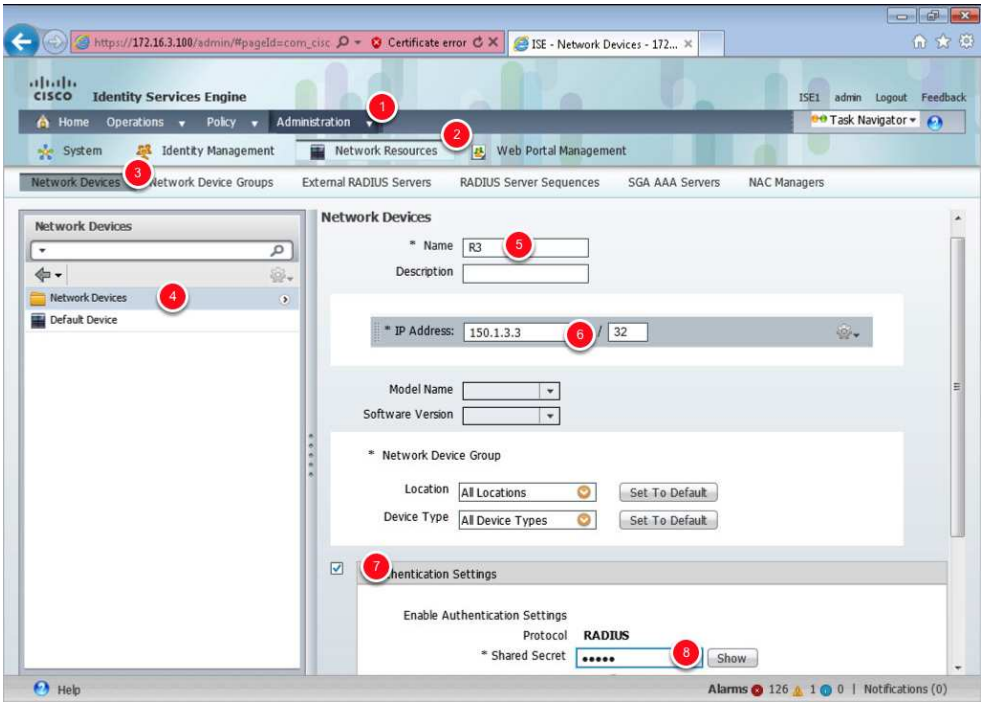
https://t.me/learningnets

```
aaa authentication login default group ISE
aaa authorization auth-proxy default group ISE
!
line console 0
 login authentication CONSOLE
!
aaa group server radius ISE
 server-private 172.16.3.100 key cisco
 ip radius source-interface Loopback0
!
ip admission name RADIUS_ISE proxy http
ip http server
!
zone security TRUSTED
zone security UNTRUSTED
!
class-map type inspect match-any PROTOCOLS
 match protocol tcp
 match protocol udp
 match protocol icmp
!
class-map type inspect match-all USER_GROUP
 match user-group CCIE_GROUP
!
class-map type inspect match-all PROTOCOLS_AUTHENTICATED
 match class-map PROTOCOLS
 match class-map USER_GROUP
!
policy-map type inspect TRUSTED_TO_UNTRUSTED
 class type inspect PROTOCOLS
  inspect
!
policy-map type inspect UNTRUSTED_TO_TRUSTED
 class type inspect PROTOCOLS_AUTHENTICATED
  inspect
!
zone-pair security TRUSTED_TO_UNTRUSTED source TRUSTED destination UNTRUSTED
 service-policy type inspect TRUSTED_TO_UNTRUSTED
!
zone-pair security UNTRUSTED_TO_TRUSTED source UNTRUSTED destination TRUSTED
 service-policy type inspect UNTRUSTED_TO_TRUSTED
!
interface FastEthernet0/0
 zone-member security UNTRUSTED
 ip admission RADIUS_ISE
!
```

https://t.me/learningnets

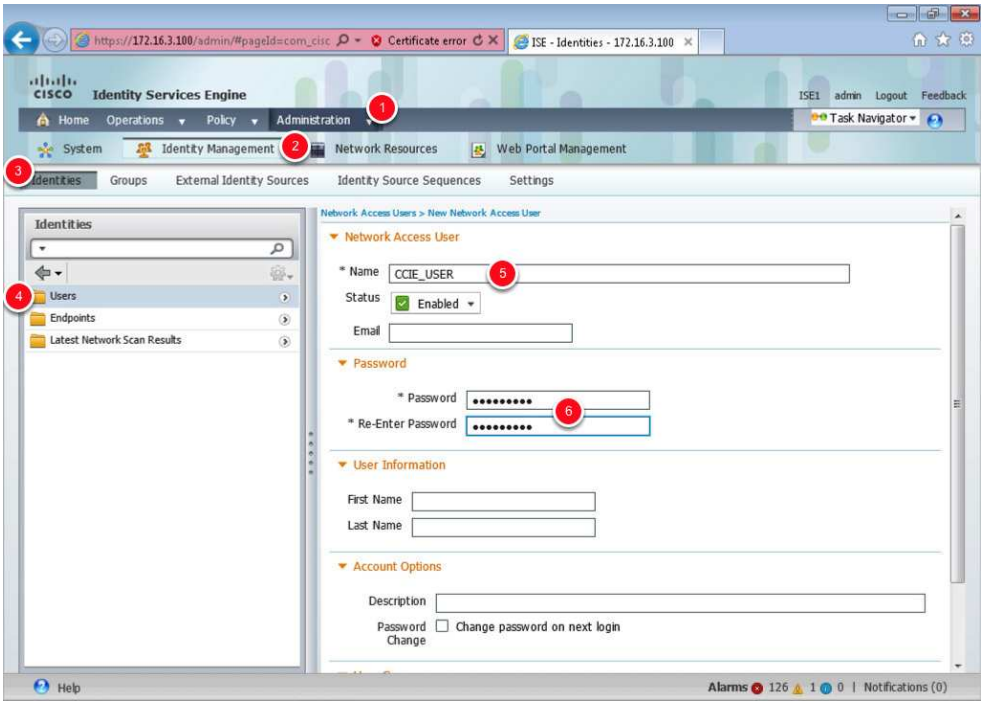
```
interface FastEthernet0/1.34
 zone-member security TRUSTED
!
interface FastEthernet0/1.123
 zone-member security TRUSTED
!
interface BVI99
 zone-member security TRUSTED
```

Add R3 as a NAD device into ISE:

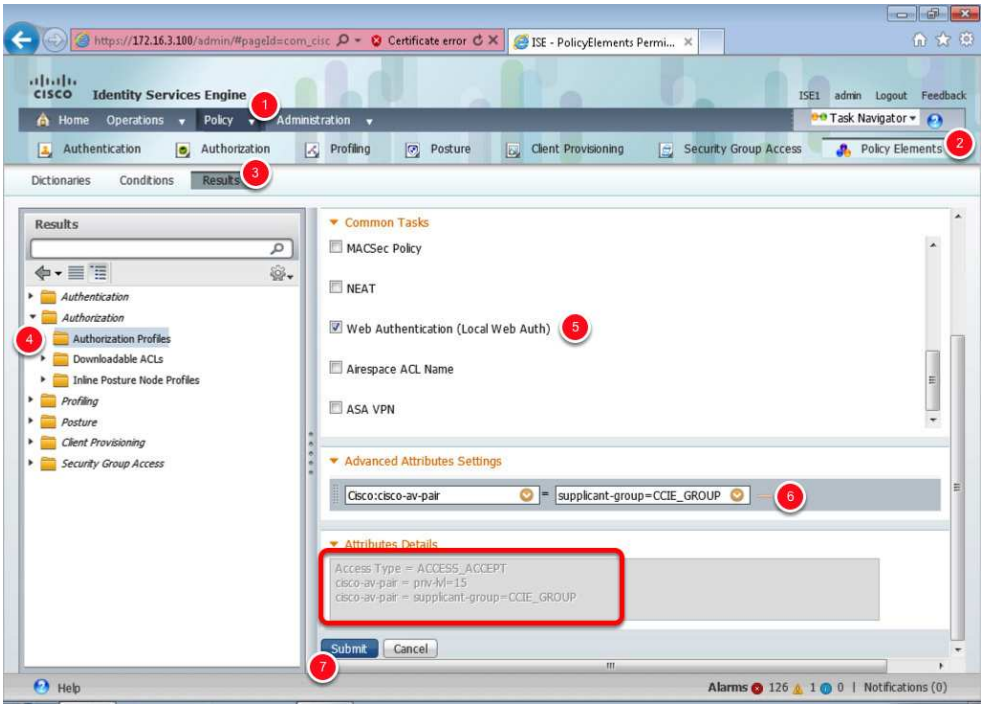


Configure a user in the ISE internal database, which will be used for authentication via auth-proxy service:

https://t.me/learningnets

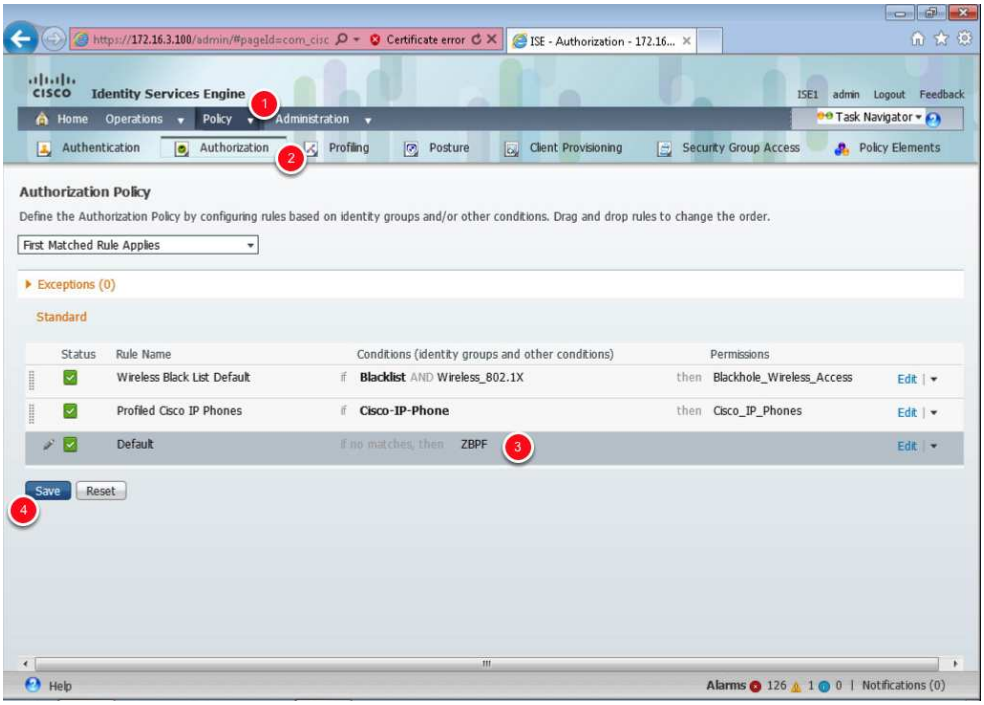


Configure the RADIUS authorization profile required for the user-based firewall functionality, and assign priv-lvl of 15 and the group specified at the router level in the firewall policies; note that the privilege level is assigned using the built-in macro **Web Authentication**. It could also be manually configured via the Cisco AV Pair.



Modify the default authorization policy or create a new one, basically applying the previously configured profile:

https://t.me/learningnets



Verification

Verify that traffic within the **TRUSTED** zone is allowed and traffic from the **TRUSTED** to the **UNTRUSTED** zone is inspected:

```
R2#telnet 150.1.4.4
Trying 150.1.4.4 ... Open


User Access Verification

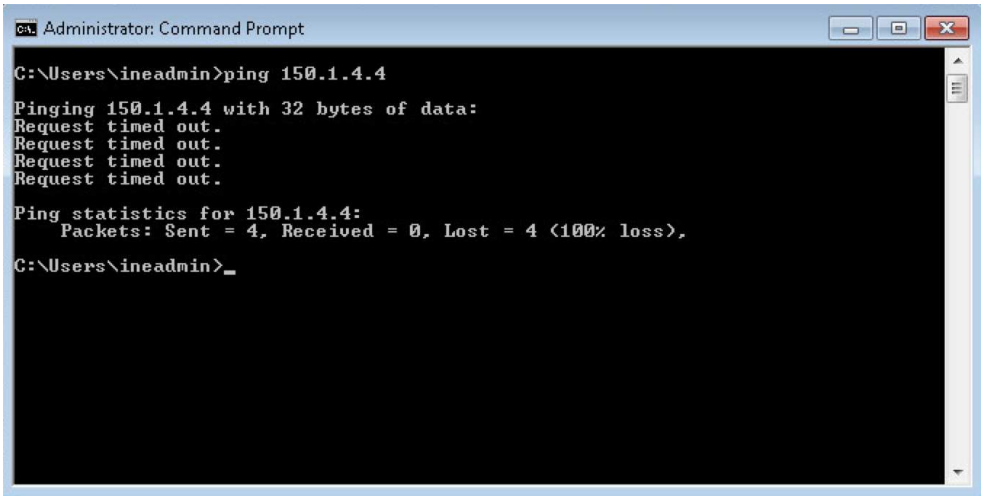
Password: R4>
!
!R2#ping 136.1.27.51

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.27.51, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/8 ms
!
!R3#show policy-firewall session
      Established Sessions = 1 Session 49F14DE0 (136.1.99.2:8)=>(136.1.27.51:0) icmp SIS_OPEN
          Created 00:00:02, Last heard 00:00:02
          ECHO request Bytes sent (initiator:responder) [360:360]
```

Verify that traffic from the **UNTRUSTED** to the **TRUSTED** zone does not work

https://t.me/learningnets

before user authentication:



Initiate an HTTP session from **TEST-PC-A** to any destination. For example R4, R3 will intercept the session and prompt for user authentication; if successfully authenticated, traffic from the **UNTRUSTED** zone to the **TRUSTED** zone should be allowed:

```
R3#show policy-firewall session zone-pair UNTRUSTED_TO_TRUSTED
 Zone-pair: UNTRUSTED_TO_TRUSTED
  Service-policy inspect : UNTRUSTED_TO_TRUSTED
    Class-map : PROTOCOLS_AUTHENTICATED(match-all) Established Sessions =  2
        Session 49F1A1E0 (136.1.27.51:1566)=>(150.1.4.4:80) tcp SIS_OPEN/TCP_CLOSEWAIT
          Created 00:00:34, Last heard 00:00:34
          Bytes sent (initiator:responder) [307:192]
Session 49F1A560 (136.1.27.51:1567)=>(150.1.4.4:80) tcp SIS_OPEN/TCP_ESTAB

          Created 00:00:34, Last heard 00:00:34
          Bytes sent (initiator:responder) [0:0]
    Class-map : class-default(match-any)
```

Verify the authentication proxy cache:

```
R3#show ip auth-proxy cache username CCIE_USER
Authentication Proxy Cache
 User Name                 : CCIE_USER
Client IP                  : 136.1.27.51
Client Port                : 1562
Timeout                    : 60
Time  Remaining            : 60              Connection state          : ESTAB

EPM information : Authproxy
Admission feature          : Authproxy
AAA Policies               :Supplicant-Group         : CCIE_GROUP
```

https://t.me/learningnets

```
EOU information
--------------------------------------------------------------------------
Address         Interface              AuthType   Posture-Token Age(min)
--------------------------------------------------------------------------


EPM information : EOU


Existing Firewall Sessions Information:

        Session 49F1A1E0 (136.1.27.51:1566)=>(150.1.4.4:80) tcp SIS_OPEN/TCP_CLOSEWAIT
Created 00:01:59, Last heard 00:00:59
Bytes sent (initiator:responder) [307:192]

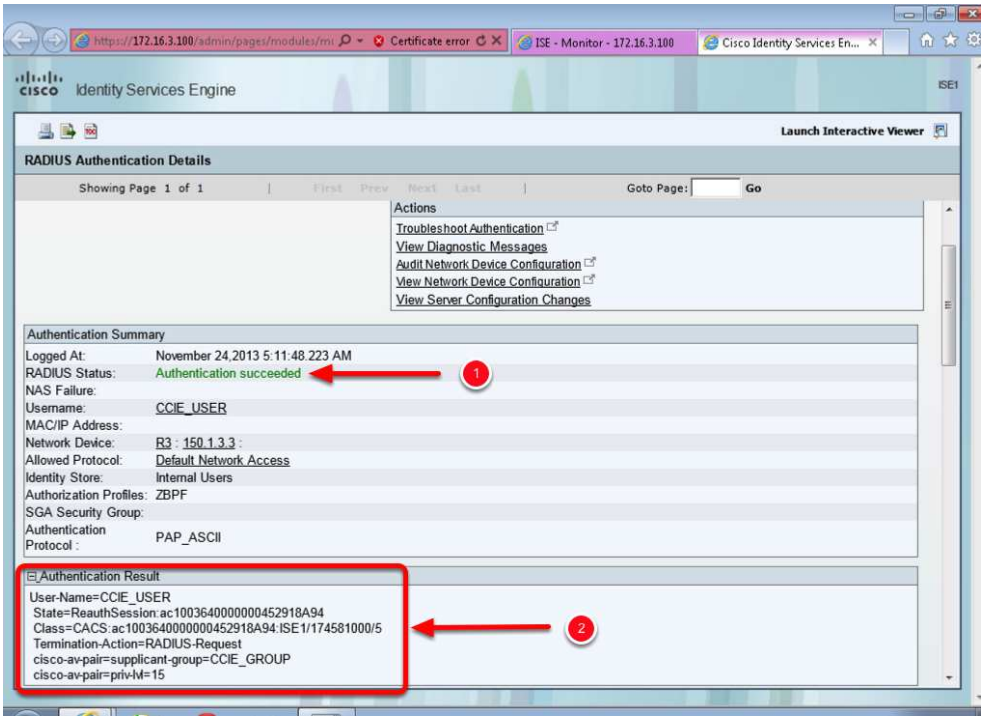
        Session 49F1A560 (136.1.27.51:1567)=>(150.1.4.4:80) tcp SIS_OPEN/TCP_ESTAB
Created 00:01:59, Last heard 00:00:59
Bytes sent (initiator:responder) [0:0]
```

Verify the user group mappings for the client IP address:

```
R3#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.27.51         0148.f8b3.2e24.23       Nov 25 2013 12:00 AM    Automatic
!
!R3#show epm session ip 136.1.27.51
Admission feature       : Authproxy
AAA Policies            :Supplicant-Group         : CCIE_GROUP
```

Verify the ISE Authentication logs:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Secure Group Firewall

Task

- Configure a zone-based firewall policy on R1:
  - Implicitly allow ICMP/TCP/UDP traffic between VLAN 19 and VLAN 128 and inspect ICMP/TCP/UDP traffic toward VLAN 123.
  - ICMP/TCP/UDP traffic from behind VLAN 123 should be authenticated first through the ISE server.
  - Firewall policies for traffic behind VLAN 123 are based on SGT; however, no SGT value should be received from ISE for R3's Loopback0.
  - R1 authenticates users from behind VLAN 123 via telnet.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Security Group Access (SGA) is a new vision of Cisco for scalable access control enforcement, moving away from the usual IP-based access-lists and firewall rules toward a more context and application aware policy. This works as follows:

- Each client/user/device is manually/statically or automatically/dynamically assigned a Security Group Tag (SGT) when traffic enters the network (at access switch level for wired access, wireless LAN controller for wireless, and ASA firewall for VPN).

https://t.me/learningnets

- If the device controlling access to the network is not hardware SGT capable (it cannot rewrite the layer 2 header and insert the SGT tag), but it is software capable, it can signal the IP-to-SGT mapping to an upstream capable device via a dedicated signaling protocol, namely SXP (SGT Exchange Protocol).
- TAG is preserved from hop-to-hop toward the destination being accessed (through MACsec), or, alternatively, the IP-to-SGT mapping is being dynamically learned by the device (switch/router/firewall) restricting access to protected resources.
- The resources being accessed, usually data-center level services, can be manually/statically assigned SGT values, because these devices will not authenticate to the network to receive an SGT; this allows you to build SGT-based firewall rules for both source and destination. Otherwise, the destination will remain IP based.
- Devices controlling access to resources (switch/router/firewall) will be statically configured with rules based on SGT values, or will download the ruleset from ISE.

SGT is a 16-bit value that will be inserted in the layer 2 header of the Ethernet frame and can be seen as a technology for labelling packets at layer 2, with the ultimate goal of having an SGT value for each source and/or destination IP address. For users/terminals and clients in general, the SGT will dynamically be received from ISE upon successful MAB, 802.1x, web authentication, or auth-proxy successful authentication (basically any authentication method available through RADIUS). For situations in which a device cannot be authenticated or the network device allowing access to the networks is not Cisco or SGT capable, a static mapping can be configured for it through multiple methods (based on the access device), one or more options being available:

- IP to SGT mapping
- VLAN to SGT mapping
- Subnet to SGT mapping
- Layer2 port to SGT mapping
- Layer3 port to SGT mapping

Basically, within an SGT deployment four types of devices are involved:

https://t.me/learningnets

- **Classification**, capable of doing the IP-to-SGT mapping.
- **Transport**, capable of tagging packets with SGT (hardware requirement) or announcing the IP-to-SGT mapping through SXP.
- **Enforcement**, capable of restricting network access based on the SGT values (SGACL on switches and SGFW on ASA and ISR).
- **Policy Management**, dynamically assigns SGT values and owns the SGACL table downloaded by SGACL-capable devices; this is only ISE.

ZBPF rules can be configured based on the SGT values using the class-map-level command `match security-group source tag <value>` ; except for this, the following is a regular auth-proxy and zone-based firewall configuration.

Configuration

R1:

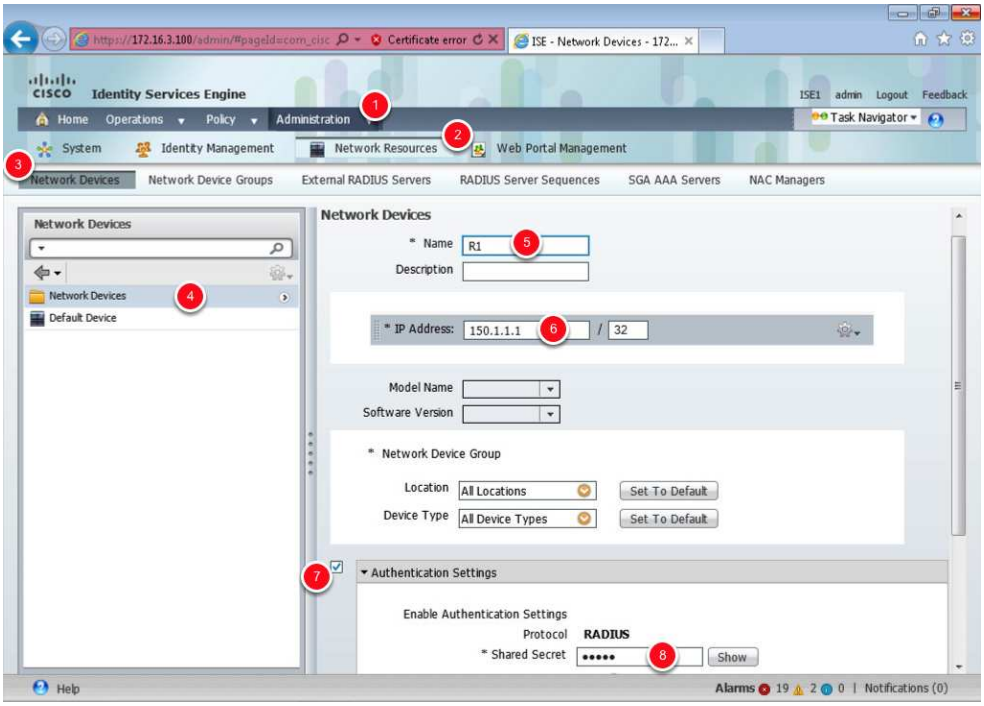
```
aaa new-model
aaa authentication login CONSOLE none
aaa authentication login default group ISE
aaa authorization auth-proxy default group ISE
!
line console 0
 login authentication CONSOLE
!
aaa group server radius ISE
 server-private 172.16.3.100 key cisco
 ip radius source-interface Loopback0
!
ip admission name RADIUS_ISE proxy telnet
!
cts role-based sgt-map host 150.1.3.3 sgt 10
!
zone security TRUSTED
zone security UNTRUSTED
!
class-map type inspect match-any PROTOCOLS
 match protocol tcp
 match protocol udp
 match protocol icmp
!
class-map type inspect match-any SECURE_GROUP
 match security-group source tag 2
 match security-group source tag 10
!
```

https://t.me/learningnets

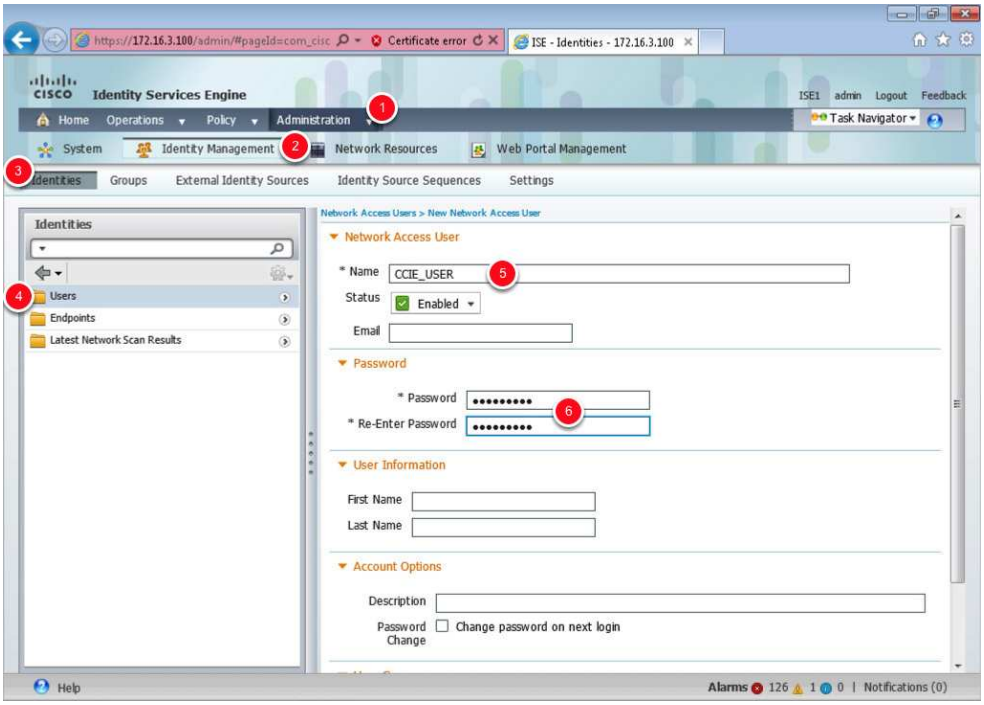
```
class-map type inspect match-all PROTOCOLS_AUTHENTICATED
 match class-map PROTOCOLS
 match class-map SECURE_GROUP
!
policy-map type inspect TRUSTED_TO_UNTRUSTED
 class type inspect PROTOCOLS
  inspect
!
policy-map type inspect UNTRUSTED_TO_TRUSTED
 class type inspect PROTOCOLS_AUTHENTICATED
  inspect
!
zone-pair security TRUSTED_TO_UNTRUSTED source TRUSTED destination UNTRUSTED
 service-policy type inspect TRUSTED_TO_UNTRUSTED
!
zone-pair security UNTRUSTED_TO_TRUSTED source UNTRUSTED destination TRUSTED
 service-policy type inspect UNTRUSTED_TO_TRUSTED
!
interface GigabitEthernet0/1.123
 zone-member security UNTRUSTED
 ip admission RADIUS_ISE
!
interface GigabitEthernet0/1.128
 zone-member security TRUSTED
!
interface GigabitEthernet0/0
 zone-member security TRUSTED
```

Configure R1 as a NAD device into ISE:

https://t.me/learningnets

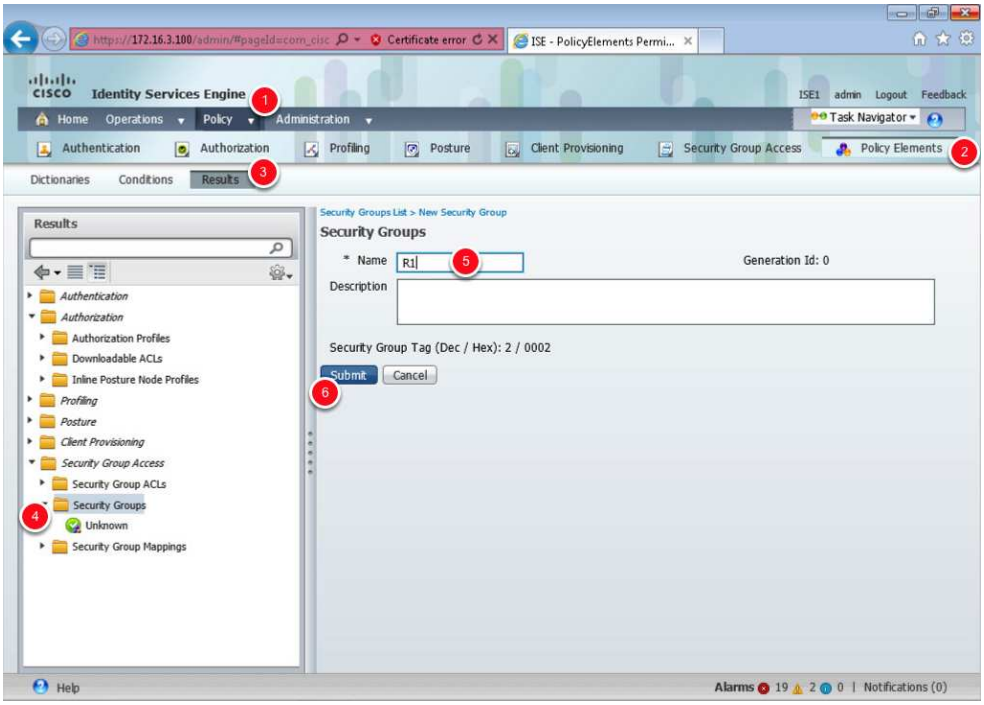


Configure a user in the ISE internal database, which will be used for authentication via auth-proxy service:

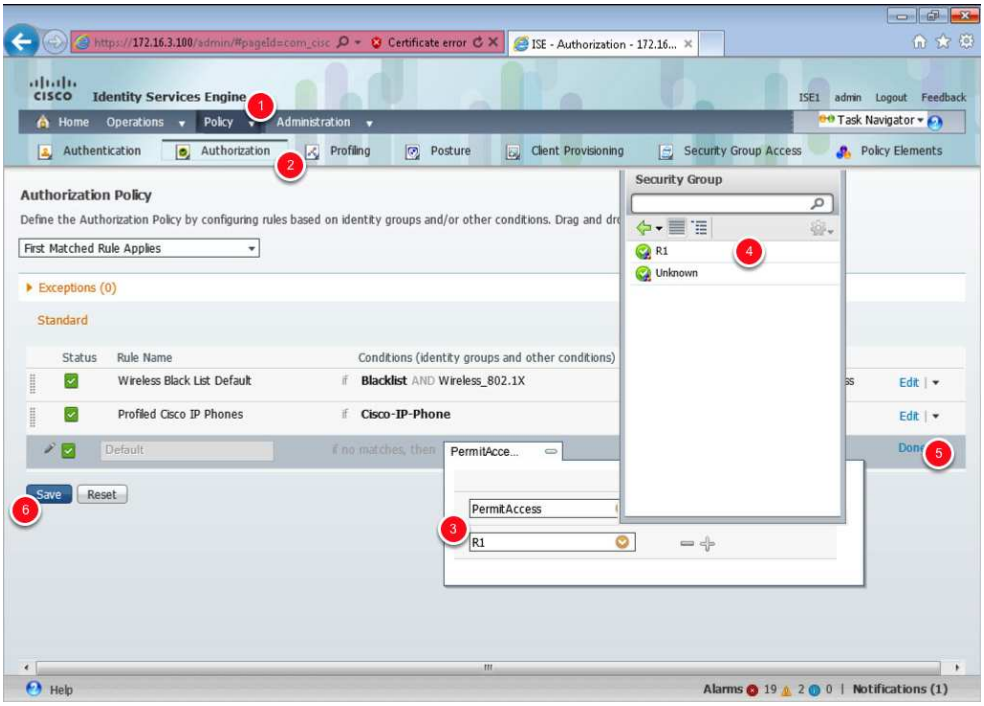


Create a new security group on ISE, and note that ISE automatically assigns the SGT value, in order from 1 to 65535:

https://t.me/learningnets



Configure a new authorization policy or modify the default one, and assign the previously configured security group as part of the authorization attributes:



Verification

Verify that traffic from the **TRUSTED** zone to the **UNTRUSTED** zone is allowed and inspected:

```
SW1#ping 136.1.27.51

Type escape sequence to abort.
```

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 136.1.27.51, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!R1#show policy-firewall session zone-pair TRUSTED_TO_UNTRUSTED
 Zone-pair: TRUSTED_TO_UNTRUSTED
  Service-policy inspect : TRUSTED_TO_UNTRUSTED
    Class-map : PROTOCOLS(match-any) Established Sessions =  1
 Session 316DD620 (136.1.128.9:8)=>(136.1.27.51:0) icmp SIS_OPEN

          Created 00:00:03, Last heard 00:00:03
          ECHO request
          Bytes sent (initiator:responder) [360:360]
    Class-map : class-default(match-any)
```

Verify that traffic from the **UNTRUSTED** zone to the **TRUSTED** zone is not allowed unless sourced from R3's Loopback0, which has a static SGT mapping that is used in the firewall configuration:

```
R3#ping 136.1.128.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.128.9, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R3#ping 136.1.128.9 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.128.9, timeout is 2 seconds:
Packet sent with a source address of 150.1.3.3 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
!
!R1#show policy-firewall session zone-pair UNTRUSTED_TO_TRUSTED
 Zone-pair: UNTRUSTED_TO_TRUSTED
  Serv:e-policy inspect : UNTRUSTED_TO_TRUSTED
    Class-map : PROTOCOLS_AUTHENTICATED(match-all) Established Sessions =  1
 Session 316DE420 (150.1.3.3:8)=>(136.1.128.9:0) icmp SIS_OPEN

          Created 00:00:02, Last heard 00:00:02
          ECHO request
          Bytes sent (initiator:responder) [360:360]
    Class-map : class-default(match-any)
```

Initiate a telnet session from R3; for example, to trigger the auth-proxy service,

https://t.me/learningnets

receive the SGT from ISE so that traffic will be allowed from this source as well:

```
R3#telnet 136.1.128.9
Trying 136.1.128.9 ... Open
 Firewall authentication
Username:CCIE_USER
Password:
Firewall authentication Success.
Connection will be closed if remote server does not respond
Connecting to remote server...

[Connection to 136.1.128.9 closed by foreign host]
!
!R3#telnet 136.1.128.9
Trying 136.1.128.9 ... Open



User Access Verification

Password: SW1>
!
!R1#show policy-firewall session zone-pair UNTRUSTED_TO_TRUSTED
 Zone-pair: UNTRUSTED_TO_TRUSTED
  Service-policy inspect : UNTRUSTED_TO_TRUSTED
    Class-map : PROTOCOLS_AUTHENTICATED(match-all) Established Sessions =  1
 Session 316DDD20 (136.1.99.3:31411)=>(136.1.128.9:23) tcp SIS_OPEN/TCP_ESTAB
          Created 00:00:26, Last heard 00:00:24
          Bytes sent (initiator:responder) [40:69]
    Class-map : class-default(match-any)
!
!R3#ping 136.1.128.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.128.9, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Verify the auth-proxy cache and the correct SGT received from ISE, as well as the statically configured one on R1:

```
R1#show ip auth-proxy cache username CCIE_USER
Authentication Proxy Cache

Authentication Method   : UnknownUser Name                : CCIE_USER
Client IP               : 136.1.99.3
```

https://t.me/learningnets

```
Client Port             : 20352
Timeout                 : 60
Time  Remaining         : 60              Connection state          : ESTAB

EPM information : Authproxy
    Admission feature:  AUTHPROXY AAA Policies:
SGT:  0002-0

EOU information
--------------------------------------------------------------------------
Address         Interface              AuthType   Posture-Token Age(min)
--------------------------------------------------------------------------


EPM information : EOU

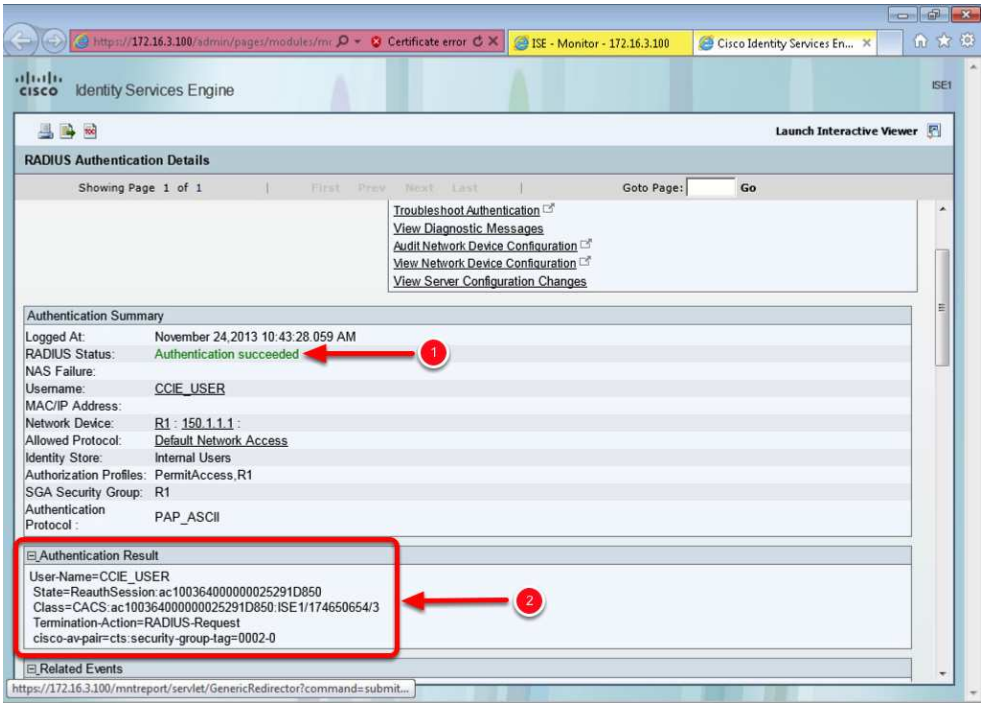

Existing Firewall Sessions Information:
!
!R1#show epm session ip 136.1.99.3
    Admission feature:  AUTHPROXY AAA Policies:
SGT:  0002-0
!
!R1#show cts role-based sgt-map all
Active IP-SGT Bindings Information

IP Address              SGT     Source
============================================136.1.99.3               2       LOCAL
150.1.3.3               10      CLI


IP-SGT Active Bindings Summary
============================================
Total number of CLI      bindings = 1
Total number of LOCAL    bindings = 1
Total number of active   bindings = 2
```

Verify the ISE Authentication logs:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

ZBPF Legacy User Based Firewall

Task

- Configure a zone-based firewall policy on R3:
  - Implicitly allow ICMP/TCP/UDP traffic between VLAN 123 and VLAN 34 and inspect ICMP/TCP/UDP traffic toward VLAN 27.
  - ICMP/TCP/UDP traffic from VLAN 27 should be authenticated first through the ACS server.
  - R3 authenticates users from VLAN 27 via telnet.
  - R3 should restrict all traffic from VLAN 27 only before authentication, but no ACL should be received from the ACS server.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The following configuration is also based on auth-proxy and zone-based firewall, but additional concepts are introduced, which complicates the overall implementation without any real added value, as we'll see. This is called the legacy method of implementing user-based firewall and is useful only in cases when the authentication RADIUS server cannot differentiate between multiple NAD devices and therefore provide different authorization services (mainly different dACL) based on the NAD doing the request, not only the user. This was the case with older ACS

https://t.me/learningnets

versions, such as 3.x and 4.x, but is no longer an issue with the new ACS 5.x, which follows a policy-oriented architecture and is therefore extremely granular from a configuration perspective. However, the main advantage of using this method is to allow an authenticated user access to certain resources and services based on an ACL; however, that ACL is configured locally (not received from the RADIUS server) and applied dynamically upon successful authentication.

Three new concepts are being introduced:

- A class-map of type control tag, configured with the command `class-map type control tag match-all <class_name>` ; its only scope is to match on the tag received from the RADIUS server as a Cisco AV Pair with the value `tag-name=<tag_value>` , using the command `match tag <tag_value> .`
- A policy-map of type control tag, configured with the command `policy-map type control tag <policy_name>` ; within the policy-map, multiple class-maps of type control can be specified (to match multiple tags), and each class-map will match on a set of attributes grouped through an identity policy using the command `identity policy <identity_name>` . The policy-map is attached to the auth-proxy global rule.
- One or more multiple identity policies can be configured using the command `identity policy <identity_name>` ; each one will be assigned the ACL governing user access after successful authentication using the command `access-group <acl_name>` and the user group that will be used in the zone-based firewall configuration using the command `user-group <group_name> .`

Basically in this setup, the user-group attribute used in the firewall configuration is locally configured (not received from RADIUS); the resulting firewall policy for an authenticated user will be based on the tag value received from RADIUS. This is an example of hierarchical configuration where:

- The TAG received from RADIUS is matched within a class-map.
- The policy-map makes use of the class-map and binds the TAG with an identity policy.
- The identity policy defines user authorization (ACL based) and links it to the firewall policy based on the locally defined user group value.
- The firewall policies are created based on the user group values, thus assigning each user different rules.

https://t.me/learningnets

Although this setup can be achieved with both ISE and ACS, given the fact that ISE and ACS solutions for user-based firewall were described in previous tasks, only one option is shown in here, namely with ACS. Except for the fact that ACS and ISE have different GUI and internal architecture, the end goal is the same: to assign the Cisco AV Pair attribute with the correct tag value as part of the authorization process, upon successful authentication.

Configuration

R3:

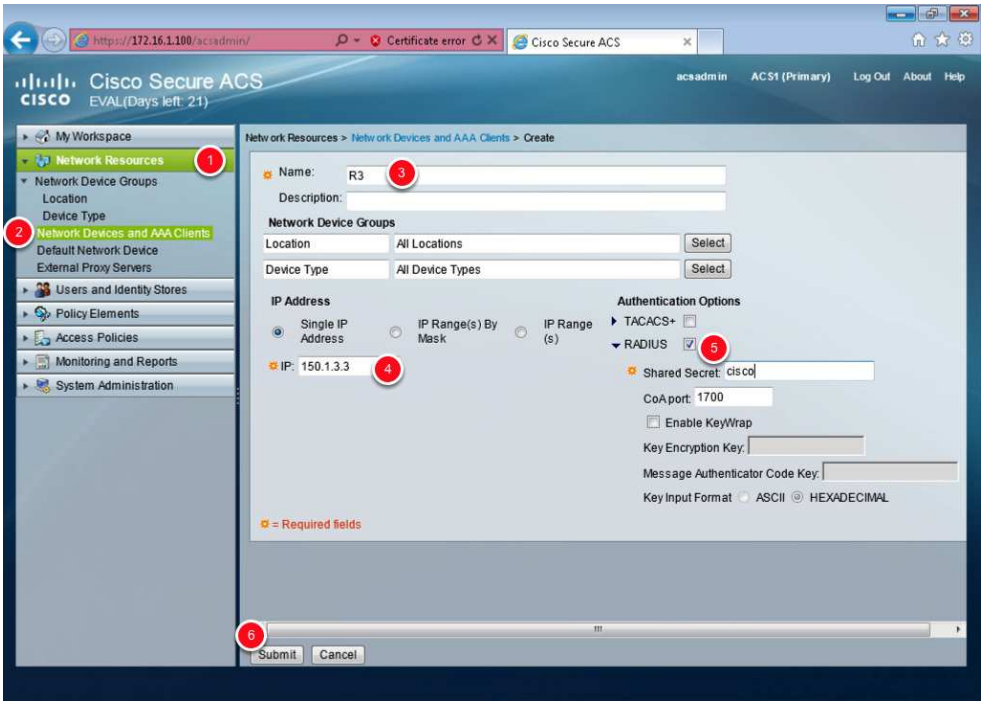
```
aaa new-model
aaa authentication login CONSOLE none
aaa authentication login default group ACS
aaa authorization auth-proxy default group ACS
!
line console 0
 login authentication CONSOLE
!
aaa group server radius ACS
 server-private 172.16.1.100 key cisco
 ip radius source-interface Loopback0
!
ip access-list extended DEFAULT_ACL
 permit tcp any any eq 23
!
ip access-list extended CCIE_USER_ACL
 permit ip any any
!
identity policy CCIE_USER_IDENTITY
 access-group CCIE_USER_ACL
 user-group CCIE_GROUP
!
class-map type control tag match-all CCIE_USER_CLASS
 match tag CCIE_USER_TAG
!
policy-map type control tag ADMISSION
 class type control tag CCIE_USER_CLASS
  identity policy CCIE_USER_IDENTITY
!
ip admission name RADIUS_ACS proxy telnet service-policy type tag ADMISSION
!
zone security TRUSTED
zone security UNTRUSTED
!
```

https://t.me/learningnets

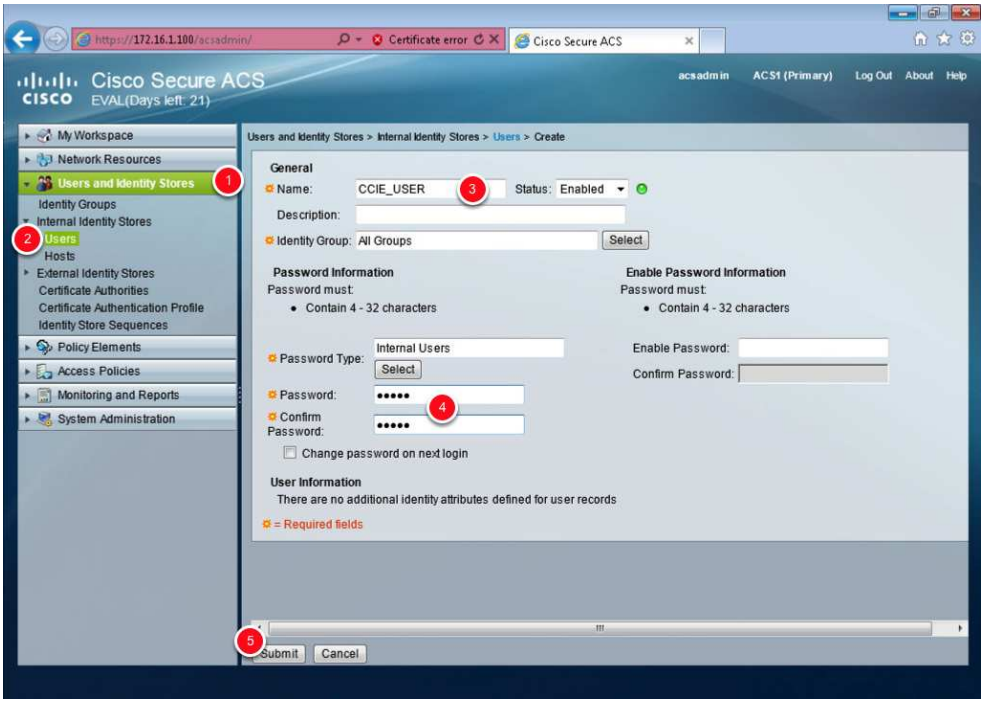
```
class-map type inspect match-any PROTOCOLS
 match protocol tcp
 match protocol udp
 match protocol icmp
!
class-map type inspect match-all USER_GROUP
 match user-group CCIE_GROUP
!
class-map type inspect match-all PROTOCOLS_AUTHENTICATED
 match class-map PROTOCOLS
 match class-map USER_GROUP
!
policy-map type inspect TRUSTED_TO_UNTRUSTED
 class type inspect PROTOCOLS
  inspect
!
policy-map type inspect UNTRUSTED_TO_TRUSTED
 class type inspect PROTOCOLS_AUTHENTICATED
  inspect
!
zone-pair security TRUSTED_TO_UNTRUSTED source TRUSTED destination UNTRUSTED
 service-policy type inspect TRUSTED_TO_UNTRUSTED
!
zone-pair security UNTRUSTED_TO_TRUSTED source UNTRUSTED destination TRUSTED
 service-policy type inspect UNTRUSTED_TO_TRUSTED
!
interface FastEthernet0/0
 zone-member security UNTRUSTED
 ip admission RADIUS_ACS
 ip access-group DEFAULT_ACL in
!
interface FastEthernet0/1.34
 zone-member security TRUSTED
!
interface FastEthernet0/1.123
 zone-member security TRUSTED
!
interface BVI99
 zone-member security TRUSTED
```

Add R3 as a NAD device into ACS:

https://t.me/learningnets

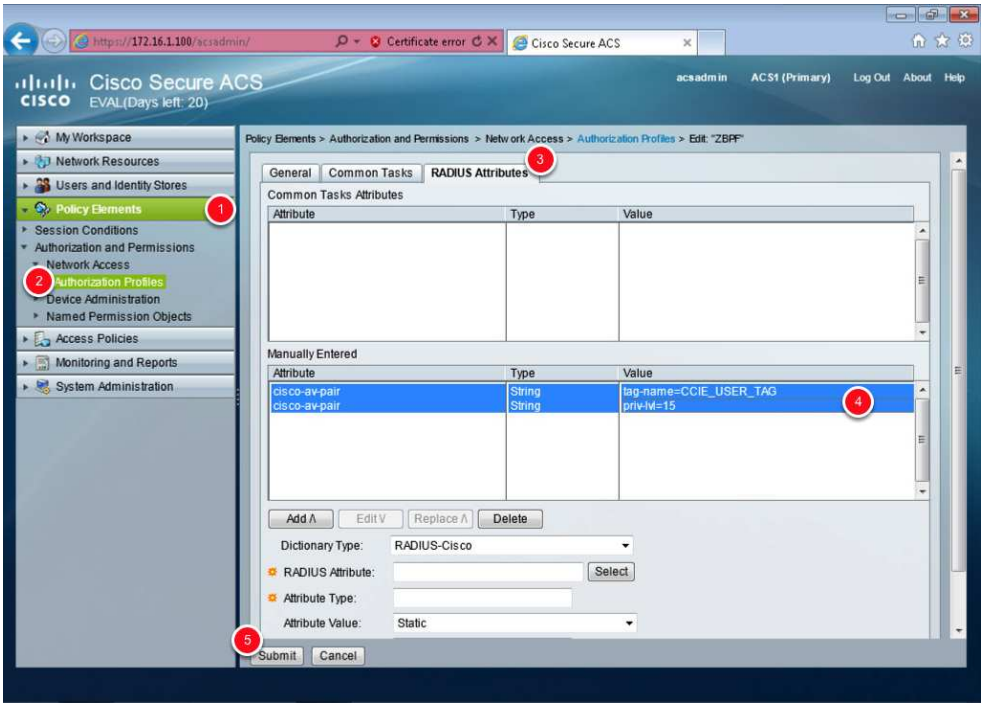


Configure a user in the ACS internal database, which will be used for authentication via auth-proxy service:

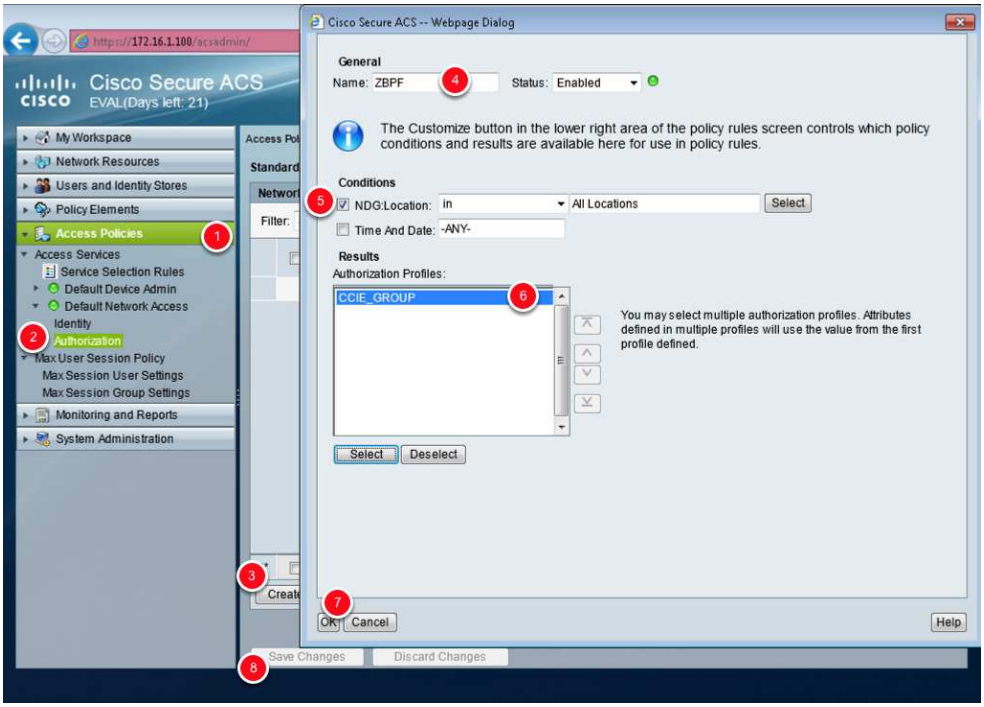


Configure the RADIUS authorization profile required for the user-based firewall functionality, assign priv-lvl of 15 and tag specified at the router level in the control tag class-map:

https://t.me/learningnets



Configure the RADIUS authorization policy, basically applying the previously configured profile:



Verification

Verify that traffic within the **TRUSTED** zone is allowed and traffic from **TRUSTED** to **UNTRUSTED** is inspected:

```
R2#telnet 150.1.4.4
Trying 150.1.4.4 ... Open
```

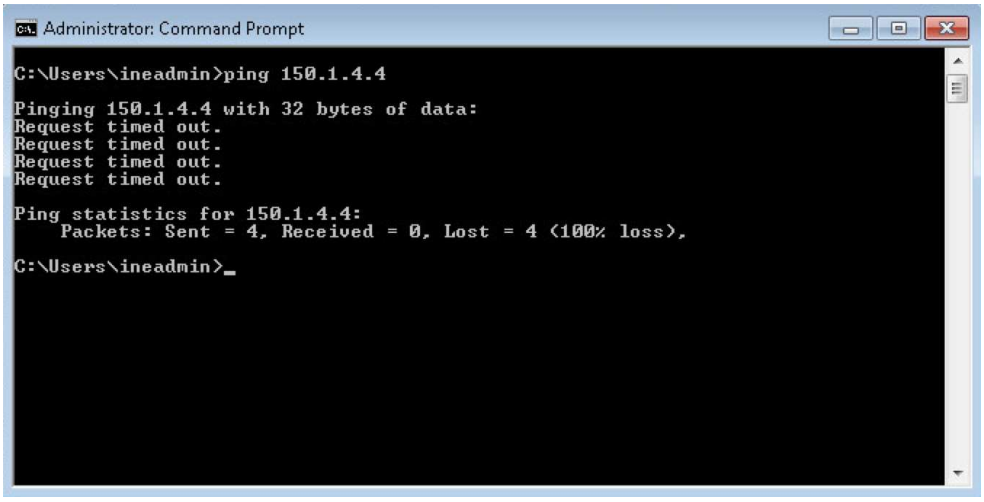
https://t.me/learningnets

```
User Access Verification

Password: R4>
!
!R2#ping 136.1.27.51

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.27.51, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/8 ms
!
!R3#show policy-firewall session
      Established Sessions = 1 Session 49F14DE0 (136.1.99.2:8)=>(136.1.27.51:0) icmp SIS_OPEN
          Created 00:00:02, Last heard 00:00:02
          ECHO request Bytes sent (initiator:responder) [360:360]
```

Verify that traffic from the **UNTRUSTED** to the **TRUSTED** zone does not work before user authentication:



Verify the ACL applied inbound at the interface enforcing auth-proxy, before authentication:

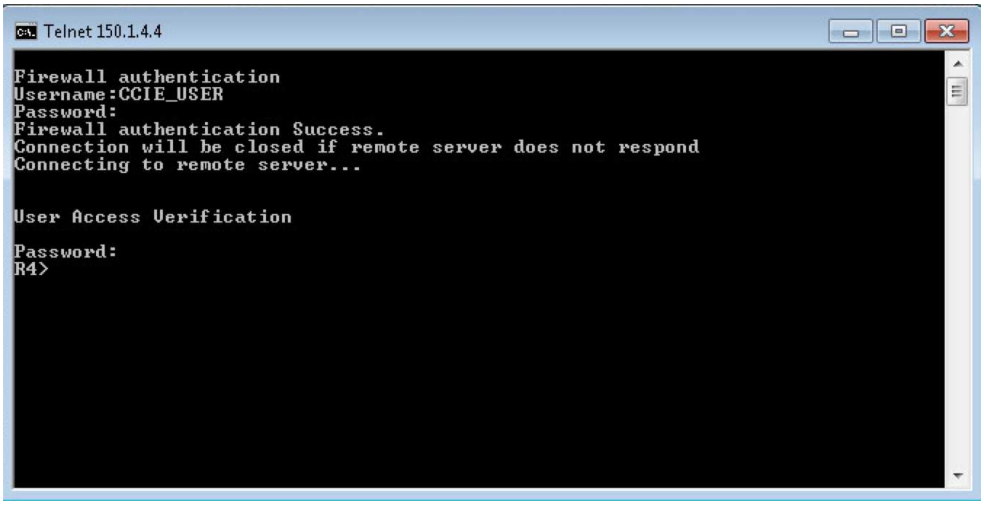
https://t.me/learningnets

```
R3#show ip access-lists DEFAULT_ACL

Extended IP access list DEFAULT_ACL
    10 permit tcp any any eq telnet
    20 permit ip any any (100 matches)
!
!R3#show ip access-lists interface fastEthernet0/0

Extended IP access list DEFAULT_ACL in
    10 permit tcp any any eq telnet
    20 permit ip any any (118 matches)
```

Telnet from **TEST-PC-A** to any destination; for example, R4, R3 will intercept the session and prompt for user authentication. If successfully authenticated, traffic from the **UNTRUSTED** zone to the **TRUSTED** zone should be allowed.




```
R3#show policy-firewall session zone-pair UNTRUSTED_TO_TRUSTED
 Zone-pair: UNTRUSTED_TO_TRUSTED
  Service-policy inspect : UNTRUSTED_TO_TRUSTED
    Class-map : PROTOCOLS_AUTHENTICATED(match-all) Established Sessions =  1
 Session 49F15860 (136.1.27.51:1272)=>(150.1.4.4:23) tcp SIS_OPEN/TCP_ESTAB


          Created 00:00:05, Last heard 00:00:03
          Bytes sent (initiator:responder) [38:65]
    Class-map : class-default(match-any)
```

Verify the ACL applied inbound at the interface enforcing auth-proxy; after authentication, note the dynamic entries identifying the ACL mapped at the identity policy level:

https://t.me/learningnets

```
R3#show ip access-lists DEFAULT_ACL

Extended IP access list DEFAULT_ACL permit ip host 136.1.27.51 any (61 matches)
    10 permit tcp any any eq telnet
    20 permit ip any any (130 matches)
!
!R3#show ip access-lists interface fastEthernet0/0
Extended IP access list DEFAULT_ACL in permit ip host 136.1.27.51 any (66 matches)

    10 permit tcp any any eq telnet
    20 permit ip any any (130 matches)
```

Verify the authentication proxy cache:

```
R3#show ip auth-proxy cache username CCIE_USER
Authentication Proxy Cache
 User Name                : CCIE_USER
Client IP                 : 136.1.27.51
Client Port               : 3842
Timeout                   : 60
Time  Remaining           : 60              Connection state          : INTERCEPT

EPM information : Authproxy
Admission feature         : Authproxy Tag Received              : CCIE_USER_TAG
Policy map used           : ADMISSION  Class map matched        : CCIE_USER_CLASS


EOU information
--------------------------------------------------------------------------
Address         Interface              AuthType   Posture-Token Age(min)
--------------------------------------------------------------------------


EPM information : EOU


Existing Firewall Sessions Information:

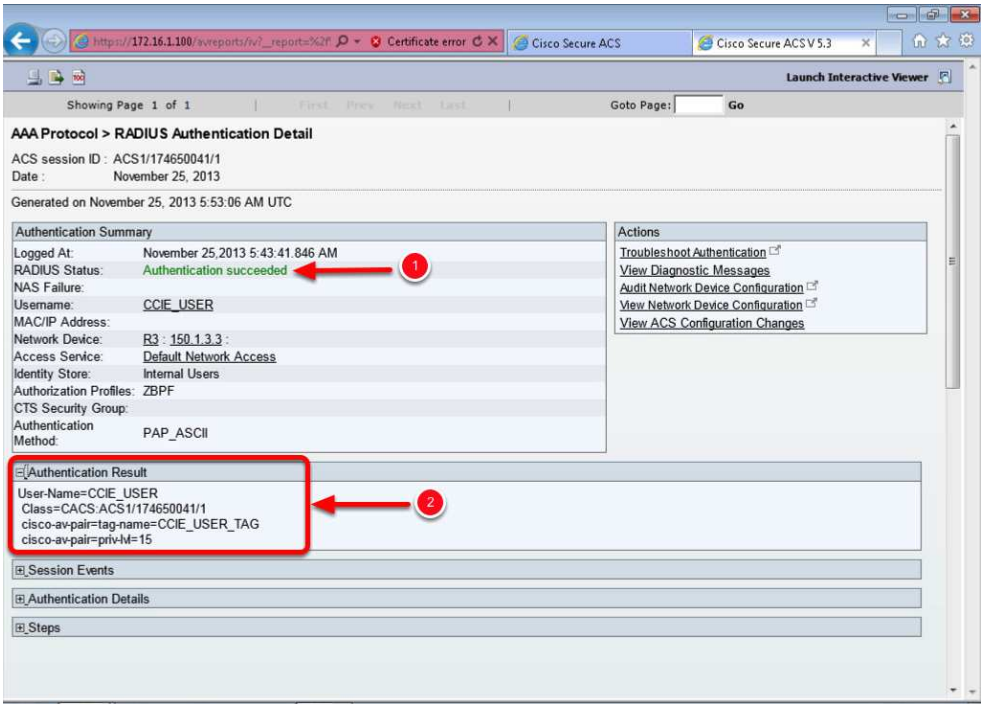
        Session 49BEF700 (136.1.27.51:3842)=>(150.1.4.4:23) tcp SIS_OPEN/TCP_ESTAB
Created 00:04:33, Last heard 00:04:30
Bytes sent (initiator:responder) [32:65]
```

https://t.me/learningnets

Verify the user group mappings for the client IP address and association between tags received from RADIUS and local class-map of type control tag:

```
R3#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.27.51         0148.f8b3.2e24.23       Nov 26 2013 12:00 AM    Automatic
!
!R3#show epm session ip 136.1.27.51
Admission feature       : Authproxy Tag Received            : CCIE_USER_TAG
Policy map used         : ADMISSION Class map matched       : CCIE_USER_CLASS
!
!R3#show user-group
Usergroup : CCIE_GROUP
--------------------------------------------------------------------------
User Name       Type     Interface              Learn         Age (min)
--------------------------------------------------------------------------
136.1.27.51     IPv4     FastEthernet0/0        Dynamic          6
```

Verify the ACS RADIUS Authentication Logs:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Static NAT

Task

- Configure static NAT on R1, considering only Gi0/1.128 as NAT inside interface:
  - SW1's VLAN 128 IP address should be visible as both 136.1.19.250 and 136.1.99.250.
  - TCP port 8080 traffic toward R1's VLAN 123 IP address should be redirected to SW1's VLAN 128 IP address on port 80.
  - SW1's Loopback0 should be statically NAT'ed based on the destination being accessed:
    - For 136.1.19.0/24, use 136.1.99.100.
    - For 136.1.0.0/16, use 136.1.99.150.
    - For 150.1.0.0/16, use 136.1.99.190.
    - Ensure that outside-to-inside traffic is subject to these rules as well.
- For each newly created translation, R1 should fire a log message.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

When configuring NAT on IOS, there are two main steps: configuring the NAT rules (static, dynamic) and designating interfaces as outside or inside. For a NAT rule to be triggered, traffic must be routed between one inside and one outside interface; so

https://t.me/learningnets

traffic forwarded between outside interfaces or between inside interfaces will never trigger any configured NAT rules. Historically, the inside interface was the one toward the private IP address space (the internal one) and the outside interface was the one toward the public IP address space (the external one), but technically you can use these two however you want. Still, depending on how you define interfaces, configuration may require additional options, and this is mainly because of order of operation:

- For inside to outside NAT, routing is performed before NAT.
- For outside to inside NAT, routing is performed after NAT.

The following terms must be understood clearly when dealing with IOS NAT:

- **Inside Local**, the IP address of a host reachable over an inside NAT interface (the real IP address)
- **Inside Global**, the IP address of **Inside Local** as visible on the outside NAT interface (the mapped IP address)
- **Outside Local**, the IP address of **Outside Global** as visible on the inside NAT interface (the mapped IP address)
- **Outside Global**, the IP address of a host reachable over an outside NAT interface (the real IP address)

Given the above statements, we can say that **Inside Local** and **Outside Local** are visible and routable on the inside NAT domain, whereas **Inside Global** and **Outside Global** are visible and routable on the outside NAT domain.

A static NAT translation creates a permanent static entry in the NAT table that is bidirectional, so traffic can be initiated both from outside to inside and from inside to outside. It is configured with the global command `ip nat inside source static [tcp|udp|esp] <inside_local> [port] [<inside_global>|interface <interface>] [port]` :

- If the TCP/UDP keyword is specified, the port option is mandatory and translation is valid only for the specific protocol/port socket.
- If the ESP keyword is specifird, the inside global IP address is not an option; only interface can be configured.
- Otherwise, there is a one-to-one IP mapping.

https://t.me/learningnets

In certain cases, you might want a single inside local address to be visible on the outside via two different inside global addresses, or a single inside global address to be visible on the inside via two different inside local addresses. Most often, this is when you have multiple inside or outside interfaces and you therefore need the same address translated differently. For IOS to allow such configuration, the `extendable` keyword must be used when defining the translation.

Static policy NAT is when you want the same inside local address to be translated on a different inside global address based on the destination being accessed. For this you must define a route-map for each specific destination (through an extended ACL, not standard) and attach it to the NAT entry. If you want the static policy NAT configuration to also be valid for when traffic is initiated from outside, you must have the `reversible` keyword attached. To better undertand this, take the following example:

```
ip nat inside source static 150.1.9.9 136.1.99.100 route-map VLAN19_SUBNET
ip nat inside source static 150.1.9.9 136.1.99.150 route-map ALL_SUBNET
```

Based on this configuration, traffic initiated from outside toward 136.1.99.100 or 136.1.99.150 will work regardless of the source, because route-map statements are not checked for outside-to-inside sessions unless the reversible keyword is configured; so for outside to inside, it basically ignores the route-map statements. However, review the following example:

```
ip nat inside source static 150.1.9.9 136.1.99.100 route-map VLAN19_SUBNET reversible
ip nat inside source static 150.1.9.9 136.1.99.150 route-map ALL_SUBNET reversible
```

Based on this configuration, traffic initiated from outside toward 136.1.99.100 will work only if the traffic source is matching the route-map. Specifically, it matches on the destination definitions of the extended ACL from the route-map.

> Route-map level **`match interface`** and **`match ip next-hop`** commands are not supported with reversible route-maps.

For inside-to-outside flow with static NAT, if by looking at the configuration an inside local address will match on several NAT statements, precedence is taken based on the order:

https://t.me/learningnets

- Static PAT
- Static NAT
- Static policy NAT

An important conclusion is that static policy NAT will basically never be matched for inside-to-outside traffic if static NAT is matched, but it will be matched for outside-to-inside traffic because here the decision is taken based on which inside global address you access. In general, it is recommended that any traffic match a unique NAT statement to avoid such unpredictible scenarios.

Configuration

R1:

```
interface GigabitEthernet0/1.128
 ip nat inside
!
interface GigabitEthernet0/1.123
 ip nat outside
!
interface GigabitEthernet0/0
 ip nat outside
!
access-list 100 permit ip any 136.1.19.0 0.0.0.255
access-list 150 deny ip any 136.1.19.0 0.0.0.255
access-list 150 permit ip any 136.1.0.0 0.0.255.255
access-list 190 permit ip any 150.1.0.0 0.0.255.255
!
route-map VLAN19_SUBNET permit 10
 match ip address 100
!
route-map ALL_SUBNET permit 10
 match ip address 150
!
route-map LOOPBACK_SUBNET permit 10
 match ip address 190
!
ip nat inside source static 150.1.9.9 136.1.99.100 route-map VLAN19_SUBNET reversible
ip nat inside source static 150.1.9.9 136.1.99.150 route-map ALL_SUBNET reversible
ip nat inside source static 150.1.9.9 136.1.99.190 route-map LOOPBACK_SUBNET reversible
!
ip nat inside source static 136.1.128.9 136.1.19.250 extendable
ip nat inside source static 136.1.128.9 136.1.99.250 extendable
ip nat inside source static tcp 136.1.128.9 80 interface GigabitEthernet0/1.123 8080
```

https://t.me/learningnets

```
!
ip nat log translations syslog
```

Verification

Verify NAT configuration, and note that policy NAT rules (configured with route-maps) are ordered in alphabetical order based on route-map names:

```
R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.1:8080    136.1.128.9:80     ---                ---
--- 136.1.19.250       136.1.128.9        ---                ---
--- 136.1.99.100       150.1.9.9          ---                ---
--- 136.1.99.150       150.1.9.9          ---                ---
--- 136.1.99.190       150.1.9.9          ---                ---
--- 136.1.99.250       136.1.128.9        ---                ---
!
!R1#show ip nat statistics
Total active translations: 7 (5 static, 2 dynamic; 2 extended)
Peak translations: 7, occurred 00:00:08 agoOutside interfaces:
 GigabitEthernet0/0, GigabitEthernet0/1.123
Inside interfaces:
 GigabitEthernet0/1.128
Hits: 5  Misses: 0
CEF Translated packets: 5, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:
-- Inside Source[Id: 0] route-map ALL_SUBNET
 [Id: 0] route-map LOOPBACK_SUBNET
 [Id: 0] route-map VLAN19_SUBNET


Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

Initiate traffic from SW1's VLAN 128 interface, and verify correct translation. You have no control over which IP address will be selected for translation with the given configuration; it will always be the lower IP address, thus 136.1.19.250:

```
SW1#ping 150.1.3.3

Type escape sequence to abort.
```

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:.!!!!

Success rate is 80 percent (4/5), round-trip min/avg/max = 1/3/9 ms
!
!R1#show ip nat translations icmp
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.19.250:2    136.1.128.9:2      150.1.3.3:2        150.1.3.3:2
```

Initiate traffic from SW1's Loopback0 toward Loopback0 subnet, and verify correct translation:

```
SW1#ping 150.1.10.10 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.10.10, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!R1#show ip nat translations icmp
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.190:3    150.1.9.9:3        150.1.10.10:3      150.1.10.10:3
```

Initiate traffic from SW1's Loopback0 towards 136.1.0.0/16 subnet, verify correct translation:

```
SW1#ping 136.1.27.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.27.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/3/9 ms
!
!R1#show ip nat translations icmp
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.150:4    150.1.9.9:4        136.1.27.3:4       136.1.27.3:4
```

Initiate traffic from SW1's Loopback0 toward 136.1.19.0/24 subnet, and verify correct translation:

```
SW1#ping 136.1.19.10 source loopback0
```

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.10, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
!
!R1#show ip nat translations icmp
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.100:8    150.1.9.9:8        136.1.19.10:8      136.1.19.10:8
```

Verify that port forwarding works, and therefore traffic on port 8080 to R1 is redirected to SW1 on port 80:

```
R3#telnet 136.1.99.1 8080
Trying 136.1.99.1, 8080 ... Open
HELO
HTTP/1.1 400 Bad Request
Date: Mon, 25 Nov 2013 23:20:53 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 136.1.99.1 closed by foreign host]
!
!R1#show ip nat translations tcp
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.1:8080    136.1.128.9:80     136.1.99.3:60423   136.1.99.3:60423

tcp 136.1.99.1:8080    136.1.128.9:80     ---                ---
```

Verify that SW1's VLAN 128 IP address is reachable over both IP addresses:

```
R3#ping 136.1.19.250

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.250, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
!R3#ping 136.1.99.250

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.99.250, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
!
```

https://t.me/learningnets

```
!R1#show ip nat translations icmp

Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.19.250:47   136.1.128.9:47     136.1.99.3:47      136.1.99.3:47
icmp 136.1.99.250:48   136.1.128.9:48     136.1.99.3:48      136.1.99.3:48
```

Check that outside-to-inside traffic is allowed based on the source IP address, because the route-map is being checked; test from 136.1.0.0/16 subnet:

```
R3#telnet 136.1.99.100
Trying 136.1.99.100 ... % Connection refused by remote host
!
!R3#telnet 136.1.99.190
Trying 136.1.99.190 ... % Connection refused by remote host
!
!R3#telnet 136.1.99.150
Trying 136.1.99.150 ... Open


User Access Verification

Password: SW1>
!
!R1#show ip nat translations tcp
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.150:23    150.1.9.9:23       136.1.99.3:41706   136.1.99.3:41706

tcp 136.1.99.1:8080    136.1.128.9:80     ---                ---
```

Check that outside-to-inside traffic is allowed based on the source IP address, because the route-map is being checked; test from 150.1.0.0/16 subnet:

```
R3#telnet 136.1.99.100 /source-interface loopback0
Trying 136.1.99.100 ... % Connection refused by remote host
!
!R3#telnet 136.1.99.150 /source-interface loopback0
Trying 136.1.99.150 ... % Connection refused by remote host
!
!R3#telnet 136.1.99.190 /source-interface loopback0
Trying 136.1.99.190 ... Open


User Access Verification
```

https://t.me/learningnets

```
Password: SW1>


!
!R1#show ip nat translations tcp
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.190:23    150.1.9.9:23       150.1.3.3:33044    150.1.3.3:33044

tcp 136.1.99.1:8080    136.1.128.9:80     ---                ---
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Dynamic NAT

Task

- Configure dynamic NAT on R1, considering only Gi0/1.128 as NAT inside interface:
  - SW1's VLAN 128 IP address should always be translated into 136.1.99.10, but allowed no more than three concurrent translations.
  - SW1's /24 Loopback0 subnet should be translated based on the destination being accessed:
    - For 136.1.0.0/16, use the range 136.1.99.100-136.1.99.105; outside-to-inside traffic should be allowed after the translation exists.
    - For 150.1.0.0/16, use the range 136.1.99.200-136.1.99.205.
- Expire the simple NAT translations after 1 hour of inactivity.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

https://t.me/learningnets

Dynamic NAT translations are created on demand, but initially the NAT table is empty. It dynamically binds one inside local address to one inside global address, which means that initially traffic must flow in the inside-to-outside direction to trigger the router to create a translation; only after the translation is created will the router allow new sessions to flow in the outside-to-inside direction (but this depends on the configuration). To configure dynamic NAT, the following steps are required:

1. Designate and configure inside/outside NAT interfaces.
2. Identify traffic subject to NAT through a standard/extended ACL, and optionally define a route-map matching the ACL.
3. Define the NAT pool, essentially the range of inside global IP addresses.
4. Configure NAT rules, binding the ACL/route-map with the NAT pool.

For regular dynamic NAT, you can use standard or extended ACL. There is no functional difference between these two options, and there is no need to use a route-map. When the router needs to create a dynamic entry in the NAT table, it will actually create two:

- One simple translation entry, identifying only the inside local and inside global addresses; this is required to basically allocate the inside global address from the pool for the period defined with the command `ip nat translation timeout <seconds>`, with a default value of 24 hours; the existence of this entry allows outside-to-inside sessions to be initiated later without any additional configurations.
- One extended translation entry for each flow matching the simple translation entry, identifying both source (inside local/global) and destination (outside local/global). Matching of the simple translation entry is solely based on the inside local address; this is why there is no difference between using a standard or extended ACL. These translation have different and shorter expiring timers, based on the protocol being used (TCP/UDP/ICMP, etc.).

For policy dynamic NAT, you must use extended ACL and also a route-map. With this configuration, you would want the same source (inside local) to be dynamically NAT'ed to different IP addresses (inside global) based on the destination being accessed. This is implemented as follows: after the route-map is used, only extended translations are created for the respective rule—no more simple translations. However, because of the missing simple translation, outside-to-inside traffic will not be allowed. To fix this, use the `reversible` keyword when configuring the NAT rule; this essentially instructs the router to again create both simple and extended translations even though a route-map is used.

The first example below is for a regular dynamic NAT configuration, resulting in both simple and extended translations. The second example is for a policy dynamic NAT

https://t.me/learningnets

configuration, resulting in only an extended translation:

```
R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.200:29   136.1.128.9:29     150.1.3.3:29       150.1.3.3:29
--- 136.1.99.200       136.1.128.9        ---                ---
!
!R1#show ip nat translations

Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.200:31   136.1.128.9:31     150.1.3.3:31       150.1.3.3:31
```

As a side observation, the `reversible` keyword has no use without a route-map, because an ACL-based dynamic NAT configuration implicitly allows outside-to-inside traffic based on the simple translation entry; even if configured, it will not appear in the configuration. With dynamic NAT, by default IOS will randomly select an IP address from the pool as the inside global and associate it with one inside local address. This behavior can be changed so that the host portion from the inside local address is preserved and matches the host portion from the inside global address by using the `type match-host` option when defining the NAT pool.

For inside-to-outside flow with dynamic NAT, if by looking at the configuration an inside local address will match on several NAT statements, precedence is taken based on the order:

- Regular dynamic NAT
- Policy dynamic NAT

This is purely based on how NAT statements are alphabetically ordered in the configuration (not the order of configuration for NAT statements). In general, it is recommended that any traffic match a unique NAT statement to avoid such unpredictable scenarios.

Configuration

```
R1:

interface GigabitEthernet0/1.128
 ip nat inside
!
interface GigabitEthernet0/1.123
 ip nat outside
!
```

https://t.me/learningnets

```
interface GigabitEthernet0/0
 ip nat outside
!
access-list 10 permit host 136.1.128.9
access-list 100 permit ip 150.1.9.0 0.0.0.255 136.1.0.0 0.0.255.255
access-list 190 permit ip 150.1.9.0 0.0.0.255 150.1.0.0 0.0.255.255
!
ip nat pool ACL10_POOL 136.1.99.10 136.1.99.10 prefix-length 24
ip nat pool ACL100_POOL 136.1.99.100 136.1.99.105 prefix-length 24
ip nat pool ACL190_POOL 136.1.99.200 136.1.99.205 netmask 255.255.255.0
!
route-map ACL100_RMAP permit 10
 match ip address 100
!
route-map ACL190_RMAP permit 10
 match ip address 190
!
ip nat inside source list 10 pool ACL10_POOL
ip nat inside source route-map ACL100_RMAP pool ACL100_POOL reversible
ip nat inside source route-map ACL190_RMAP pool ACL190_POOL
!
ip nat translation timeout 3600
ip nat translation max-entries host 136.1.128.9 3
```

Verification

Verify the NAT configuration; because of the dynamic NAT nature, there are no permanent translations so there is no need to check the NAT table:

```
R1#show ip nat statistics
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Peak translations: 9, occurred 03:27:17 agoOutside interfaces:
 GigabitEthernet0/0, GigabitEthernet0/1.123
Inside interfaces:
 GigabitEthernet0/1.128
Hits: 0  Misses: 0
CEF Translated packets: 0, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:
-- Inside Source
[Id: 16] access-list 10 pool ACL10_POOL refcount 0
 pool ACL10_POOL: netmask 255.255.255.0
        start 136.1.99.10 end 136.1.99.10
        type generic, total addresses 1, allocated 0 (0%), misses 0
```

https://t.me/learningnets

```
[Id: 17] route-map ACL100_RMAP pool ACL100_POOL refcount 0
 pool ACL100_POOL: netmask 255.255.255.0
        start 136.1.99.100 end 136.1.99.105
        type generic, total addresses 6, allocated 0 (0%), misses 0
[Id: 18] route-map ACL190_RMAP pool ACL190_POOL refcount 0
 pool ACL190_POOL: netmask 255.255.255.0
        start 136.1.99.200 end 136.1.99.205
        type generic, total addresses 6, allocated 0 (0%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0nat-limit statistics:
host 136.1.128.9: max allowed 2, used 0, missed 0

Queued Packets: 0
```

Initiate three ICMP sessions from SW1's VLAN 128 IP address within 60 seconds (the default timeout for ICMP translations). Note that the third session is not allowed because of maximum number of translations are reached; this is because IOS first builds a simple translation to reserve the inside global address for the inside local address, and then it builds extended translations for each source/destination pair.

```
SW1#ping 150.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
!
!SW1#ping 136.1.19.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.10, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
!
!SW1#ping 136.1.99.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.99.3, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.10:41    136.1.128.9:41     150.1.3.3:41       150.1.3.3:41
icmp 136.1.99.10:42    136.1.128.9:42     136.1.19.10:42     136.1.19.10:42
--- 136.1.99.10        136.1.128.9        ---                ---
```

https://t.me/learningnets

```
!
!R1#show ip nat statistics | section nat-limit
nat-limit statistics: host 136.1.128.9: max allowed 3, used 3, missed 10
```

The extended translations will expire automatically after 60 seconds and be removed from the NAT table; however the simple translation will remain in the NAT table for the configured timeout of 3600 seconds (in case of inactivity), with a default of 24 hours:

```
R1#show ip nat translations verbose
Pro Inside global      Inside local       Outside local      Outside global
--- 136.1.99.10        136.1.128.9        ---                ---    create 00:07:47, use 00:01:04
timeout:3600000, left 00:58:55
, Map-Id(In): 16,
    flags:
limited, use_count: 0, entry-id: 148, lc_entries: 0
```

Initiate traffic from SW1's Loopback0 subnet toward 136.1.0.0/16, and verify correct NAT translation:

```
SW1#ping 136.1.19.10 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.10, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!SW1#ping 136.1.99.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.99.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 .!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
!
!R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.100:49   150.1.9.9:49       136.1.19.10:49     136.1.19.10:49
icmp 136.1.99.100:50   150.1.9.9:50       136.1.99.3:50      136.1.99.3:50
--- 136.1.99.100       150.1.9.9          ---                ---
```

When the dynamic simple translation has been created, initiate traffic from outside

https://t.me/learningnets

to inside toward the inside global IP address:

```
R3#telnet 136.1.99.100
Trying 136.1.99.100 ... Open


User Access Verification

Password: SW1>show tcp brief
TCB       Local Address           Foreign Address        (state)
05F79E84  150.1.9.9.23            136.1.99.3.29369       FINWAIT2
05F7BD68  150.1.9.9.23            136.1.99.3.36625       ESTAB
05F7B8AC  136.1.128.9.23          136.1.99.3.59401       FINWAIT2
05C3D2AC  136.1.128.9.23          136.1.99.3.56711       FINWAIT2
05F7A340  136.1.128.9.23          136.1.99.3.45257       FINWAIT2
!
!R1#show ip nat translations tcp
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.100:23    150.1.9.9:23       136.1.99.3:36625   136.1.99.3:36625
```

Initiate traffic from SW1's Loopback0 subnet toward 150.1.0.0/16, verify correct NAT translation, and note that there are no simple translations created, just extended; because of this, and because the keyword reversible was not configured, traffic from outside to inside toward the inside global IP address is not allowed:

```
SW1#ping 150.1.3.3 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.3.3, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
!
!SW1#ping 150.1.10.10 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.10.10, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 136.1.99.200:52   150.1.9.9:52       150.1.3.3:52       150.1.3.3:52
icmp 136.1.99.200:53   150.1.9.9:53       150.1.10.10:53     150.1.10.10:53
```

https://t.me/learningnets

```
!
!R3#telnet 136.1.99.200
Trying 136.1.99.200 ... % Connection refused by remote host
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Overload NAT

Task

- Configure NAT overload on R1, considering only Gi0/1.128 as NAT inside interface:
  - VLAN 128 subnet should use a pool with the range 136.1.99.10-136.1.99.15; outside-to-inside traffic should be allowed when translation exists.
  - SW1's /24 subnet of Loopback0 should use R1's Loopback0 IP address.
- NAT entries for telnet sessions must be removed after 2 minutes of inactivity.
- NAT entries for gracefully closed TCP sessions must be kept no longer than 30 seconds.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

From a configuration perspective, NAT overload (PAT) is very similar to dynamic NAT. The only difference is the addition of the `overload` keyword when defining the NAT rule. NAT overload does not create a one-to-one mapping between an inside local and an inside global IP address; it always creates a many-to-one mapping between multiple inside local addresses and one inside global address. This is achieved through port multiplexing, hence the name *overload*. The implicit outcome is that, by default, NAT overload only creates extended translations and no simple

https://t.me/learningnets

translations. Because of this, policy PAT no longer requires a route-map; using an extended ACL will work as well.

The route-map is a requirement only if the `reversible` keyword is needed, so outside-to-inside traffic is allowed following the creation of the translation. Just as with dynamic NAT, the `reversible` keyword will result in the creation of both simple and extended translations. The outside-to-inside traffic has limited applicability because the simple translation will exist only for the inside host that triggered the translation (and only for the protocol used for that session, like TCP/UDP/ICMP). For this reason, if true support for outside-to-inside traffic is a requirement, use static NAT or dynamic policy NAT.

PAT allows for the use of the `interface` keyword to define the inside global address; the specified interface need not be NAT enabled, nor must it be the exit interface for the traffic. If a PAT pool is specified instead, IOS will always use a single IP address from the pool until all ports are exhausted, and then it will select the next available IP from the pool.

Configuration

https://t.me/learningnets

```
R1:

interface GigabitEthernet0/1.128
 ip nat inside
!
interface GigabitEthernet0/1.123
 ip nat outside
!
interface GigabitEthernet0/0
 ip nat outside
!
access-list 10 permit 136.1.128.0 0.0.0.255
access-list 20 permit 150.1.9.0 0.0.0.255
!
ip nat pool ACL10_POOL 136.1.99.10 136.1.99.15 prefix-length 24
!
route-map ACL10_RMAP permit 10
 match ip address 10
!
ip nat inside source route-map ACL10_RMAP pool ACL10_POOL overload reversible
ip nat inside source list 20 interface Loopback0 overload
!
ip nat translation port-timeout tcp 23 120
ip nat translation finrst-timeout 30
```

Verification

Verify the NAT configuration; because of PAT functionality, there are no permanent translations so there is no need to check the NAT table:

```
R1#show ip nat statistics
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Peak translations: 0Outside interfaces:
 GigabitEthernet0/0, GigabitEthernet0/1.123
Inside interfaces:
 GigabitEthernet0/1.128

Hits: 0  Misses: 0
CEF Translated packets: 0, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:
-- Inside Source
[Id: 34] access-list 20 interface Loopback0 refcount 0
```

https://t.me/learningnets

```
[Id: 33] route-map ACL10_RMAP pool ACL10_POOL refcount 0
 pool ACL10_POOL: netmask 255.255.255.0
        start 136.1.99.10 end 136.1.99.15
        type generic, total addresses 6, allocated 0 (0%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0
nat-limit statistics:
Queued Packets: 0
```

Initiate traffic from the VLAN 128 subnet, verify correct NAT translation, and check that outside-to-inside traffic is allowed:

```
SW1#telnet 150.1.3.3
Trying 150.1.3.3 ... Open


User Access Verification

Password: R3>
!
!R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.10:50708  136.1.128.9:50708  150.1.3.3:23       150.1.3.3:23
tcp 136.1.99.10:0      136.1.128.9:0      ---                ---
!
!R3#telnet 136.1.99.10
Trying 136.1.99.10 ... Open


User Access Verification

Password: SW1>
!
!R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.10:23     136.1.128.9:23     136.1.99.3:47864   136.1.99.3:47864

tcp 136.1.99.10:50708  136.1.128.9:50708  150.1.3.3:23       150.1.3.3:23
tcp 136.1.99.10:0      136.1.128.9:0      ---                ---
```

Verify the translation timeout of 120 seconds for the telnet session:

https://t.me/learningnets

```
R1#show ip nat translations verbose

Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.10:50708  136.1.128.9:50708  150.1.3.3:23       150.1.3.3:23
    create 00:00:20, use 00:00:18 timeout:120000, left 00:01:41
, Map-Id(In): 33,
    flags:
extended, use_count: 0, entry-id: 254, lc_entries: 0
tcp 136.1.99.10:0      136.1.128.9:0      ---                ---
    create 00:00:20, use 00:00:20 timeout:3600000, left 00:59:39, Map-Id(In): 33,
    flags:
extended, extendable, global-only-entry, routemap-out2in, use_count: 1, entry-id: 253, lc_entries: 0
```

Exit the telnet session (so that it is gracefully closed), and verify the timeout of 30 seconds for the session:

```
R1#show ip nat translations verbose
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.10:50708  136.1.128.9:50708  150.1.3.3:23       150.1.3.3:23
    create 00:00:06, use 00:00:04 timeout:120000, left 00:00:25
, Map-Id(In): 33,
    flags:
extended, timing-out, use_count: 0, entry-id: 254, lc_entries: 0
tcp 136.1.99.10:0      136.1.128.9:0      ---                ---
    create 00:00:06, use 00:00:06 timeout:3600000, left 00:59:53, Map-Id(In): 33,
    flags:
extended, extendable, global-only-entry, routemap-out2in, use_count: 1, entry-id: 258, lc_entries: 0
```

Initiate traffic from SW1's Loopback0 subnet, verify correct NAT translation, and check that outside-to-inside traffic is not allowed, because there is no simple translation created:

https://t.me/learningnets

```
SW1#telnet 150.1.3.3 /source-interface loopback0

Trying 150.1.3.3 ... Open


User Access Verification

Password: R3>

!
!R1#show ip nat translations

Pro Inside global      Inside local       Outside local      Outside global
tcp 150.1.1.1:54949    150.1.9.9:54949    150.1.3.3:23       150.1.3.3:23
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Outside NAT

Task

- Configure outside NAT on R1, considering Gi0/1.128 as NAT inside and Gi0/1.123 as NAT outside:
  - SW1 should reach R3's Loopback0 address at 136.1.128.3.
  - VLAN 27 traffic should be dynamically translated using the range 136.1.128.100-136.1.128.102.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Note the following important points about outside NAT:

- There is no route-map option for static NAT, so there is no policy NAT.
- NAT overload is not supported, so the `interface` keyword is not available.
- The `reversible` keyword is not supported.

When traffic flows from inside to outside, first IOS performs routing, and then NAT. At the same time that a session is translated, it must be passed between one inside and one outside interface. In our example, when SW1 initiates a session toward the NAT'ed IP address of R3's Loopback0 (136.1.128.3):

https://t.me/learningnets

- R1 will receive the packet (because of the proxy ARP for the NAT'ed IP address).
- 136.1.128.3 is locally connected on its inside NAT interface, so it will not route the packet on its outside interface.
- Because packet is not routed, it will not be translated.

To fix this issue, use the option `add-route` when defining the outside static NAT rule. At that moment, the router will install a /32 static route for the outside local address (which is 136.1.128.3) through the outside global address (which is 150.1.3.3); this will allow the router to first route the packet on its outside interface and then perform NAT. The same problem occurs when configuring dynamic outside NAT, but in this case the static route is added on demand, when the dynamic NAT entry is created; because the outside local address is unknown initially, it is randomly selected from the NAT pool upon NAT entry creation. With dynamic outside NAT, the `add-route` option is attached to the NAT entry rule, not to the NAT pool.

Configuration

```
R1:
interface GigabitEthernet0/1.128
 ip nat inside
!
interface GigabitEthernet0/1.123
 ip nat outside
!
access-list 10 permit 136.1.27.0 0.0.0.255
!
ip nat pool ACL10_POOL 136.1.128.100 136.1.128.102 prefix-length 30
ip nat outside source list 10 pool ACL10_POOL add-route
ip nat outside source static 150.1.3.3 136.1.128.3 add-route
```

Verification

Verify the existence of static NAT in the NAT table and NAT configuration:

```
R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
--- ---                ---                136.1.128.3        150.1.3.3
!
!R1#show ip nat statistics
Total active translations: 1 (1 static, 0 dynamic; 0 extended)
Peak translations: 0 Outside interfaces:
```

https://t.me/learningnets

```
 GigabitEthernet0/1.123
 Inside interfaces:
  GigabitEthernet0/1.128

Hits: 0  Misses: 0
CEF Translated packets: 0, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:
-- Outside Source
[Id: 5] access-list 10 pool ACL10_POOL refcount 0
 pool ACL10_POOL: netmask 255.255.255.0
        start 136.1.128.100 end 136.1.128.102
        type generic, total addresses 3, allocated 0 (0%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

Verify the existence of the static route for the static NAT; for the dynamic NAT it will be added on demand:

```
R1#show ip route static | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 10 subnets, 2 masks
S        136.1.128.3/32 [1/0] via 150.1.3.3
```

Initiate traffic from SW1 toward the NAT'ed IP address of R3's Loopback0, and verify the extended NAT entry:

https://t.me/learningnets

```
SW1#telnet 136.1.128.3

Trying 136.1.128.3 ... Open


User Access Verification

Password: R3>

!
!R1#show ip nat translations

Pro Inside global      Inside local       Outside local      Outside global
--- ---                ---                136.1.128.3        150.1.3.3
tcp 136.1.128.9:26078  136.1.128.9:26078  136.1.128.3:23     150.1.3.3:23
```

Initiate traffic from R3's VLAN 27 subnet toward SW1, and verify the dynamic entries in the NAT table (simple and extended) and the created static route:

```
R3#telnet 150.1.9.9 /source-interface fastEthernet0/0
Trying 150.1.9.9 ... Open


User Access Verification

Password: SW1>
!
!R1#show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
--- ---                ---                136.1.128.101      136.1.27.3
tcp 150.1.9.9:23       150.1.9.9:23       136.1.128.101:41824 136.1.27.3:41824
!
!R1#show ip route static | b Gateway
Gateway of last resort is not set

      136.1.0.0/16 is variably subnetted, 9 subnets, 2 masks
S        136.1.128.3/32 [1/0] via 150.1.3.3S        136.1.128.101/32 [1/0] via 136.1.27.3
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS NVI NAT

Task

- Configure NAT on R1 without defining any inside/outside domains:
  - R3's Loopback0 should be statically mapped to 136.1.128.3.
  - VLAN 99 traffic should be overloaded using R1's Loopback0.
  - VLAN 27 traffic should be dynamically NAT'ed using the pool range 136.1.128.150-136.1.128.155; no simple translations are allowed.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

NAT NVI moves from the old paradigm of having inside and outside NAT interfaces, with different order of processing for inside-to-outside and outside-to-inside traffic. With NVI NAT, interfaces passing traffic will be just NAT enabled using the command `ip nat enable`, and routing will always be performed twice: first time to route packets to the NVI interface, after which NAT happens, and second to route packets towards the final destination. For NAT NVI to work, traffic must pass through NAT-enabled interfaces. Another limitation of regular NAT was that with interfaces configured as inside/outside, no NAT rules can apply for inside-to-inside or outside-to-outside traffic. With NVI NAT, this limitation does not exist, because there are no longer two different domains (inside and outside); there is a single NAT

https://t.me/learningnets

domain identified by all NAT-enabled interfaces.

Static, dynamic, and overload NAT is supported with NVI NAT (the same behavior as with regular NAT), but the following should be considered:

- There is no route-map option for static NAT, so there is no policy NAT.
- The `reversible` keyword is not supported.

With NAT overload, if the `interface` keyword is used, it can be different from the interface traffic is exiting on, but it must be NAT enabled. Cisco recommends that NVI NAT be used for inter-VRF routing, when the session source is in one VRF and the session destination is in another VRF; for all other case (intra-VRF, or from one VRF to global), use regular NAT. Although you might find many bugs if used in unsupported scenarios, it might still work depending on the IOS version.

Configuration

```
R1:
interface GigabitEthernet0/1.128
 ip nat enable
!
interface GigabitEthernet0/1.123
 ip nat enable
!
interface Loopback0
 ip nat enable
!
access-list 100 permit ip 136.1.99.0 0.0.0.255 any
access-list 150 permit ip 136.1.27.0 0.0.0.255 any
!
route-map ACL150_RMAP permit 10
 match ip address 150
!
ip nat source static 150.1.3.3 136.1.128.3
ip nat source list 100 interface Loopback0 overload
!
ip nat pool ACL150_POOL 136.1.128.150 136.1.128.155 prefix-length 24
ip nat source route-map ACL150_RMAP pool ACL150_POOL overload
```

Verification

https://t.me/learningnets

Verify the static NAT translation and NAT configuration:

```
R1#show ip nat nvi translations
Pro Source global      Source local       Destin  local      Destin  global
--- 136.1.128.3        150.1.3.3          ---                ---
!
!R1#show ip nat nvi statistics
Total active translations: 1 (1 static, 0 dynamic; 0 extended)NAT Enabled interfaces:
 GigabitEthernet0/1.123, GigabitEthernet0/1.128, Loopback0

Hits: 0  Misses: 0
CEF Translated packets: 0, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:
-- Source [Id: 6] access-list 100 interface Loopback0 refcount 0[Id: 3] route-map ACL150_RMAP pool ACL150_POOL refco
 pool ACL150_POOL: netmask 255.255.255.0
        start 136.1.128.150 end 136.1.128.155
        type generic, total addresses 6, allocated 0 (0%), misses 0
```

Verify the static NAT translation; the session should be allowed bidirectionally:

https://t.me/learningnets

```
R3#telnet 150.1.9.9 /source-interface loopback0

Trying 150.1.9.9 ... Open


User Access Verification

Password: SW1>

!
!SW1#telnet 136.1.128.3 /source-interface loopback0

Trying 136.1.128.3 ... Open


User Access Verification

Password: R3>

!
!R1#show ip nat nvi translations

Pro Source global      Source local       Destin  local      Destin  global
tcp 136.1.128.3:50851  150.1.3.3:50851    150.1.9.9:23       150.1.9.9:23

--- 136.1.128.3        150.1.3.3          ---                ---
tcp 150.1.9.9:60103    150.1.9.9:60103    136.1.128.3:23     150.1.3.3:23
```

Verify NAT overload; only extended translations should exist:

https://t.me/learningnets

```
R3#telnet 150.1.9.9

Trying 150.1.9.9 ... Open


User Access Verification

Password: SW1>

!
!R1#show ip nat nvi translations

Pro Source global      Source local       Destin  local      Destin  global
tcp 150.1.1.1:27128    136.1.99.3:27128   150.1.9.9:23       150.1.9.9:23

--- 136.1.128.3        150.1.3.3          ---                ---
```

Verify dynamic NAT. Only extended translations should exist because of the use of route-map; reverse-traffic is not allowed.

```
R3#telnet 150.1.9.9 /source-interface fastEthernet0/0
Trying 150.1.9.9 ... Open


User Access Verification

Password: SW1>
!
!R1#show ip nat nvi translations
Pro Source global      Source local       Destin  local      Destin  global
tcp 136.1.128.150:48870 136.1.27.3:48870  150.1.9.9:23       150.1.9.9:23
--- 136.1.128.3        150.1.3.3          ---                ---
!
!SW1#telnet 136.1.128.150
Trying 136.1.128.150 ... % Connection refused by remote host
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Stateful NAT

Task

- Configure R1 and R2 to share their NAT table entries, with R1 being the primary router:
  - Consider GI0/1.128 as NAT inside and Gi0/1.123 as NAT outside.
  - SW1's Loopback0 should be statically mapped to 136.1.99.9.
  - VLAN 128 traffic should be overloaded using the range 136.1.99.100 136.1.99.105.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Stateful NAT allows two NAT devices to function as a translation group in ACTIVE-STANDBY mode. The ACTIVE device in the translation group handles traffic requiring translation, and it informs the STANDBY translator of active flows as they occur. The STANDBY translator can use information from the active translator to prepare duplicate translation table entries, so if the ACTIVE translator fails, the traffic can rapidly be switched to the STANDBY. The traffic flow continues because the same network address translations are used, and the state of those translations has been synchronized.

The ACTIVE and STANBY device roles can be automatically detected based on the

https://t.me/learningnets

HSRP roles (so this option requires using HSRP), or they can be manually specified. This task maskes use of HSRP, but configuring devices manually would be similar to the following example:

```
R1:
ip nat stateful id 11
 primary 136.1.128.1
 peer 136.1.128.2
 mapping-id 10

R2:

ip nat stateful id 22
 backup 136.1.128.2
 peer 136.1.128.1
 mapping-id 10
```

Configuration for stateful NAT is very similar to regular NAT, with some simple additional commands:

- Stateful NAT is globally activated with the command `ip nat stateful id <nr>` ; the ID must be unique between routers.
- Primary/backup roles must be defined manually as shown above or identified from HSRP using the command `redundancy <hsrp_group_name> .`
- An identical mapping ID must be configured on both routers using the command `mapping-id <nr>` ; this ID will be attached to the regular NAT commands to identify which NAT-able entries will be synchronized from ACTIVE to STANDBY.
- Static NAT entries requiring stateful NAT will have the `redundancy <hsrp_group_name> mapping-id <nr>` option attached; because the simple translation is always active in the NAT table, the `redundancy` keyword is required so that only the ACTIVE router will have control over it and respond to ARP requests.
- Dynamic or overload NAT entries requiring stateful NAT will have the `mapping-id <nr>` option attached.
- All other NAT-related commands must be identical between the two routers to ensure the same service level, regardless of which router owns the ACTIVE role.

The following are two important restrictions of stateful NAT:

https://t.me/learningnets

- It does not work with NVI NAT, only with regular NAT (inside/outside domains).
- NAT overload with the `interface` keyword is not supported, simply because the IP address of the interface between the two routers is different.

Configuration

R1:

```
interface GigabitEthernet0/1.128
 ip nat inside
 standby 10 ip 136.1.128.12
 standby 10 priority 150
 standby 10 preempt
 standby 10 name HSRPNAT
!
interface GigabitEthernet0/1.123
 ip nat outside
!
ip nat stateful id 11
 redundancy HSRPNAT
 mapping-id 30
!
access-list 100 permit ip 136.1.128.0 0.0.0.255 any
!
ip nat pool ACL100_POOL 136.1.99.100 136.1.99.105 prefix-length 24
ip nat inside source static 150.1.9.9 136.1.99.9 redundancy HSRPNAT mapping-id 30
ip nat inside source list 100 pool ACL100_POOL mapping-id 30 overload
```

R2:

```
interface GigabitEthernet0/1.128
 ip nat inside
 standby 10 ip 136.1.128.12
 standby 10 preempt
 standby 10 name HSRPNAT
!
interface GigabitEthernet0/1.123
 ip nat outside
!
ip nat stateful id 22
 redundancy HSRPNAT
 mapping-id 30
!
```

https://t.me/learningnets

```
access-list 100 permit ip 136.1.128.0 0.0.0.255 any
!
ip nat pool ACL100_POOL 136.1.99.100 136.1.99.105 prefix-length 24
ip nat inside source static 150.1.9.9 136.1.99.9 redundancy HSRPNAT mapping-id 30
ip nat inside source list 100 pool ACL100_POOL mapping-id 30 overload
```

Verification

Verify the stateful failover NAT:

```
R1#show ip snat distributed

Stateful NAT Connected Peers
SNAT: Mode IP-REDUNDANCY :: ACTIVE
: State READY
    : Local Address 136.1.128.1 : Local NAT id 11
    : Peer Address 136.1.128.2 : Peer NAT id 22
: Mapping List 30

!
!R2#show ip snat distributed

Stateful NAT Connected Peers
SNAT: Mode IP-REDUNDANCY :: STANDBY
    : State READY
    : Local Address 136.1.128.2 : Local NAT id 22
    : Peer Address 136.1.128.1 : Peer NAT id 11
: Mapping List 30
```

Initiate traffic for static NAT, and verify that translation is replicated from ACTIVE to STANDBY; on STANDBY you should see the `created-by-remote` keyword:

```
R4#telnet 136.1.99.9
Trying 136.1.99.9 ... Open


User Access Verification

Password: SW1>
!
!R1#show ip nat translations tcp verbose
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.9:23      150.1.9.9:23       136.1.99.4:31153   136.1.99.4:31153
```

https://t.me/learningnets

```
    create 00:00:13, use 00:00:11 timeout:86400000, left 23:59:48,
    flags: extended, use_count: 0 nat_id: 11 nat_entry_num: 28 nat_mapping_id[in]: 30
 nat_mapping_id[out]: 0, entry-id: 37, lc_entries: 0
!
!R2#show ip nat translations tcp verbose
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.9:23      150.1.9.9:23       136.1.99.4:31153   136.1.99.4:31153
    create 00:00:44, use 00:00:44 timeout:86400000, timing-out,
    flags: extended, created-by-remote, use_count: 0 nat_id: 11 nat_entry_num: 28 nat_mapping_id[in]: 30
 nat_mapping_id[out]: 0, entry-id: 13, lc_entries: 0
```

Initiate traffic for NAT overload, and verify that translation is replicated from ACTIVE to STANDBY; on STANDBY you should see the **`created-by-remote`** keyword. Because of overload NAT, only extended translation will be created:

```
SW1#telnet 150.1.3.3
Trying 150.1.3.3 ... Open


User Access Verification

Password: R3>
!
!R1#show ip nat translations tcp verbose
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.100:62149 136.1.128.9:62149  150.1.3.3:23       150.1.3.3:23
    create 00:00:55, use 00:00:51 timeout:86400000, left 23:59:08, Map-Id(In): 9,
    flags: extended, use_count: 0 nat_id: 11 nat_entry_num: 29 nat_mapping_id[in]: 30
 nat_mapping_id[out]: 0, entry-id: 38, lc_entries: 0
!
!R2#show ip nat translations tcp verbose
Pro Inside global      Inside local       Outside local      Outside global
tcp 136.1.99.100:62149 136.1.128.9:62149  150.1.3.3:23       150.1.3.3:23
    create 00:01:30, use 00:01:30 timeout:86400000, timing-out, Map-Id(In): 3,
    flags: extended, created-by-remote, use_count: 0 nat_id: 11 nat_entry_num: 29 nat_mapping_id[in]: 30
 nat_mapping_id[out]: 0, entry-id: 14, lc_entries: 0
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Policy Routing

Task

- Configure policy routing on R1 for traffic received on its GigabitEthernet0/0 interface:
  - If destination is 136.1.0.0/16, forward it to R2's VLAN 128 interface.
  - If destination is 150.1.0.0/16, forward it to R2's VLAN 123 interface.
  - For other destinations, forward it to R3's VLAN 123 interface only if there is no route available.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

When a router receives an IPv4/IPv6 packet, it will perform a lookup for the destination in the routing table and forward it accordingly. With policy-based routing (PBR), we can instruct the router to completely ignore the routing table, effectively bypassing it, and forward the packet to a particular next-hop based on:

- Source IP address
- Destination IP address
- Source and destination IP address

From a configuration perspective, the steps are:

https://t.me/learningnets

1. Identify the traffic with a standard or extended ACL.
2. Configure a route-map matching the previously configured ACL as a condition.
3. Within the route-map, set the action by configuring the next-hop; use the command `set ip next-hop <IP_address> .`
4. Apply the route-map on the interface on which traffic is received inbound, using the command `ip policy route-map <rmap_name> .`

Optionally, we can configure the router to first perform a regular route lookup for the destination, and only if there is no match in the routing table to use the PBR configuration; this requires the route-map level command `set ip default next-hop <IP_address>` and can be seen as an alternative to having a default route present in the routing table.

Configuration

R1:

```
access-list 100 permit ip any 136.1.0.0 0.0.255.255
access-list 150 permit ip any 150.1.0.0 0.0.255.255
!
route-map POLICY_ROUTING permit 10
 match ip address 100
 set ip next-hop 136.1.128.2
!
route-map POLICY_ROUTING permit 20
 match ip address 150
 set ip next-hop 136.1.99.2
!
route-map POLICY_ROUTING permit 30
 set ip default next-hop 136.1.99.3
!
interface GigabitEthernet0/0
 ip policy route-map POLICY_ROUTING
```

Verification

Verify that policy routing is configured and check matched before issuing any traffic:

```
R1#show ip policy
Interface      Route map Gi0/0           POLICY_ROUTING
!
!R1#show route-map
```

https://t.me/learningnets

```
route-map POLICY_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.128.2 Policy routing matches: 0 packets, 0 bytes
route-map POLICY_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.99.2 Policy routing matches: 0 packets, 0 bytes
route-map POLICY_ROUTING, permit, sequence 30
  Match clauses:
  Set clauses:
    ip default next-hop 136.1.99.3 Policy routing matches: 0 packets, 0 bytes
```

Initiate traffic toward 136.1.0.0/16 subnet, note the path it takes, and verify that packets are matching the route-map entry:

https://t.me/learningnets

```
SW2#traceroute 136.1.27.3


Type escape sequence to abort.
Tracing the route to 136.1.27.3

  1 136.1.19.1 8 msec 0 msec 0 msec 2 136.1.128.2 0 msec 0 msec 0 msec

  3 136.1.99.3 8 msec *  0 msec
!
!R1#show route-map

route-map POLICY_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.128.2 Policy routing matches: 9 packets, 540 bytes

route-map POLICY_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.99.2
  Policy routing matches: 0 packets, 0 bytes
route-map POLICY_ROUTING, permit, sequence 30
  Match clauses:
  Set clauses:
    ip default next-hop 136.1.99.3
  Policy routing matches: 0 packets, 0 bytes
```

Initiate traffic toward 150.1.0.0/16 subnet, note the path it takes, and verify that packets are matching the route-map entry:

```
SW2#traceroute 150.1.3.3

Type escape sequence to abort.
Tracing the route to 150.1.3.3

  1 136.1.19.1 0 msec 0 msec 0 msec 2 136.1.99.2 0 msec 0 msec 0 msec
  3 136.1.99.3 9 msec *  0 msec
!
!R1#show route-map
route-map POLICY_ROUTING, permit, sequence 10
  Match clauses:
```

https://t.me/learningnets

```
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.128.2
  Policy routing matches: 9 packets, 540 bytes
route-map POLICY_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.99.2 Policy routing matches: 6 packets, 360 bytes

route-map POLICY_ROUTING, permit, sequence 30
  Match clauses:
  Set clauses:
    ip default next-hop 136.1.99.3
  Policy routing matches: 0 packets, 0 bytes
```

Configure a new Loopback on R3 outside the two subnet ranges, and configure on SW2 a static route for it so that traffic reaches R1:

```
R3:
interface Loopback1
 ip address 160.1.33.33 255.255.255.255


SW2:

ip route 160.1.33.33 255.255.255.255 136.1.19.1
```

Initiate traffic toward the newly created Loopback, note the path it takes, and verify that packets are matching the route-map entry:

```
R1#show ip route 160.1.33.33 255.255.255.255
% Subnet not in table
!
!SW2#traceroute 160.1.33.33

Type escape sequence to abort.
Tracing the route to 160.1.33.33

  1 136.1.19.1 0 msec 8 msec 0 msec 2 136.1.99.3 0 msec *  0 msec
!
!R1#show route-map
route-map POLICY_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
```

https://t.me/learningnets

```
  Set clauses:
    ip next-hop 136.1.128.2
  Policy routing matches: 9 packets, 540 bytes
route-map POLICY_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.99.2
  Policy routing matches: 6 packets, 360 bytes
route-map POLICY_ROUTING, permit, sequence 30
  Match clauses:
  Set clauses:
    ip default next-hop 136.1.99.3 Policy routing matches: 6 packets, 360 bytes
```

Configure a static route on R1 for the newly created Loopback. Initiate traffic again, and note the path it takes; this time policy routing should not be triggered. Even though route-map counters increase, this is because of the order of operation; if you issue a **`debug ip policy`** on R1, you'll see from the messages that packets are not policy routed. Logs will be similar to **`IP: s=136.1.19.10 (GigabitEthernet0/0), d=160.1.33.33, len 28, policy rejected -- normal forwarding .`**

```
R1:
ip route 160.1.33.33 255.255.255.255 136.1.99.2
!
!SW2#traceroute 160.1.33.33

Type escape sequence to abort.
Tracing the route to 160.1.33.33

  1 136.1.19.1 0 msec 0 msec 8 msec 2 136.1.99.2 0 msec 0 msec 0 msec
3 136.1.99.2 !H  *  !H
!
!R1#show route-map
route-map POLICY_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.128.2
  Policy routing matches: 9 packets, 540 bytes
route-map POLICY_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.99.2
```

https://t.me/learningnets

```
  Policy routing matches: 6 packets, 360 bytes
route-map POLICY_ROUTING, permit, sequence 30
  Match clauses:
  Set clauses:
    ip default next-hop 136.1.99.3 Policy routing matches: 15 packets, 900 bytes
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Local Policy Routing

Task

- R1-generated traffic should be engineered as follows:
  - If destination is SW1's Loopback0, forward it out VLAN 123 to R2.
  - If destination is R3's Loopback0, forward it out VLAN 128 to R2.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Local policy routing is very similar to regular policy routing, but it affects only traffic originated by the router, not transit traffic. Configuration is identical; the only difference is that PBR policy is applied globally on the router using the command `ip local policy route-map <rmap_name>` , instead of at the interface level.

Configuration

R1:

```
access-list 100 permit ip any host 150.1.9.9
access-list 150 permit ip any host 150.1.3.3
!
route-map LOCAL_ROUTING permit 10
```

https://t.me/learningnets

```
 match ip address 100
 set ip next-hop 136.1.99.2
!
route-map LOCAL_ROUTING permit 20
 match ip address 150
 set ip next-hop 136.1.128.2
!
ip local policy route-map LOCAL_ROUTING
```

Verification

Verif the local policy routing configuration, and note counters within the route-map:

```
R1#show ip policy
Interface      Route map local           LOCAL_ROUTING
!
!R1#show route-map
route-map LOCAL_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.123.2 Policy routing matches: 0 packets, 0 bytes
route-map LOCAL_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.128.2 Policy routing matches: 0 packets, 0 bytes
```

Initiate traffic from R1 toward R3's Loopback0, and note the path and check the route-map counters:

```
R1#traceroute 150.1.3.3
Type escape sequence to abort.
Tracing the route to 150.1.3.3
VRF info: (vrf in name/id, vrf out name/id) 1 136.1.128.2 0 msec 0 msec 0 msec
  2 136.1.99.3 4 msec *  0 msec
!
!R1#show route-map
route-map LOCAL_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.123.2
```

https://t.me/learningnets

```
  Policy routing matches: 0 packets, 0 bytes
route-map LOCAL_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.128.2 Policy routing matches: 6 packets, 492 bytes
```

Initiate traffic from R1 toward SW1's Loopback0, and note the path and check the route-map counters:

```
R1#traceroute 150.1.9.9
Type escape sequence to abort.
Tracing the route to 150.1.9.9
VRF info: (vrf in name/id, vrf out name/id) 1 136.1.99.2 0 msec 0 msec 0 msec
  2 136.1.128.9 4 msec *  4 msec
!
!R1#show route-map
route-map LOCAL_ROUTING, permit, sequence 10
  Match clauses:
    ip address (access-lists): 100
  Set clauses:
    ip next-hop 136.1.99.2 Policy routing matches: 6 packets, 474 bytes

route-map LOCAL_ROUTING, permit, sequence 20
  Match clauses:
    ip address (access-lists): 150
  Set clauses:
    ip next-hop 136.1.128.2
  Policy routing matches: 6 packets, 492 bytes
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Classification and Marking

Task

- Configure an outbound policy on R1's VLAN 19 to satisfy the requirements:
  - Telnet return traffic from VLAN 27 should be marked with an IP precedence of 2.
  - ICMP packets larger than 1200 bytes should be dropped.
  - Voice packets with UDP ports in the destination range of 16384-32767 and layer3 packet size of 64 bytes should be marked with DSCP EF.
  - All other traffic received inbound on its VLAN 123 interface with a precedence of 0 should be remarked with a precedence of 1.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **configure replace nvram:startup-config** instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

The Modular Quality of Service Command Line Interface (MQC) unifies all IOS QoS features under a single language. MQC allows the implementation of a full suite of QoS tools, including classification, congestion management, traffic metering, marking, traffic shaping, and link efficiency. The main advantage of using the MQC over some legacy methods is that multiple QoS features can be applied to the same interface in the same direction.

https://t.me/learningnets

Classification in MQC uses case-sensitive class-maps for traffic matching. Each class-map performs a logical AND (match-all) or a logical OR (match-any) on its criteria. In other words, in a match-all class-map, all matches (conditions) must be true for the class to be true. Class-maps can be nested inside other maps to build complicated classification of type AND-OR. If multiple match criteria appear on the same line (such as match ip dscp or match ip precedence) they are treated as a logical OR, whereas if they appear on different lines they are treated as a logical AND.

Different IOS versions and platforms support different matches in the class-map, but as a general rule the following classification criteria are supported:

- Named and numbered access-lists, allows matching of IP addresses, TCP/UDP ports, IP protocol numbers.
- Layer 3 packet length.
- Layer 2 addresses, allows matching of source/destination MAC address.
- Packet marking, allows matching on layer2 CoS, layer3 DSCP/IP precedence.
- Network-Based Application Recognition (NBAR).
- Inverse logical matching (logical NOT).

When classification is configured in a class-map, actions are defined for the different classes in a case-sensitive policy-map. A policy map is an ordered list of class-maps with their corresponding actions, similar to how a route-map works. The router matches packets entering/leaving the interface against all class-map entries in the respective input/output policy-map on the interface in a top-down fashion. This means that the first match in a class-map is used for classification, which implies that the order of the classes called in the policy-map is significant. The policy-map actions include marking, shaping, policing, assigning queue weight, compressing, etc. Any unclassified traffic in a policymap falls into the class-default category.

Configuration

```
R1:
ip access-list extended TELNET
 permit tcp 136.1.27.0 0.0.0.255 eq 23 any
!
ip access-list extended VOICE
 permit udp any any range 16384 32767
!
class-map match-any TELNET
 match access-group name TELNET
!
class-map match-all ICMP
```

https://t.me/learningnets

```
 match protocol icmp
 match packet length min 1201
!
class-map match-all VOICE
 match access-group name VOICE
 match packet length min 64 max 64
!
class-map match-all OTHER
 match input-interface GigabitEthernet0/1
 match ip precedence 0
!
policy-map MARKING
 class TELNET
  set ip precedence 2
 class ICMP
  drop
 class VOICE
  set ip dscp ef
 class OTHER
  set ip precedence 1
!
interface GigabitEthernet0/0
 service-policy output MARKING
```

Verification

Verify the policy-map settings:

```
R1#show policy-map interface gigabitEthernet0/0 output
 GigabitEthernet0/0

  Service-policy output: MARKING
 Class-map: TELNET (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps Match: access-group name TELNET
        0 packets, 0 bytes
        5 minute rate 0 bps
      QoS Set
        precedence 2
          Packets marked 0
 Class-map: ICMP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps Match: protocol icmp
 Match: packet length min 1201
```

https://t.me/learningnets

```
      drop
 Class-map: VOICE (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps Match: access-group name VOICE
 Match: packet length min 64 max 64
      QoS Set
        dscp ef
          Packets marked 0
 Class-map: OTHER (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps Match: input-interface GigabitEthernet0/1
 Match: ip precedence 0
      QoS Set
        precedence 1
          Packets marked 0
 Class-map: class-default (match-any)

      11 packets, 1122 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any
```

Check telnet traffic marking:

```
SW2#telnet 136.1.27.3
Trying 136.1.27.3 ... Open


User Access Verification

Password: R3>
!
!R1#show policy-map interface gigabitEthernet0/0 output | section TELNET
    Class-map: TELNET (match-any)
      22 packets, 1278 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: access-group name TELNET 22 packets, 1278 bytes
        5 minute rate 0 bps
      QoS Set
        precedence 2 Packets marked 22
```

Check ICMP traffic drop:

```
R3#ping 150.1.10.10 size 1200
```

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 1200-byte ICMP Echos to 150.1.10.10, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
!
!R3#ping 150.1.10.10 size 1201

Type escape sequence to abort.
Sending 5, 1201-byte ICMP Echos to 150.1.10.10, timeout is 2 seconds:.....
Success rate is 0 percent (0/5)
!
!R1#show policy-map interface gigabitEthernet0/0 output | section ICMP
    Class-map: ICMP (match-all)  5 packets, 6075 bytes
5 minute offered rate 2000 bps, drop rate 2000 bps

      Match: protocol icmp
      Match: packet length min 1201
      drop
```

Previously issued ICMP packets with a size of 1200 were sent by default with a precedence of 0, so check the policy-map:

```
R1#show policy-map interface gigabitEthernet0/0 output | section OTHER
    Class-map: OTHER (match-all)  5 packets, 6070 bytes
5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: input-interface GigabitEthernet0/1
      Match: ip precedence 0
      QoS Set
        precedence 1 Packets marked 5
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 6: Perimeter Security and Services - IOS Firewall

IOS Time-Based Access-Lists

Task

- Configure filtering on R1 as follows:
  - Restrict access from VLAN 128 to 150.1.0.0/16 subnet on weekends.
  - Restrict access from VLAN 128 to VLAN 19 between 4PM and 10AM on weekdays.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 6 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the ASA.

Overview

Time-based ACLs allow specific entries within an access-list to be valid/active only for specified time intervals; only extended access-lists are supported. There are two types of time-ranges:

https://t.me/learningnets

- Absolute: Define non-recurring time intervals (such as from Feb 5 to Feb 9), with fixed starting and ending dates/times.
- Periodic: Define recurring time intervals that may repeat infinitely (such as every Monday from 10AM to 11AM). Based on the concept of weekdays: You may define periods based on days of the week (Mon, Tue, Fri, and so on) and specific ranges such as daily, weekends, and weekdays.

It is important to note that time-ranges start (become active) with the first second of the first minute in the range and end with the last second of the last minute in the range. For example, the interval from 00:01 to 02:00 will last from 00:01:00 to 02:00:59(hh:mm:ss). Therefore, if you want some interval to end at exactly 4:00pm, use the value 15:59 (note that IOS uses the 24-hour time format).

Configuration

R1:

```
time-range WEEKENDS
 periodic weekend 0:00 to 23:59
!
time-range WEEKDAYS
 periodic weekdays 16:00 to 23:59
 periodic weekdays 0:00 to 9:59
!
ip access-list extended TIME_BASED_ACL
 deny ip 136.1.128.0 0.0.0.255 150.1.0.0 0.0.255.255 time-range WEEKENDS
 deny ip 136.1.128.0 0.0.0.255 136.1.19.0 0.0.0.255 time-range WEEKDAYS
 permit ip any any
!
interface GigabitEthernet0/1.128
 ip access-group TIME_BASED_ACL in
```

Verification

Verify the time-ranges and which one is active, based on the system current time:

```
R1#show time-range
time-range entry: WEEKDAYS (active)
   periodic weekdays 16:00 to 23:59
   periodic weekdays 0:00 to 9:59
   used in: IP ACL entrytime-range entry: WEEKENDS (inactive)
```

https://t.me/learningnets

```
    periodic weekend 0:00 to 23:59
    used in: IP ACL entry
!
!R1#show clock
06:44:43.155 UTC Tue Dec 3 2013
```

Based on the current system time, only access to 136.1.19.0/24 should be restricted:

```
SW1#ping 136.1.19.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.10, timeout is 2 seconds:U.U.U
Success rate is 0 percent (0/5)
!
!SW1#ping 136.1.19.10 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.19.10, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
!
!R1#show ip access-lists
Extended IP access list TIME_BASED_ACL
    10 deny ip 136.1.128.0 0.0.0.255 150.1.0.0 0.0.255.255 time-range WEEKENDS (inactive)
20 deny ip 136.1.128.0 0.0.0.255 136.1.19.0 0.0.0.255 time-range WEEKDAYS (active) (5 matches)

    30 permit ip any any (93 matches)
```

Verify that access to 150.1.0.0/16 is not restricted:

https://t.me/learningnets

```
SW1#ping 150.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
!
!SW1#ping 150.1.1.1 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 150.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 150.1.9.9 !!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/9 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - CCIE Security V4 Lab Resources

Security Rack Rental and Hardware Guide

Click here for the Security V4.0 Rack Rental and Hardware Guides

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

Section 8 Introduction

This section utilizes Security Services and Wireless Security.

Lab Topics

The Cisco Unified Wireless Network (UWN) is a complex architecture with multiple components and functionalities; within the scope of the CCIE Security lab exam, only a subset of these are tested, mostly relevant to WLC (Wireless LAN Controller) and LAP (Lightweight Access Point), and UWN integration with ISE/ACS. The following high-level topics are provided by Cisco as *blueprint* topics:

- Wireless LAN Controller 2504
  - Wireless Security on AP and WLC
  - EAP methods
  - WPA/WPA-2
  - WIPS

As with previous topics, we can expand upon this blueprint with the following checklist:

Technology	Mastered?
Understanding Cisco Unified Wireless Architecture	
Understanding WPA and WPA-2 Requirements	
Lightweight Access Point Local Mode versus FlexConnect Mode	

https://t.me/learningnets

Technology	Mastered?
Lightweight Access Point Authorization	
802.1x for Lightweight Access Point	
802.1x for Wireless Clients	
Local Web Authentication versus Central Web Authentication	
Bring Your Own Device with Native Supplicant Provisioning	
Wireless Guest Services	
Wireless Rogue Device Management	
Wireless LAN Controller Administration	
Wireless Intrusion Prevention System	

Configurations

Initial configuration files are provided at the top of this page. Each lab has an initial configuration that is assumed. The initial configuration in some cases is the configuration that the previous lab ended with. If there is only one configuration file, the other devices did not change in the previous lab.

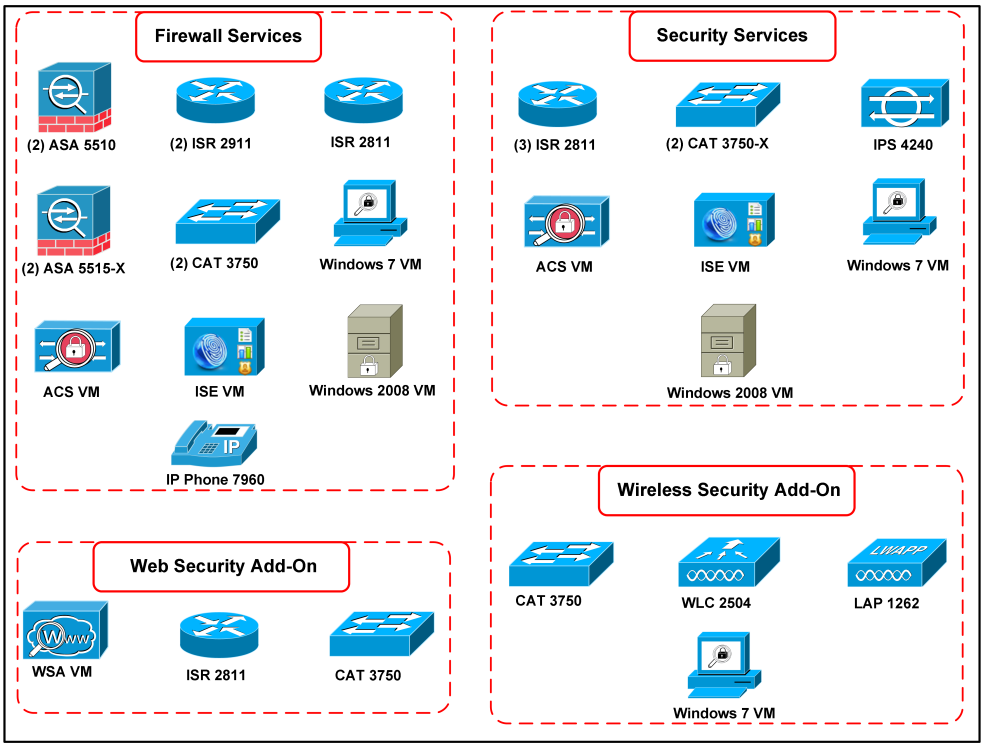
Load the Section 8 Initial Configuration Files to initialize your rack. Pay attention to the instructions in each task, because you may be required to load a different set of initial configuration files.

Section 8 General Lab Topology

The INE CCIE Security Racks are divided into modules as seen in the figure below.

https://t.me/learningnets

This section, Section 8, utilizes Security Services and Wireless Security.



In this section, you will focus on how to provide various security services using WLC and LAP. As you work on each task, focus on the small elements of building each technology. Gaining an in-depth knowledge of the workings of small tasks will lead to better troubleshooting of large scenarios in the CCIE Lab exam, as well in as production environments. Pay close attention to the topology, which may change.

Using Your Own Equipment

If you use your own rack to practice these labs, it is recommended that you use the following hardware/software versions to obtain the same results and outputs:

- 2811 is running 15.1(3)T2 with Advanced Enterprise feature set
- 2911 is running 15.2(3)T3 with Security license
- 3750 is running 12.2(55)SE7 with IP Services feature set
- 3750X is running 15.0(2)SE2 with IP Services license
- ASA5500 is running 8.2(5) or 8.4(5) along with ASDM 7.1(2)
- ASA5500-X is running 8.6(1) along with ASDM 7.1(2)
- AnyConnect is running version 3.0.11042
- VPN Client is running version 5.0.07.0440
- WSA is running version 7.1.3-031 along with AD Agent 1.0.0.32.1
- WLC is running version 7.2.111.3
- IPS is running version 7.1(7)E4
- ISE is running version 1.1.3.124

https://t.me/learningnets

- ACS is running version 5.3
- NAC Agent is running version 4.9.049

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC Initialization

Task

- Initialize the WLC using the following values:
  - Hostname **WLC**.
  - Admin username/password of **admin/default1A**.
  - IP address 172.16.5.100/24 with a default gateway of 172.16.5.1.
  - The **admin** user should have a password of **default1A**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.

Overview

While in the CCIE lab, you should receive the WLC pre-configured with basic settings. It is still recommended that you know the initialization process and its requirements. At a minimum, you must configure the following from the CLI setup wizard:

- system name (hostname)
- One user with admin rights; make sure to use admin/default1A as the username/password for INE racks
- Management interface IP address, mask, and default gateway
- Physical port used by the management interface; make sure to use port 1 for INE racks
- Virtual interface IP address (used for DHCP relay and local web authentication process)
- Mobility RF group name (used for scenarios not relevant to CCIE Security lab exam to enhance roaming, when multiple WLCs share same LAPs and clients)

https://t.me/learningnets

- One WLAN SSID name (SSID name is case sensitive)

You can leave the rest at default values, unless instructed otherwise. Make sure to leave the default country code, which is US; otherwise, the LAP may not be able to join the WLC. All these settings can be modified later from the GUI.

Configuration

Log in to the WLC console with the username/password of admin/default1A and issue the `clear config` command, followed by the `reset system` command; do not save configuration changes.

```
User:  admin
Password:*********(Cisco Controller) >clear config
Are you sure you want to clear the configuration? (y/n)y

Configuration Cleared!(Cisco Controller) >reset system

The system has unsaved changes.Would you like to save them now? (y/N)n


Configuration Not Saved!Are you sure you would like to reset the system? (y/N)y

System will now restart!Restarting system.
```

After the system restarts, make sure to start configuration within 30 seconds; otherwise, WLC will start the DHCP client and autoinstall functionality, which will load the configuration from the TFTP server:

```
Welcome to the Cisco Wizard Configuration Tool
Use the '-' character to backup

Would you like to terminate autoinstall? [yes]: yes

AUTO-INSTALL: process terminated -- no configuration loaded

System Name [Cisco_37:b2:c4] (31 characters max): WLC
Enter Administrative User Name (24 characters max): admin
Enter Administrative Password (3 to 24 characters): *********
Re-enter Administrative Password                : *********

Management Interface IP Address: 172.16.5.100
Management Interface Netmask: 255.255.255.0
```

https://t.me/learningnets

```
Management Interface Default Router: 172.16.5.1
Management Interface VLAN Identifier (0 = untagged):
Management Interface Port Num [1 to 4]: 1
Management Interface DHCP Server IP Address: 172.16.5.1
Virtual Gateway IP Address: 1.1.1.1
Mobility/RF Group Name: CCIE-SECURITY
Network Name (SSID): CCIELAB
Configure DHCP Bridging Mode [yes][NO]:
Allow Static IP Addresses [YES][no]:
Configure a RADIUS Server now? [YES][no]: no
Warning! The default WLAN security policy requires a RADIUS server.
Please see documentation for more details.

Enter Country Code list (enter 'help' for a list of countries) [US]:

Enable 802.11b Network [YES][no]:
Enable 802.11a Network [YES][no]:
Enable 802.11g Network [YES][no]:
Enable Auto-RF [YES][no]:
Configure a NTP server now? [YES][no]: no
Configure the system time now? [YES][no]: no

Warning! No AP will come up unless the time is set.
Please see documentation for more details.
Configuration correct? If yes, system will save it and reset. [yes][NO]: yes

Configuration saved!Resetting system with new configuration...
```

When finished with the configuration wizard, the system will reset and you can log in to the GUI using the configured username/password of admin/default1A.

Verification

Verify the physical and logical interface status/configuration:

```
(Cisco Controller) >show  port summary

           STP   Admin   Physical   Physical   Link   Link
Pr  Type   Stat   Mode     Mode      Status   Status  Trap     POE
-- ------- ---- ------- ---------- ---------- ------ ------- -------
1  Normal  Forw Enable  Auto       100 Full   Up     Enable  N/A
2  Normal  Disa Enable  Auto       Auto       Down   Enable  N/A
3  Normal  Disa Enable  Auto       Auto       Down   Enable  Enable  (Power Off)
4  Normal  Disa Enable  Auto       Auto       Down   Enable  Enable  (Power Off)
```

https://t.me/learningnets

```
!
!(Cisco Controller) >show interface summary

Interface Name                   Port Vlan Id  IP Address      Type    Ap Mgr Guest
------------------------------ ---- -------- --------------- ------- ------ -----
management                       1    untagged 172.16.5.100    Static  Yes    No
virtual                          N/A  N/A      1.1.1.1         Static  No     No
```

Verify what WLANs exist (one configured from the setup wizard):

```
(Cisco Controller) >show wlan summary

Number of WLANs.............................. 1

WLAN ID  WLAN Profile Name / SSID               Status    Interface Name
-------  -----------------------------------  --------  --------------------
1        CCIELAB / CCIELAB                     Enabled   management
```

Verify that LAP joined the WLC, and that the DTLS session for control-plane is UP (the AP default name is based on its Ethernet MAC address, so it may differ on your rack):

```
(Cisco Controller) >show ap summary

Number of APs................................ 1

Global AP User Name.......................... Not Configured
Global AP Dot1x User Name.................... Not Configured

AP Name             Slots  AP Model              Ethernet MAC       Location          Port  Country  Priority
------------------  -----  --------------------  -----------------  ----------------  ----  -------  --------
AP0006.f63b.9ada    2      AIR-LAP1262N-A-K9     00:06:f6:3b:9a:da  default location  1        US       1
!
!(Cisco Controller) >show ap join stats summary  all

Number of APs........................................ 1

Base Mac             AP EthernetMac       AP Name                 IP Address         Status
f0:29:29:04:21:f0    00:06:f6:3b:9a:da    AP0006.f63b.9ada        136.1.47.100       Joined
!
!(Cisco Controller) >show dtls connections

       AP Name         Local Port        Peer IP        Peer Port               Ciphersuite
```

https://t.me/learningnets

```
------------------- ------------ ---------------  ------------     ------------------------------
AP0006.f63b.9ada      Capwap_Ctrl    136.1.47.100      47534           TLS_RSA_WITH_AES_128_CBC_SHA
```

The WLC 2500 model does not support DTLS for the data-plane session, so it appears as disabled (data-plane encryption is disabled by default; it can be enabled at the LAP level):

```
(Cisco Controller) >show ap data-plane all

                        Min Data          Data          Max Data     Last
AP Name                Round Trip      Round Trip      Round Trip    Update
------------------  --------------  --------------  --------------  ------
AP0006.f63b.9ada           0.000s          0.000s       0.000s     23:18:20
!
!(Cisco Controller) >show ap link-encryption all

             Encryption  Dnstream  Upstream    Last
AP Name           State     Count     Count  Update
------------------  ---  --------  --------  ------AP0006.f63b.9ada     Dis         0          0   23:18
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC TACACS Authentication

Task

- Configure WLC to authenticate management users via TACACS:
  - User **ccieadmin** should be authorized for full access.
  - Take configuration changes performed on WLC into account.
  - If TACACS server is unavailable, authenticate using local user database.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload, or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

https://t.me/learningnets

WLC can authenticate management users locally, via TACACS+ or RADIUS, and the order of preference is configurable; this applies to both GUI and CLI access. If TACACS+ is being used, command authorization and command accounting are supported. Traditionally TACACS+ authorization is done using command sets, but because WLC is mostly managed from the GUI, command authorization is performed similar to a role-based access; basically, you will give the user access to all or specified tabs/menus from the WLC GUI. The main menus available in the WLC GUI and relevant to TACACS+ authorization are:

- MONITOR
- WLANS
- CONTROLLER
- WIRELESS
- SECURITY
- MANAGEMENT
- COMMANDS

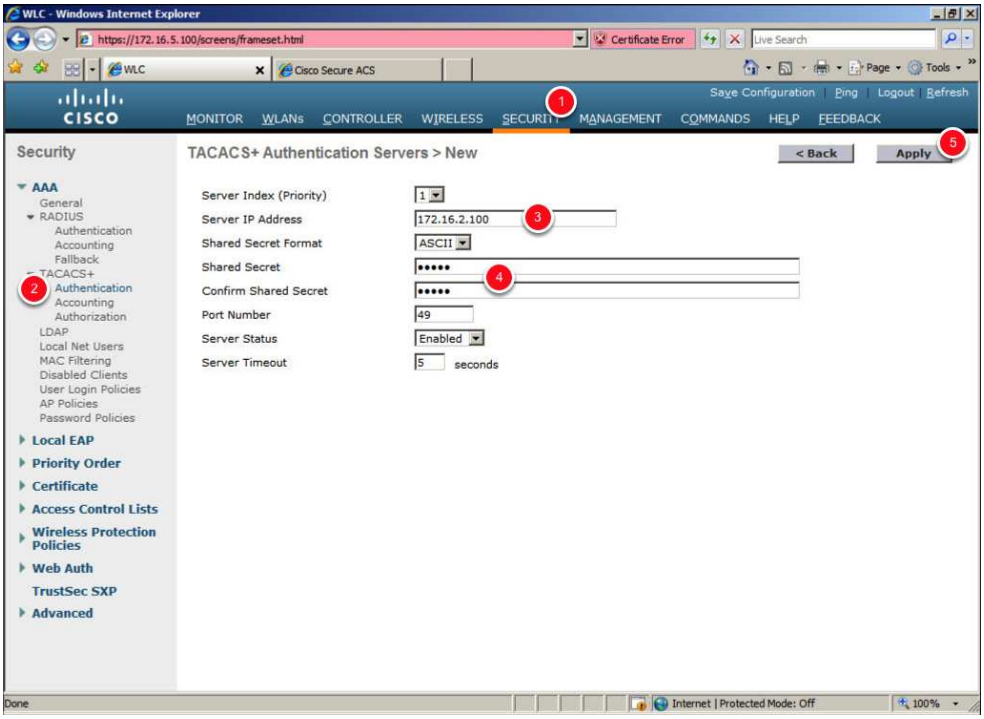
All the above names will be used as roles in the TACACS+ authorization policies, with one or multiple roles being assigned to the same user, allowing access to one or multiple menus. If full access is required, there is a dedicated role with a value of **ALL**. Without TACACS+ integration, a local management user with a role of lobby-admin can be created; this admin will have access only to create local non-management users, necessary for a WLAN that requires a username/password authentication method (for example, if the WLAN does not use layer2 authentication but uses layer3 authentication through the built-in portal of the WLC). If TACACS+ is used for managing the WLC, the lobby-admin role will be authorized using a value of **LOBBY**.

WLC supports up to three TACACS servers. The order of preference is based on the server index (priority) value defined when configuring the TACACS server; a lower index value has a higher priority/preference.

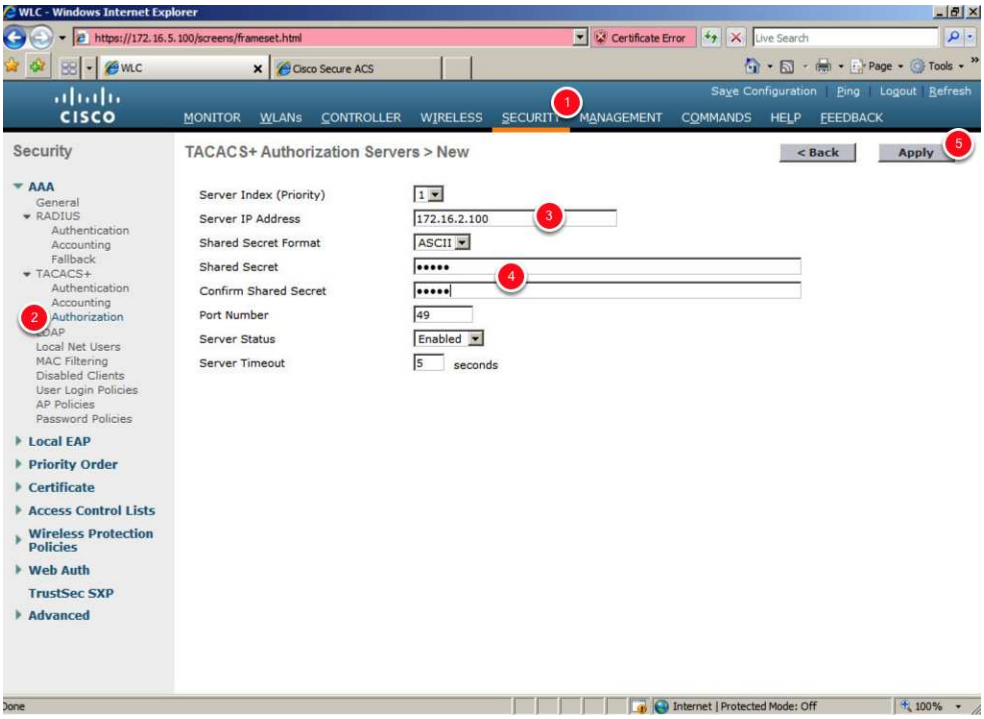
Configuration

Configure WLC to use ACS as authentication server via TACACS+:

https://t.me/learningnets

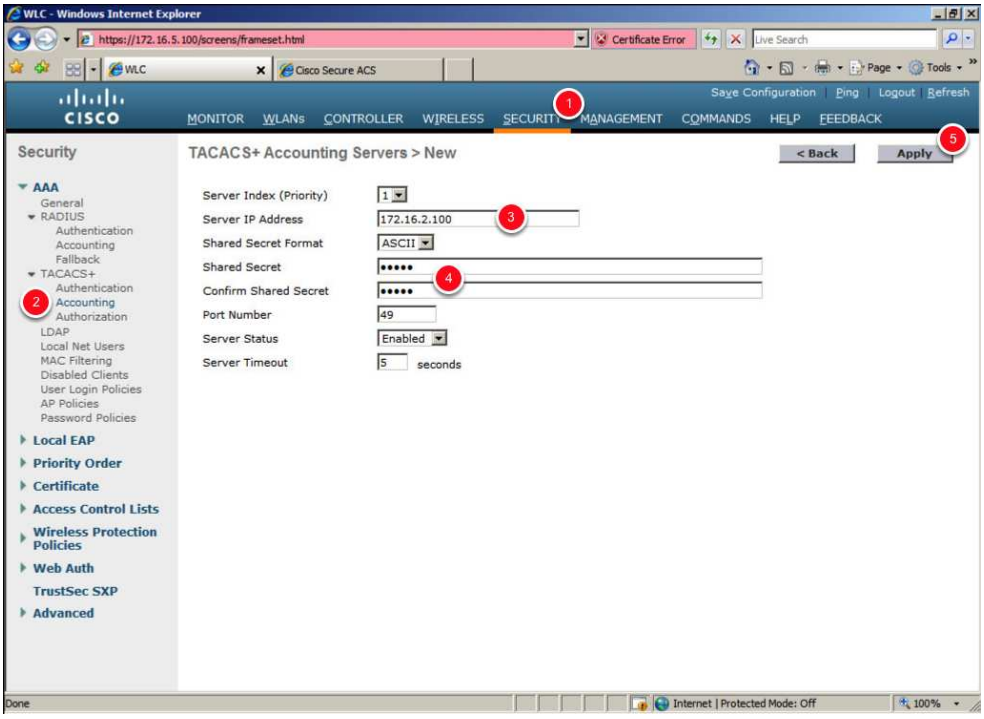


Configure WLC to use ACS as authorization server via TACACS+:

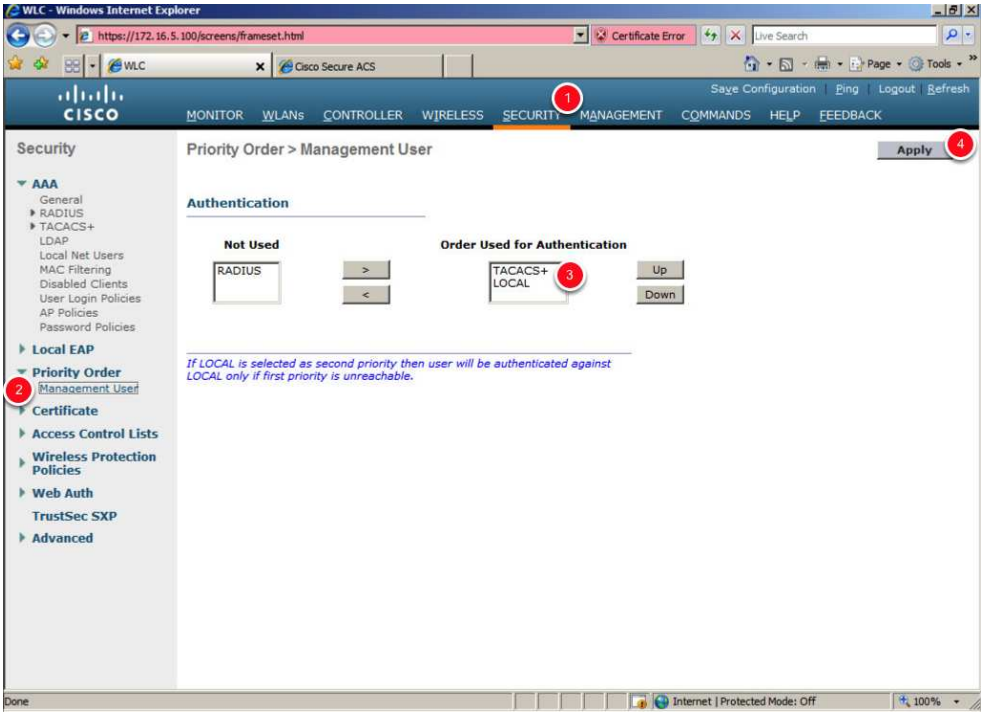


Configure WLC to use ACS as accounting server via TACACS+:

https://t.me/learningnets

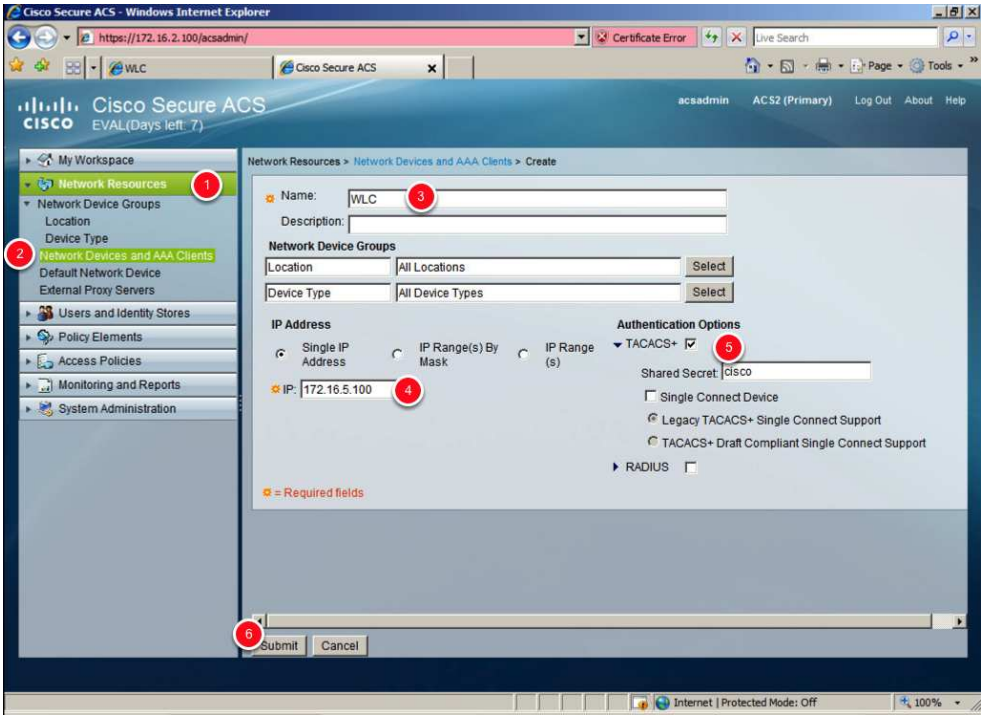


Configure WLC to first use TACACS+ server for management user authentication; if ACS server is not reachable, use the local username database:

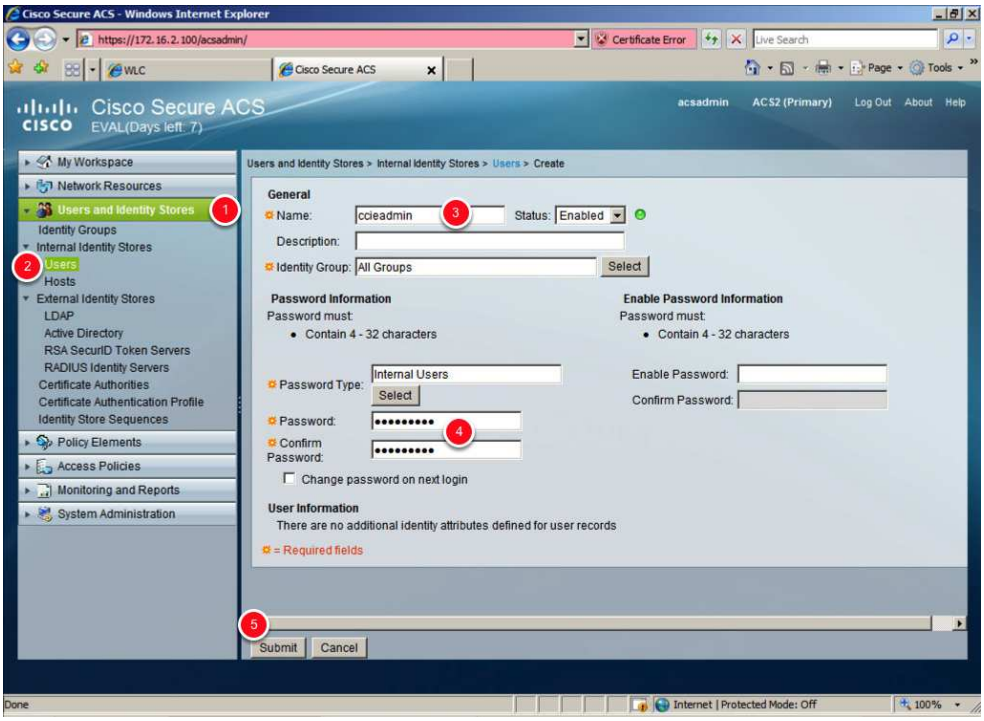


Add WLC as a NAD device into ACS, using TACACS+ protocol:

https://t.me/learningnets

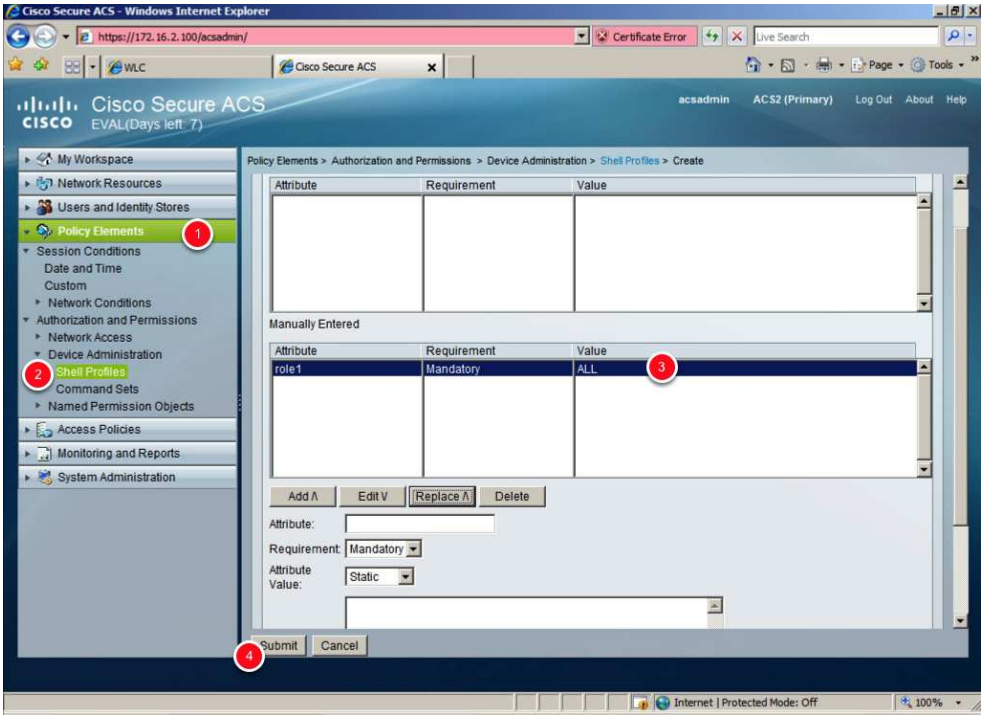


Add the **ccieadmin** user to the internal datastore of ACS:

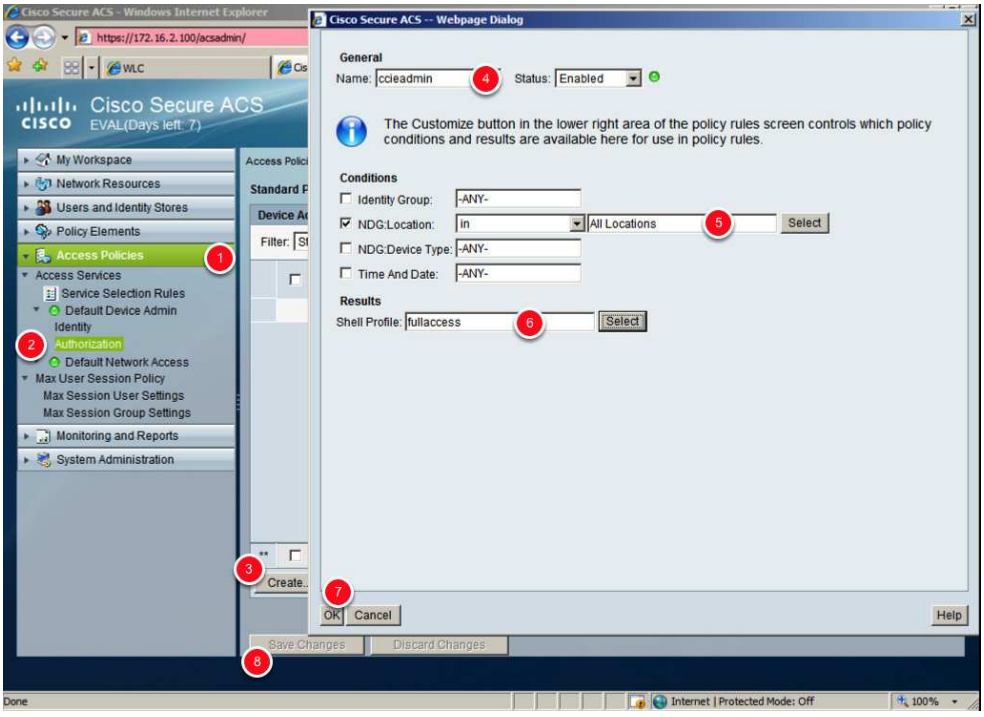


Configure a shell profile, adding a custom mandatory attribute named **role1** with a value of **ALL**:

https://t.me/learningnets



Configure a TACACS authorization access policy, and apply the previously configured shell profile:



Verification

Authenticate using the TACACS user and verify TACACS configuration (authenticate also from GUI to check that you have access to all GUI tabs):

```
User:ccieadmin
```

https://t.me/learningnets

```
Password:********(Cisco Controller) >show tacacs summary


Authentication Servers

Idx  Server Address     Port    State     Tout  MgmtTout
---  ---------------   ------  --------  ----  --------1    172.16.2.100      49       Enabled   10    2

Authorization Servers

Idx  Server Address     Port    State     Tout  MgmtTout
---  ---------------   ------  --------  ----  --------1    172.16.2.100      49       Enabled   10    2

Accounting Servers

Idx  Server Address     Port    State     Tout  MgmtTout
---  ---------------   ------  --------  ----  --------1    172.16.2.100      49       Enabled   5     5
```

Verify the authentication order used for WLC management:

```
(Cisco Controller) >show aaa auth

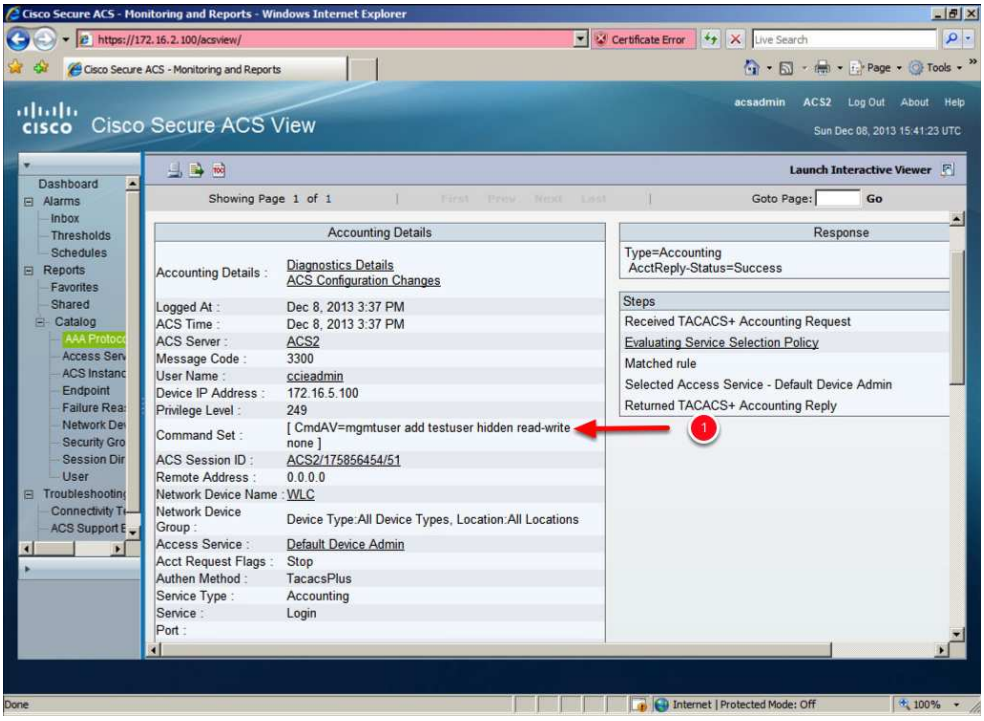

Management authentication server order:
    1..................................... tacacs
    2..................................... local
```

Configure a new management user on WLC and verify the accounting records on ACS:

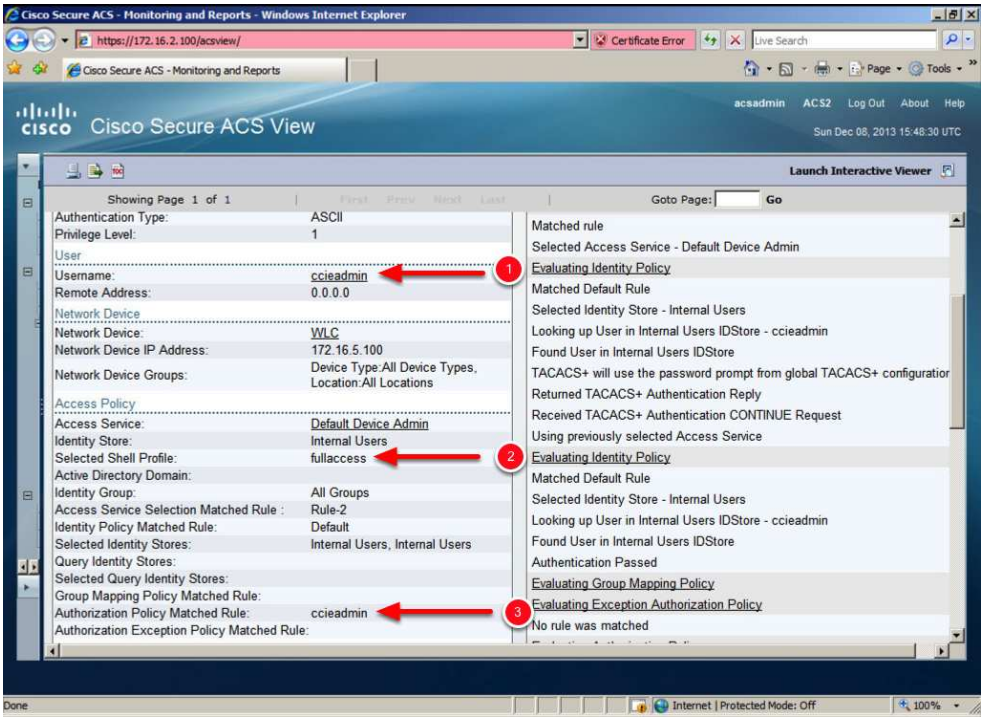
```
(Cisco Controller) >config mgmtuser add testuser testpAss1 read-write
!
!(Cisco Controller) >show mgmtuser

User Name                 Permissions    Description    Password Strength
-----------------------   ------------   --------------------  ------------------
admin                     read-write                                    Strong
cisco                     read-write                                    Strong
testuser                  read-write                                    Strong
```

https://t.me/learningnets

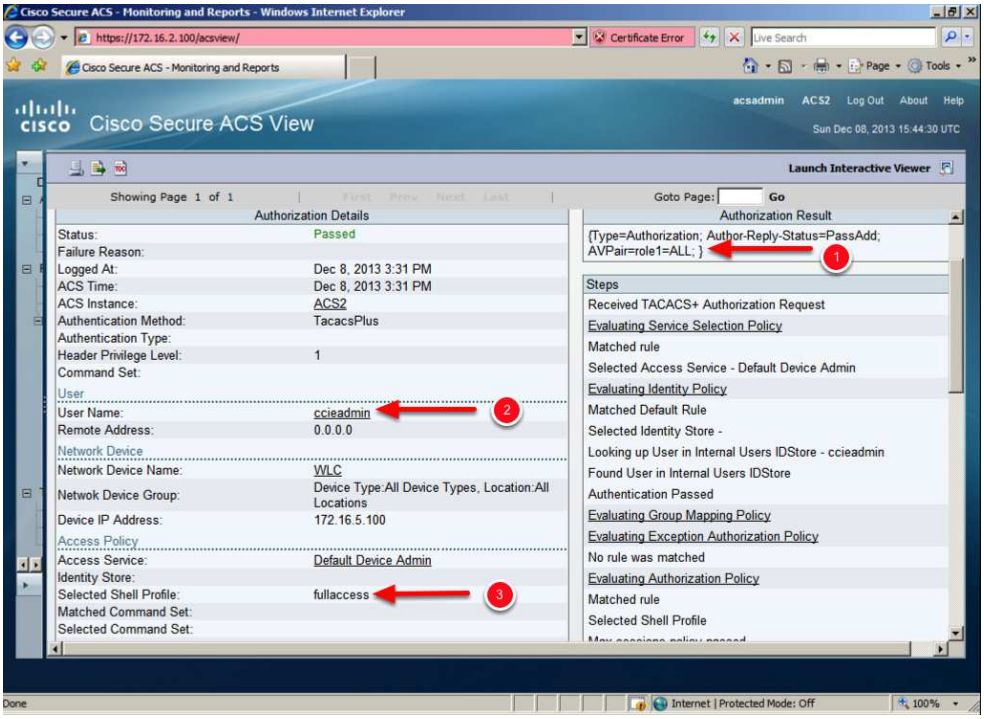


Verify the TACACS authentication records from ACS:



Verify the TACACS authorization records from ACS:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC RADIUS Authentication

Task

- Configure WLC to authenticate management users via RADIUS using ACS:
  - User **ccieadmin** should be authorized for full access.
  - RADIUS server should be used only for management scope.
  - If RADIUS server fails, WLC should probe it every three minutes and activate it only if the probe is successful.
  - If RADIUS server is unavailable, authenticate using local user database.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

WLC can authenticate management users locally, via TACACS+ or RADIUS, and the order of preference is configurable; this applies to both GUI and CLI access. If RADIUS is being used, within the authorization attributes the user role must be received. There are three possible user roles, mapping the roles available when configuring management users in the local database:

https://t.me/learningnets

- readonly
- readwrite
- lobbyadmin

The user role is assigned within RADIUS using the IETF Service-Type attribute, with following values:

- **NAS Prompt** for readonly
- **Administrative** for readwrite
- **Callback Administrative** for lobbyadmin

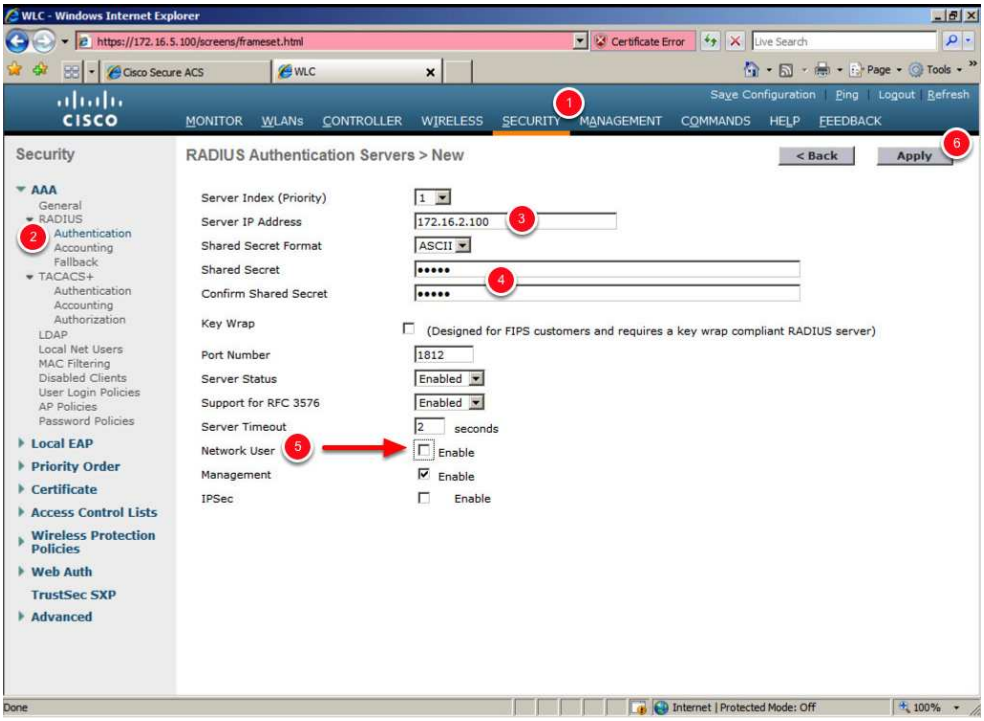
WLC supports the RADIUS server fallback feature, so that when multiple RADIUS servers are defined, it determines whether the primary server is available and switches back to the primary RADIUS server when it is available (the primary server is the one with the lower specified index value). This can be configured in two modes:

- **active** mode: If primary RADIUS server does not respond to the WLC authentication request, WLC marks it as dead and starts sending probe messages periodically (interval is configurable) until that server responds. When the server responds, WLC activates it again and stops sending probe messages. In this mode, when an authentication request is received, WLC always picks the lowest index (highest priority) active server; so a dead server is activated only if it responds to a probe message.
- **passive** mode: If primary RADIUS server does not respond to the WLC authentication request, WLC marks it as dead and sets a timer (configurable). When the timer expires, WLC unconditionally activates the server. In this mode, when an authentication request is received WLC, always picks the lowest index (highest priority) active server. However, a dead server is activated after a specified interval, without being tested for availability. If the activated server does not respond to an authentication request, it is marked again as dead for the configured interval.

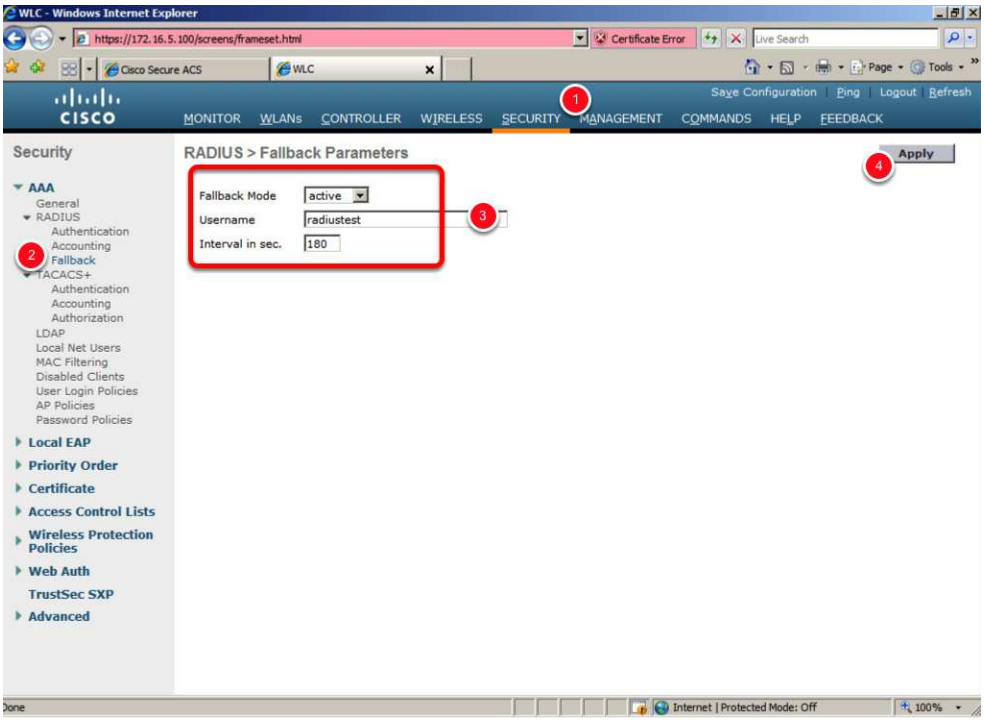
Configuration

Configure WLC to use ACS as authentication server via RADIUS (disable using RADIUS for network users):

https://t.me/learningnets

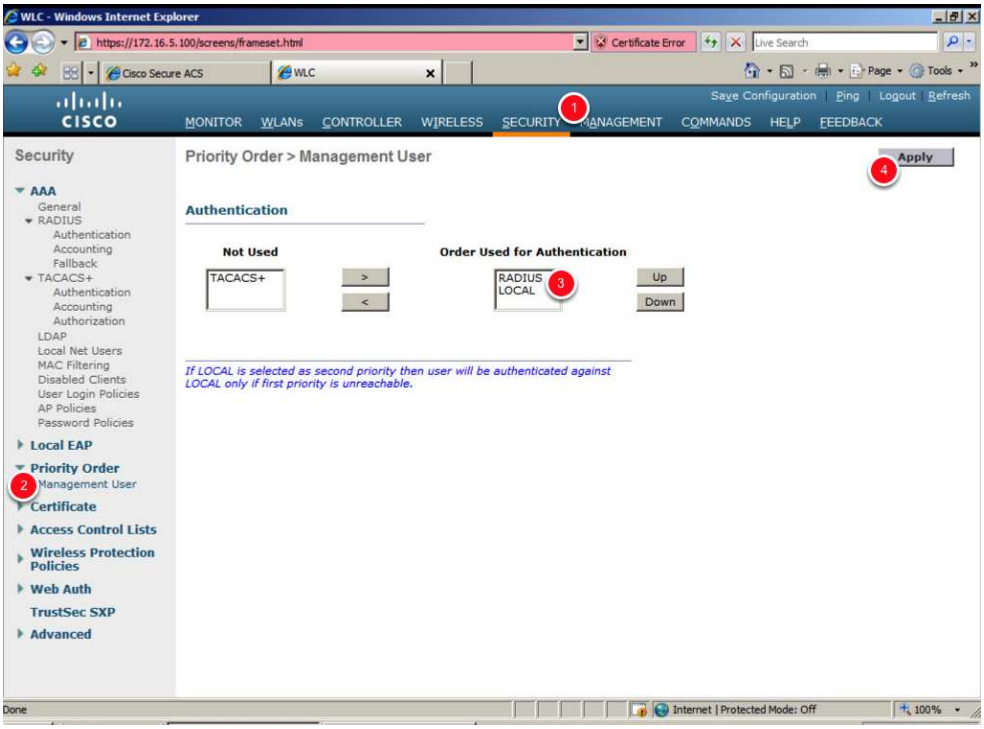


Configure RADIUS fallback in active mode:

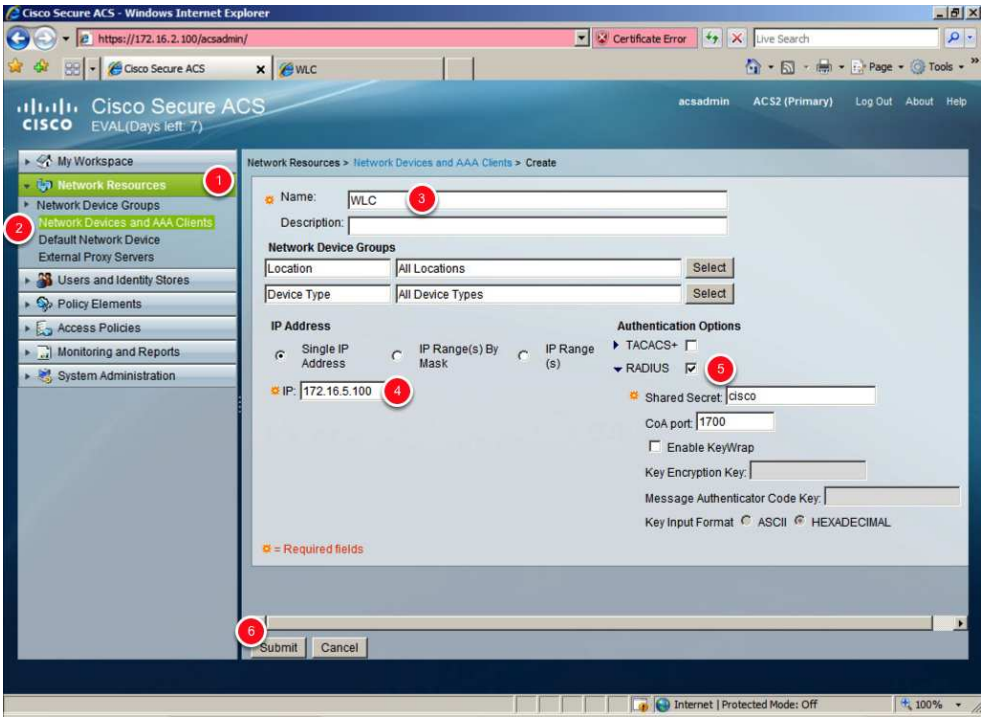


Configure WLC to first use RADIUS server for management user authentication; if ACS server is not reachable, use the local username database:

https://t.me/learningnets

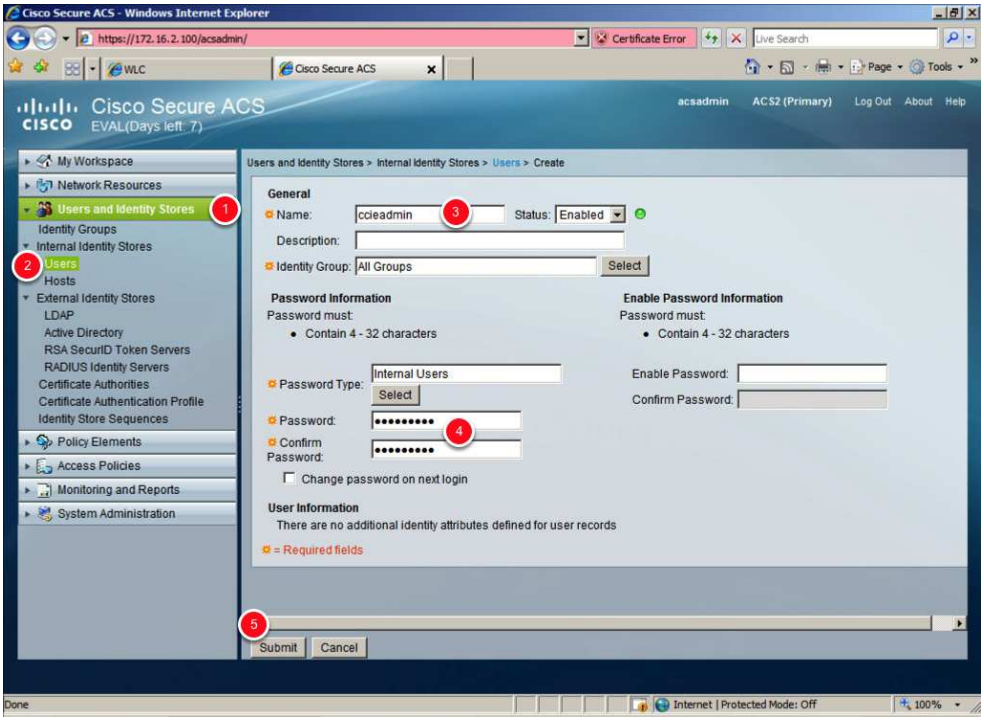


Add WLC as a NAD device into ACS, using RADIUS protocol:

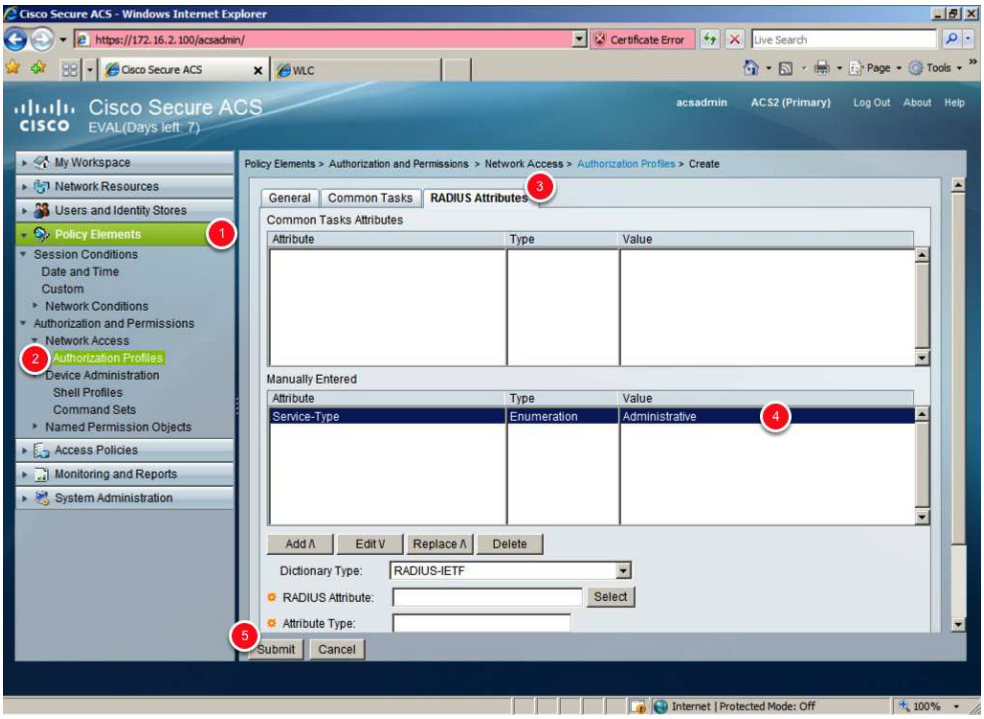


Add the **ccieadmin** user to the internal datastore of ACS:

https://t.me/learningnets

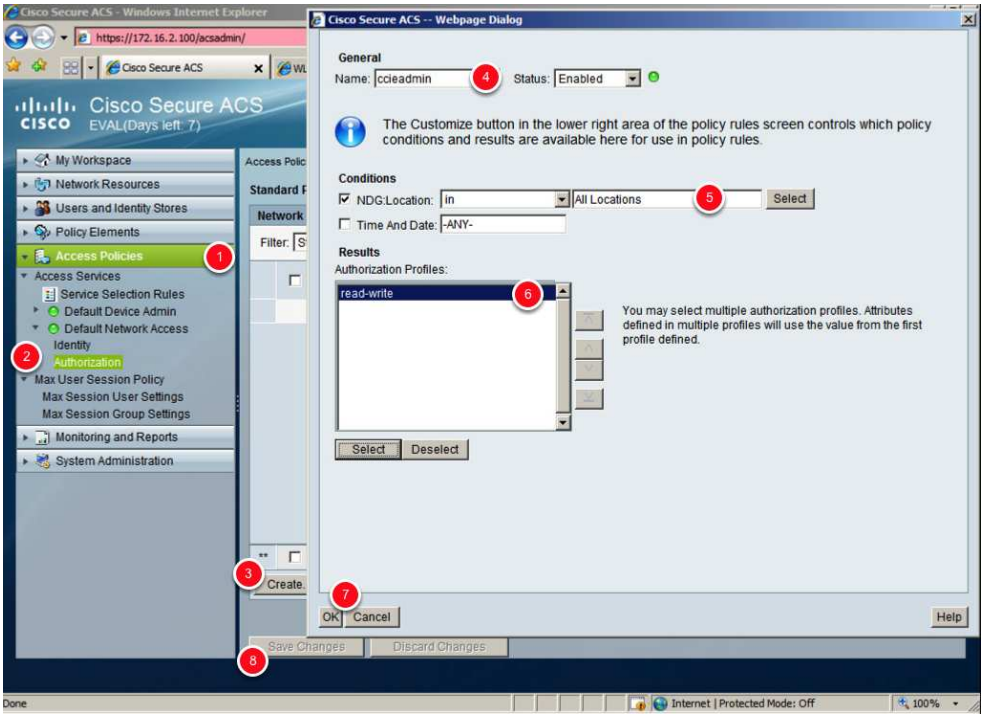


Configure a RADIUS authorization profile and assign the value **Administrative** to the IETF Service-Type attribute:



Configure a RADIUS access policy, and apply the previously configured authorization profile:

https://t.me/learningnets



Verification

Authenticate using the RADIUS user and verify RADIUS configuration (authenticate also from GUI to confirm that you have access to all GUI tabs):

```
User:  ccieadmin
Password:********(Cisco Controller) >show radius summary

Vendor Id Backward Compatibility..................... Disabled
Call Station Id Case................................. lower
Call Station Id Type................................. Mac Address
Aggressive Failover.................................. Enabled
Keywrap.............................................. DisabledFallback Test:
Test Mode.................................... Active
Probe User Name.............................. radiustest
Interval (in seconds)........................ 180
MAC Delimiter for Authentication Messages........ hyphen
MAC Delimiter for Accounting Messages............ hyphen

Authentication Servers

Idx  Type  Server Address    Port    State     Tout  MgmtTout  RFC3576  IPSec - AuthMode/Phase1/Group/Lifetime/Auth/
---  ----  ---------------  ------  --------  ----  --------  -------  ----------------------------------------
1    M     172.16.2.100     1812    Enabled   2     2         Enabled   Disabled - none/unknown/group-0/0 none/none



Accounting Servers
```

https://t.me/learningnets

```
Idx  Type  Server Address    Port    State     Tout  MgmtTout  RFC3576  IPSec - AuthMode/Phase1/Group/Lifetime/Auth/

--More-- or (q)uit
---  ----  ---------------  ------  --------  ----  --------  -------  ----------------------------------------
```

Verify the authentication order used for WLC management:

```
(Cisco Controller) >show aaa auth



Management authentication server order:
    1..................................... radius
    2..................................... local
```

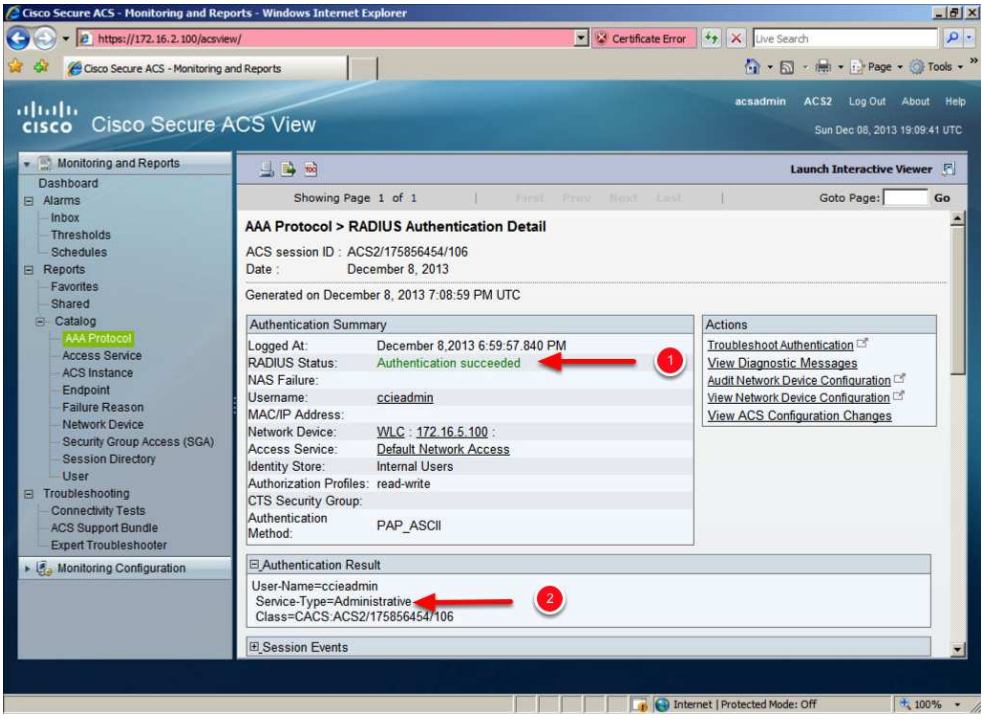
Configure a new management user on WLC to verify read-write access:

```
(Cisco Controller) >config mgmtuser add testuser testpAss1 read-write
!
!(Cisco Controller) >show mgmtuser

User Name                 Permissions    Description    Password Strength
-----------------------   ------------   ---------------------  ------------------
admin                     read-write                                    Strong
cisco                     read-write                                    Strong
testuser                  read-write                                    Strong
```

Verify the RADIUS authentication records from ACS:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC Rogue Management

Task

- Configure rogue device management as follows on WLC:
  - Access-point with MAC address of 00:11:0a:23:32:ab should be whitelisted as friendly.
  - Any access-points broadcasting the SSID of **good** with open authentication should be classified as friendly.
  - Any access-points with minimum 5 clients or broadcasting the SSID of **bad** should be classified as malicious.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command **`clear lwapp private-config`** and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

https://t.me/learningnets

Any wireless-enabled device that is visible within your coverage area but is not under your management is considered rogue. Because radio coverage of an access-point cannot be strictly controlled, it is normal to have rogue devices in most networks. The real problem occurs when the respective devices are not under your management and also:

- Are attached to your wired network, thus offering possibly unathorized access.
- Announce the same SSID's as your infrastructure, thus possibly trying to attract clients and steal their credentials.
- Function in ad-hoc mode.

Cisco Unified Wireless Network (UWN) system has a fully functional built-in mechanism to deal with rogue devices; it's called rogue device management and consists of three phases:

- Rogue Device Detection: Spectrum scanning is used to detect devices not under (UWN) management.
- Rogue Device Classification: Devices can be categorized manually or automatically through configured rules; based on this result, specific actions are taken.
- Rogue Device Mitigation: Actions taken to minimize the security threat.

In the scope of the CCIE Security lab exam, it is not necessary to cover this functionality in detail; a basic understanding and configuration knowledge is sufficient. Essentially, access-points not broadcasting the same RF group-name (used to dynamically change the channel scheme for devices within the same group to achieve the best radio performance by minimizing interference) or part of the same mobility-group (used to facilitate user roaming between access-points) are considered rogue access-points. Rogue clients are clients associated with rogue access-points. RF group-name and mobility-group are values configured when the WLC is initialized.

All rogue devices are automatically classified as **unclassified**, and the other two options are **malicious** and **friendly**. There is one exception to this: If the rogue device is identified as being connected to the wired network as well, it is automatically classified as **malicious**. To conclude, a rogue device can be indentifed based on:

- Default classification (unclassified or malicious)
- Manual administrator classification
- Rogue classification rules defined by the administrator

To detect whether a rogue access-point is attached to the wired network, three methods are being used: Rogue Location Detection Protocol (RLDP), access-points

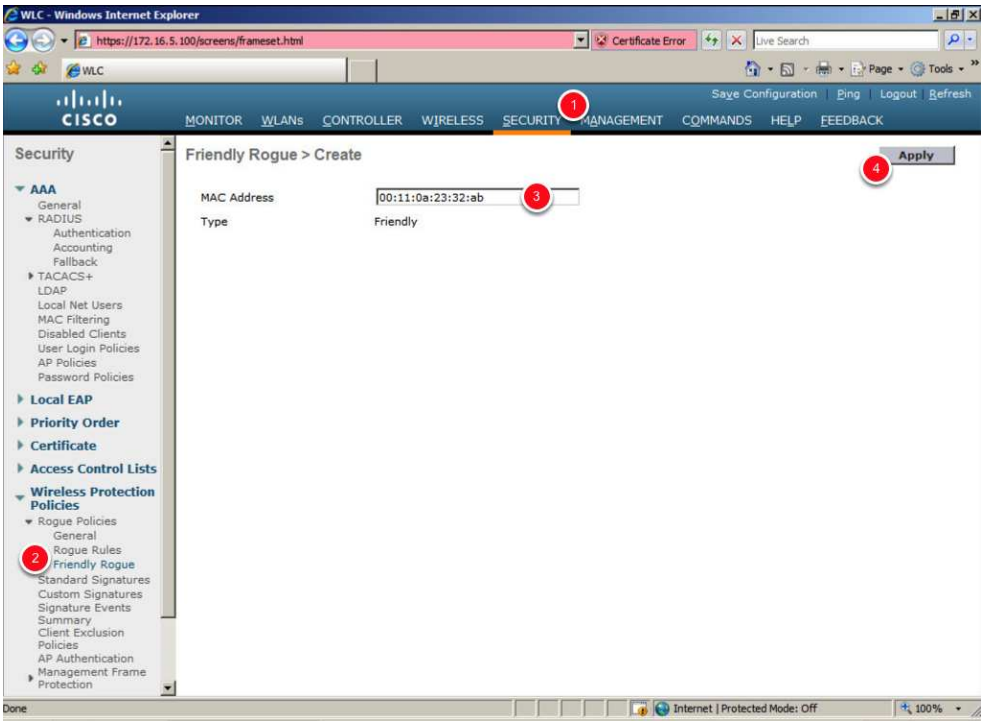
https://t.me/learningnets

running in rogue detector mode, and switchport tracing. Manual administrator classification takes precedence over rogue classification rules. There are two types of rogue classification rules: malicious and friendly. All these rules are processed top-down until first match, with malicious rules being processed first.

Mitigation actions are taken only for devices classsified as **malicious**. The most important and relevant action is rogue containment, which can be automatic or manual and consists of spoofing the rogue access-point and sending deauthentication frames.

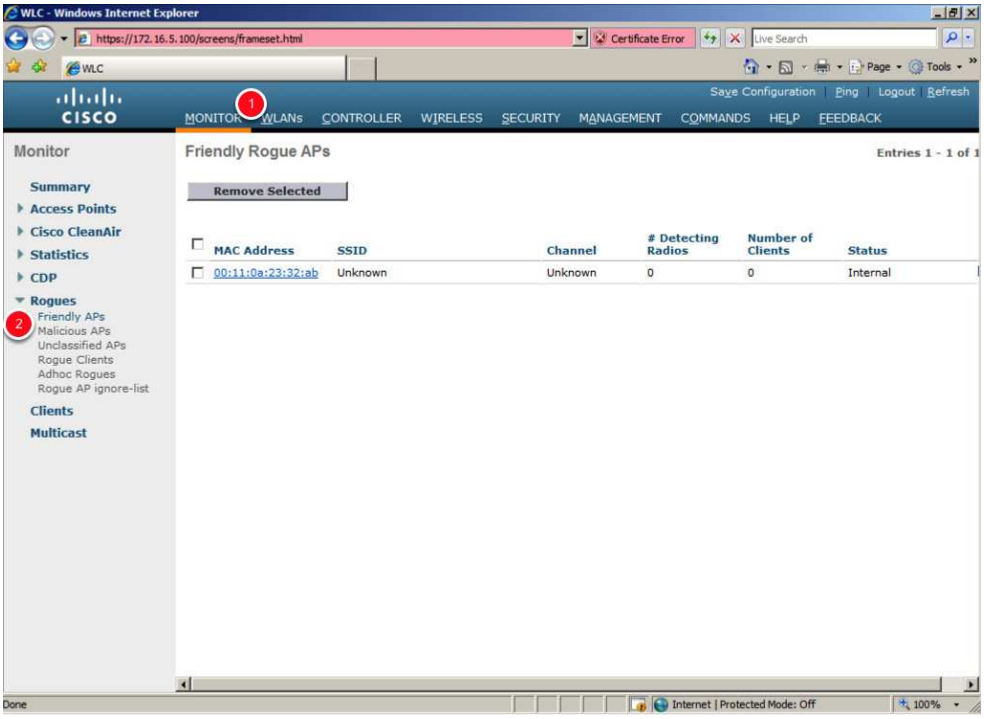
Configuration

Manually configure the access-point with the radio MAC address of 00:11:0a:23:32:ab as friendly:

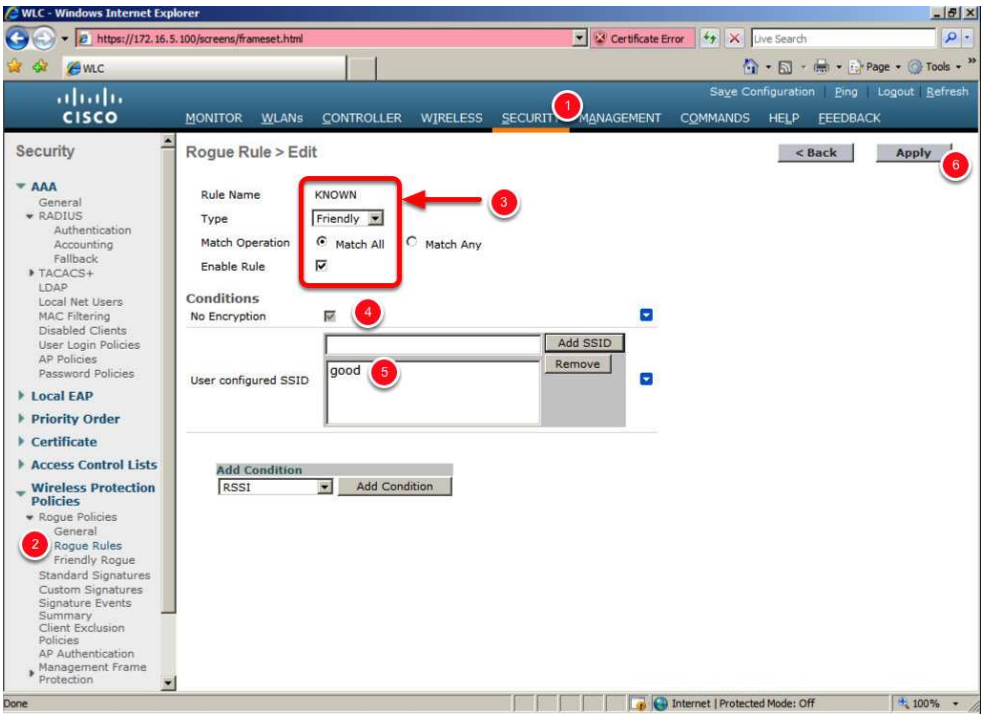


Add a friendly rogue classification rule:

https://t.me/learningnets

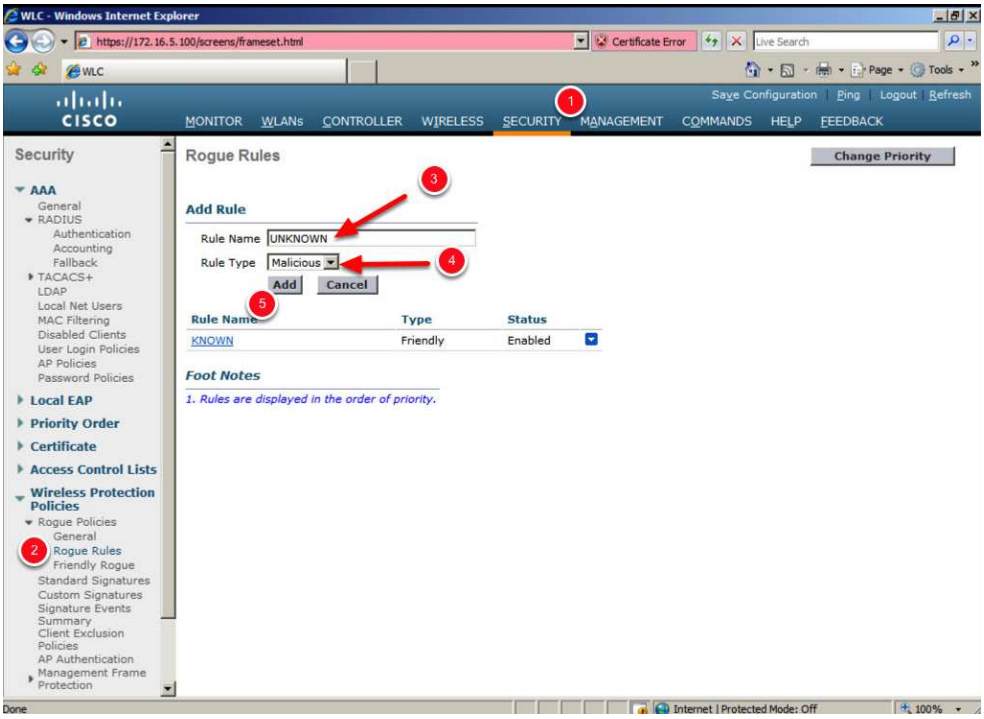


Edit the friendly rogue classification rule to match task requirements; note it is of type match-all:

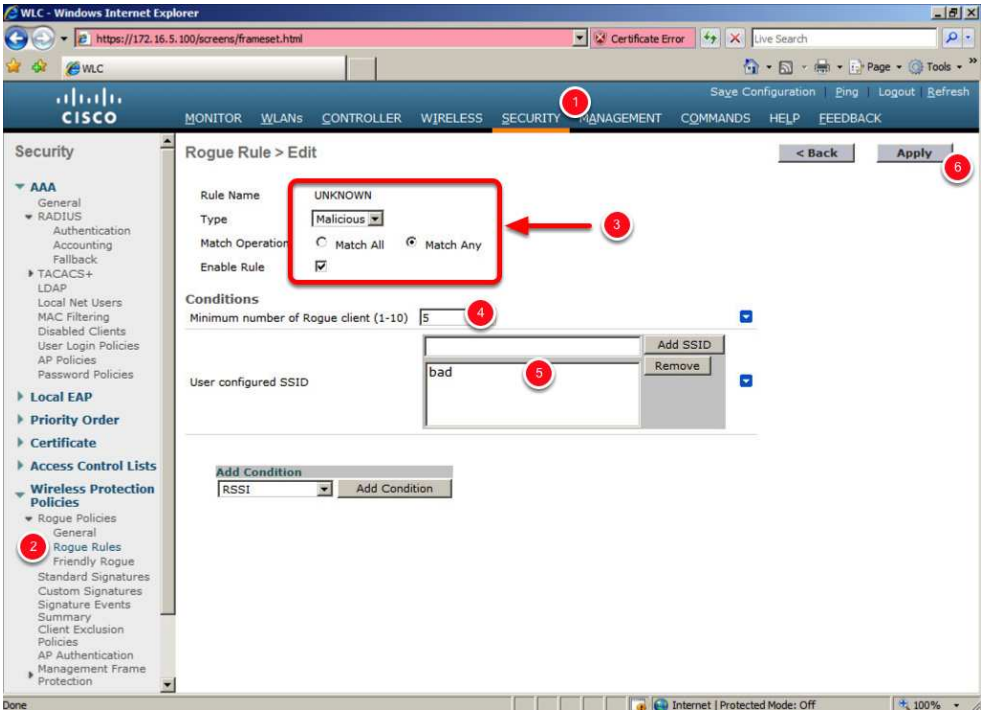


Add a malicious rogue classification rule:

https://t.me/learningnets



Edit the malicious rogue classification rule to match task requirements; note it is of type match-any:



Verification

Verify the configured rogue policies:

```
(Cisco Controller) >show rogue rule summary

Priority Rule Name                  State     Type          Match Hit Count
```

https://t.me/learningnets

```
-------- ---------------------- -------- ------------ ----- ---------
1        KNOWN                  Enabled  Friendly     All   0
2        UNKNOWN                Enabled  Malicious    Any   0
!
!(Cisco Controller) >show rogue rule detailed KNOWN

Priority..................................... 1
Rule Name.................................... KNOWN
State........................................ Enabled
Type......................................... Friendly
Match Operation.............................. All
Hit Count.................................... 0
Total Conditions............................. 2
Condition 1 type..................................... No-encryption
    value.................................... Enabled
Condition 2 type..................................... Ssid
    SSID Count............................... 1 SSID 1..................................... good
!
!(Cisco Controller) >show rogue rule detailed UNKNOWN

Priority..................................... 2
Rule Name.................................... UNKNOWN
State........................................ Enabled
Type......................................... Malicious
Match Operation.............................. Any
Hit Count.................................... 0
Total Conditions............................. 2
Condition 1 type..................................... Client-count
value.................................... 5
Condition 2 type..................................... Ssid
    SSID Count............................... 1 SSID 1..................................... bad
```

Verify that the configured friendly access-point has not been seen yet, but other access-points are automatically classified as unclassified (also auto-containment has not been enabled):

```
(Cisco Controller) >show rogue ap summary

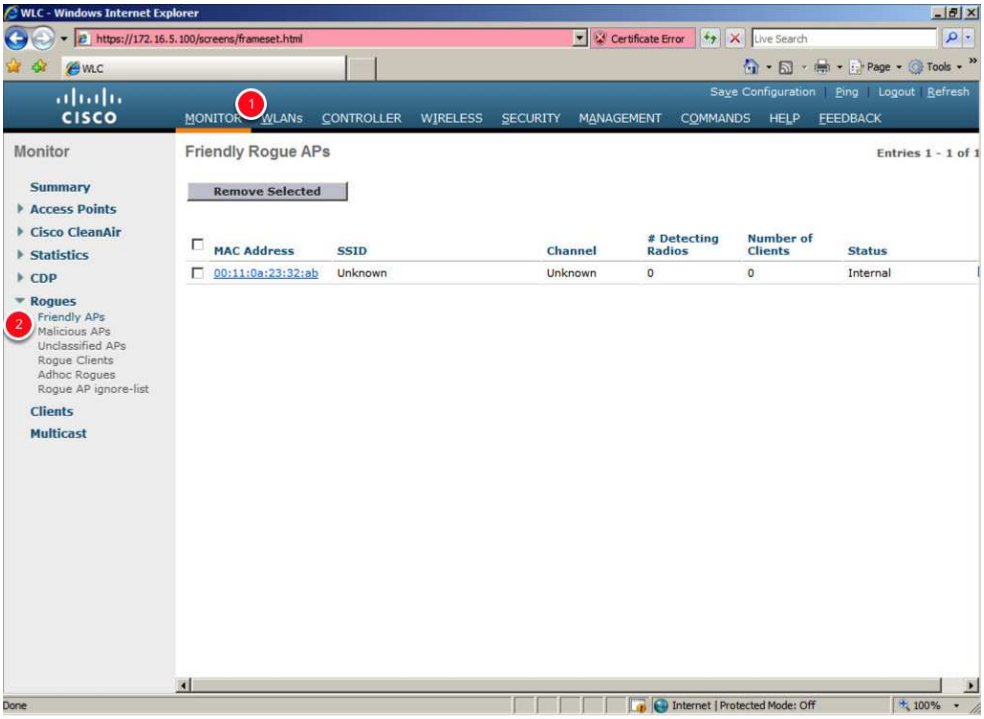
Rogue on wire Auto-Contain..................... Disabled
Rogue using our SSID Auto-Contain.............. Disabled
Valid client on rogue AP Auto-Contain.......... Disabled
Rogue AP timeout............................... 1200
Rogue Detection Report Interval................ 10
Rogue Detection Min Rssi....................... -128
```

https://t.me/learningnets

```
Rogue Detection Transient Interval.............. 0

MAC Address        Classification     # APs # Clients Last Heard
-----------------  ------------------ ----- --------- ----------------------
00:11:0a:23:32:ab  Friendly           0     0         Not Heard

58:6d:8f:9f:2e:8f  Unclassified       1     0         Sun Dec  8 14:16:05 2013
68:bd:ab:48:34:70  Unclassified       1     0         Sun Dec  8 14:10:07 2013
e8:04:62:f7:37:10  Unclassified       1     0         Sun Dec  8 14:10:07 2013
```



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

LAP Authorization

Task

- Configure WLC to authorize all LAPs against the local database.
  - Assume that all LAPs use MIC certificates.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command **`configure replace nvram:startup-config`** instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command **`clear lwapp private-config`** and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands -->Download File**).

Overview

When the LAP registers/joins the WLC, the LAP and WLC mutually authenticate using X.509 certificates; based on these certificates, the DTLS control-plane and, optionally, data-plane tunnels are established. Both LAPs and WLCs have X.509 certificates burned into protected flash from factory; installed certificates on LAPs are called manufacturing installed certificates (MIC). If, however, the LAP was converted from autonomous mode to LWAPP/CAPWAP, a self-signed certificate (SSC) is generated in the conversion process. Another option is to use your own internal PKI and deploy certificates to both WLC and LAP, in which case this authentication method is using local signed certificates (LSC).

https://t.me/learningnets

As an addition to this automatic trust established between WLC and LAP, you can manually authorize certain LAPs to join the WLC by using their Ethernet MAC addresses. This can be performed against a local MAC database (kept on the WLC) or against a remote database through RADIUS (kept on ACS or ISE); an exception to this is when using SSC certificates, in which case only local MAC database can be used. When a RADIUS server is used, the WLC will perform the LAP authorization against the ACS/ISE server in the form of username/password, where both these will have an identical value with the LAP's MAC address. Note that although this is called LAP authorization, from the RADIUS server perspective it is just a regular authentication process; no RADIUS authorization attributes need be returned in the Access-Accept message (in case of successful authentication) or Access-Reject message (in case of failed authentication).

Configuration

Identify the Ethernet MAC address of the LAP:

```
(Cisco Controller) >show ap summary

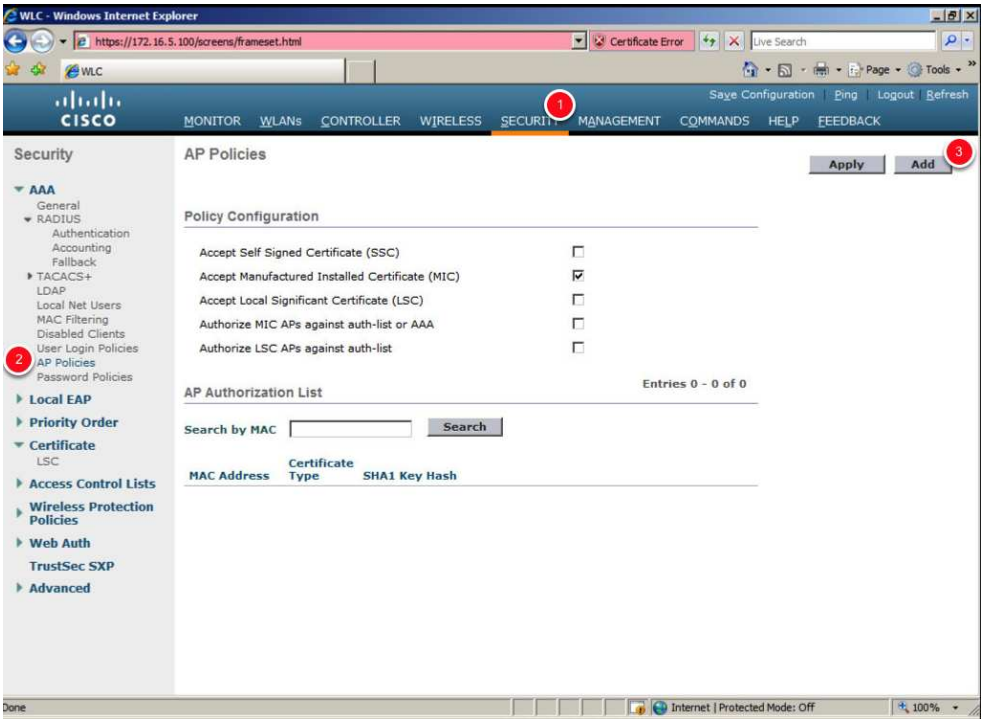

Number of APs.................................... 1

Global AP User Name.............................. Not Configured
Global AP Dot1x User Name........................ Not Configured

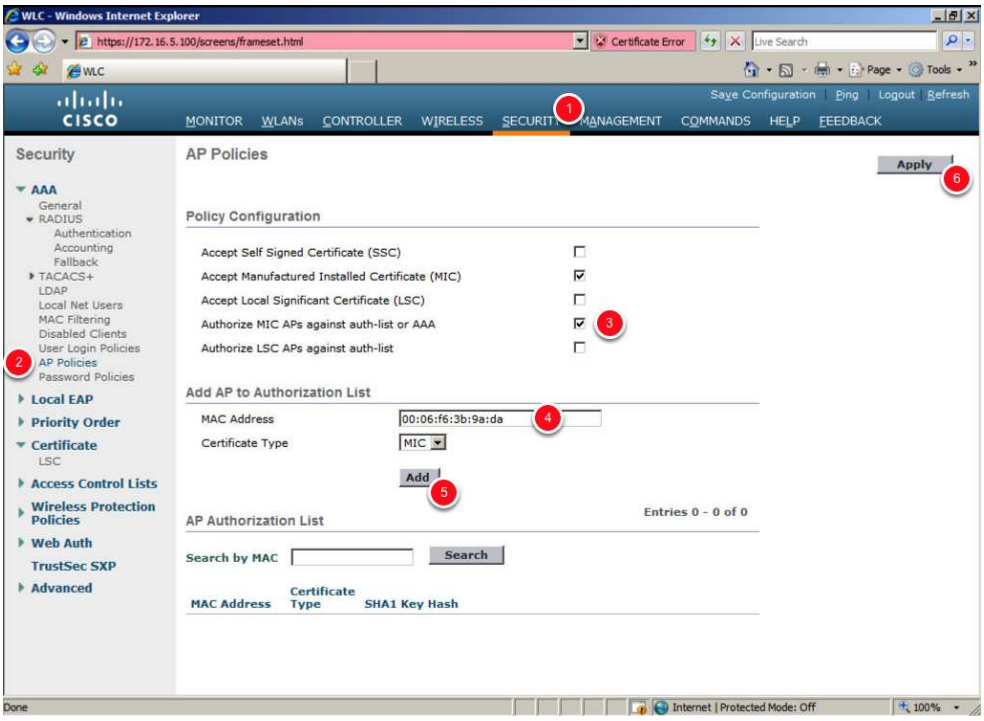
AP Name             Slots  AP Model              Ethernet MAC       Location          Port  Country  Priority
------------------  -----  --------------------  -----------------  ----------------  ----  -------  --------
AP0006.f63b.9ada     2     AIR-LAP1262N-A-K9     00:06:f6:3b:9a:da  default location  1        US       1
```

Navigate to the WLC GUI where this function is configured:

https://t.me/learningnets



Enable LAP authorization (with MIC) and add your rack's LAP Ethernet MAC address in the local database:



Verification

Verify the local MAC database and that the function is enabled:

```
(Cisco Controller) >show auth-list

Authorize MIC APs against AAA ..................... enabled
```

https://t.me/learningnets

```
Authorize LSC APs against Auth-List ..... disabled
APs Allowed to Join AP with Manufacturing Installed Certificate.... yes
  AP with Self-Signed Certificate..... no
  AP with Locally Significant Certificate..... no

Mac Addr                  Cert Type    Key Hash
-----------------------   ----------   ----------------------------------------
00:06:f6:3b:9a:da         MIC
```

Although the LAP is already joined to the WLC, the authorization is not enforced; this is done only in the registration process. To ensure that your configuration is correct, reload the LAP and see if it successfully joins/registers to the WLC:

```
(Cisco Controller) >show ap join stats summary all

Number of APs..................................... 1

Base Mac             AP EthernetMac       AP Name                IP Address         Status
f0:29:29:04:21:f0    00:06:f6:3b:9a:da    AP0006.f63b.9ada       136.1.47.100       Joined
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

LAP 802.1x Authentication

Task

- Configure SW5 to authenticate the LAP via 802.1x using the ACS server.
- All LAPs should use a username of **lapuser** for 802.1x authentication, but you rack's LAP should use a username of **LAP**.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands -->Download File**).

Overview

A LAP has a built-in 802.1x supplicant, so it can be authenticated by the switch, but it supports only EAP-FAST with anonymous PAC provisioning. Depending on the LAP functional mode (local or flexconnect), the switchport host-mode must be configured:

- In local mode, all user traffic is tunneled through the data-plane tunnel between the LAP and WLC (centrally switched), so the switch sees a single MAC address, the one belonging to the LAP; in this case, the switchport is in access mode and must be configured in single-host mode for 802.1x.

https://t.me/learningnets

- In flexconnect mode, user traffic can be tunneled between the WLC and LAP (centrally switched) or sent directly to the local switch (locally switched), and this is configurable at the SSID level; in this case, the switchport is in trunk mode and must be configured in multi-host mode for 802.1x so that only the LAP is authenticated.

EAP-FAST credentials used by LAP can be:

- Locally configured at the LAP level.
- Globally configured on WLC, affecting all joined LAPs.
- Per LAP configured on WLC, affecting only the specific LAP.

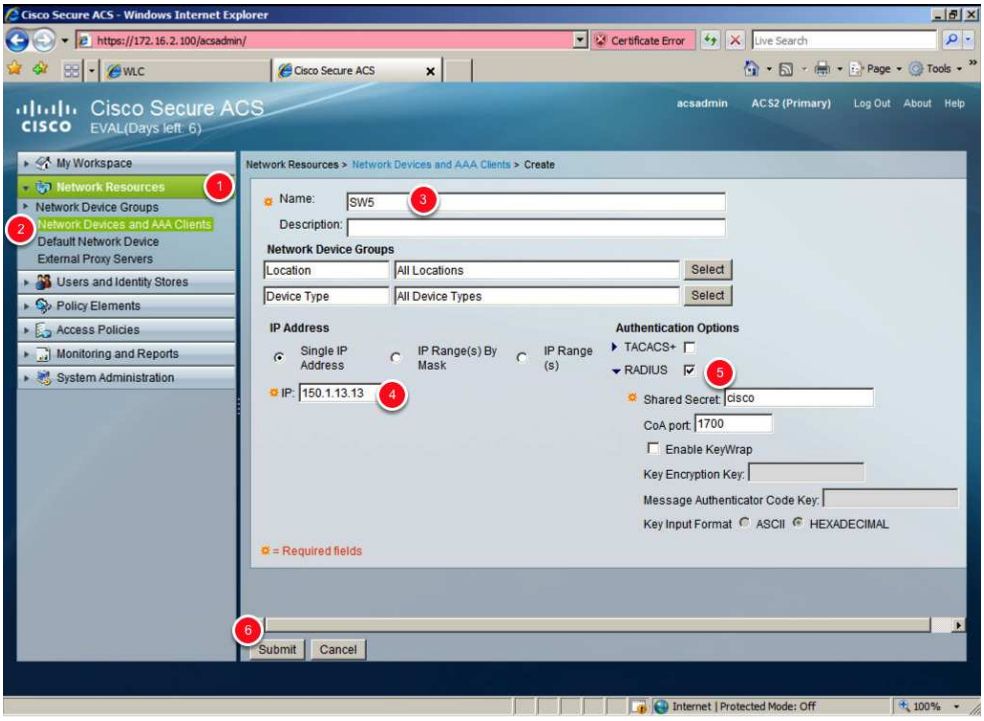
Configuration

SW5:

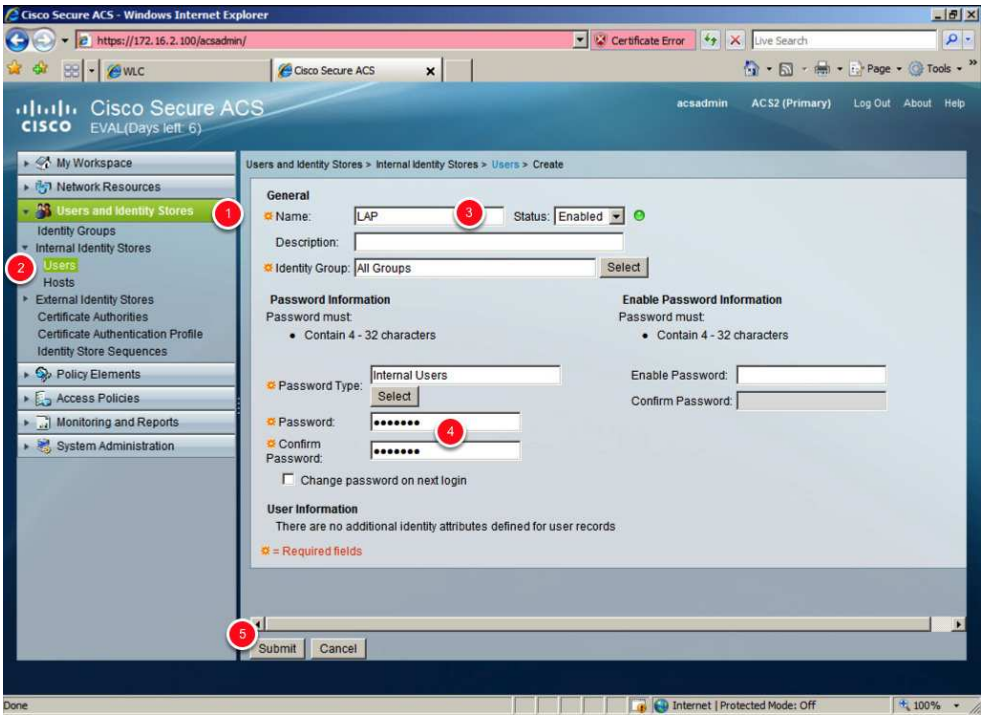
```
aaa new-model
aaa authentication dot1x default group radius
!
dot1x system-auth-control
!
radius-server host 172.16.2.100 key cisco
ip radius source-interface Loopback0
!
interface FastEthernet1/0/11
 authentication host-mode single-host
 authentication port-control auto
 dot1x pae authenticator
```

Configure SW5 as NAD into ACS:

https://t.me/learningnets

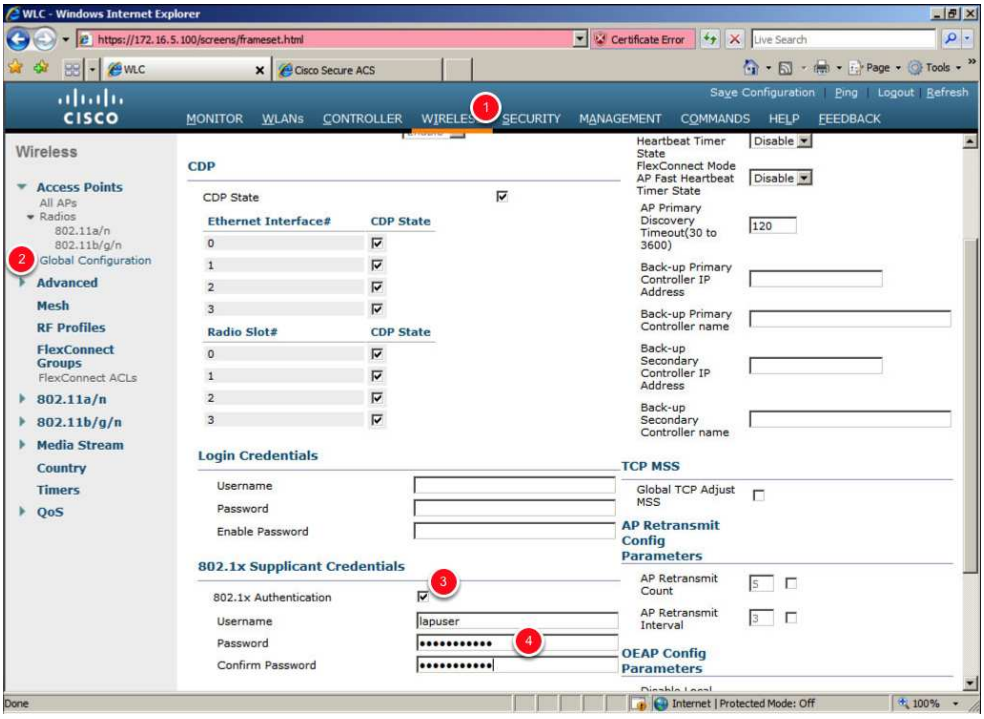


Configure the username used by LAP through EAP-FAST to authenticate against ACS:

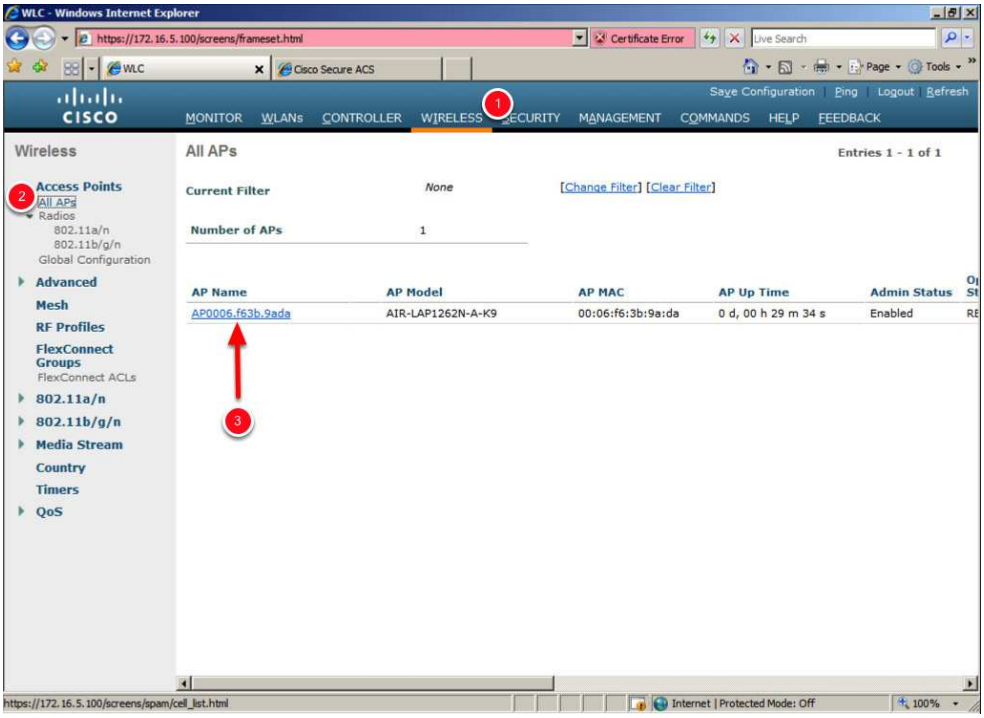


Configure WLC with 802.1x credentials for all joined LAPs (globally):

https://t.me/learningnets

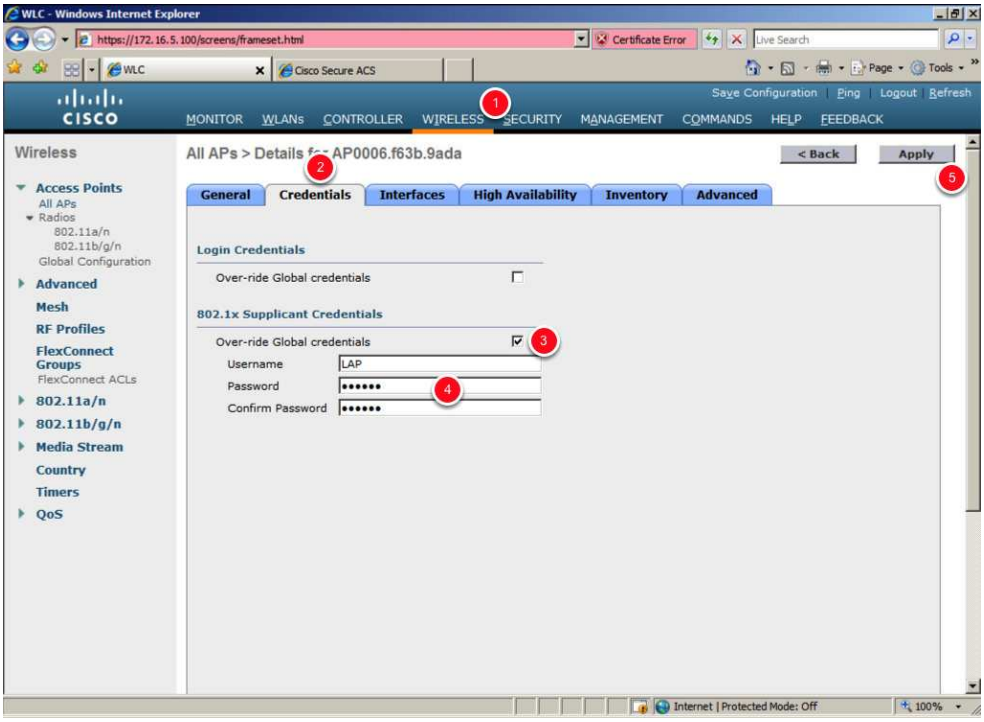


Navigate in the WLC GUI to configure your LAP, and click on its name:



Configure the LAP within your rack with specific credentials to be used for 802.1x authentication, overriding the global credentials:

https://t.me/learningnets



The default RADIUS access policy allows EAP-FAST protocol (along with anonymous in-band PAC provisioning) and authenticates against the internal user store. Additionally, the default RADIUS authorization access policy is set to **Permit Access**, which means a RADIUS Access-Accept message will be sent upon successful authentication; because we are not required to perform any specific authorizations (like dACL or VLAN), using these default values will work.

Verification

Verify that LAP was successfully authenticated by the RADIUS server:

```
SW5#show authentication sessions interface fastEthernet1/0/11
            Interface:  FastEthernet1/0/11
          MAC Address:  0006.f63b.9ada
           IP Address:  Unknown User-Name:  LAP
Status:  Authz Success
               Domain:  DATA
      Security Policy:  Should Secure
      Security Status:  Unsecure Oper host mode:  single-host
     Oper control dir:  both Authorized By:  Authentication Server
           Vlan Group:  N/A
      Session timeout:  N/A
         Idle timeout:  N/A
    Common Session ID:  88012D0D000000014F9E39EF
      Acct Session ID:  0x00000002
               Handle:  0x37000001
```

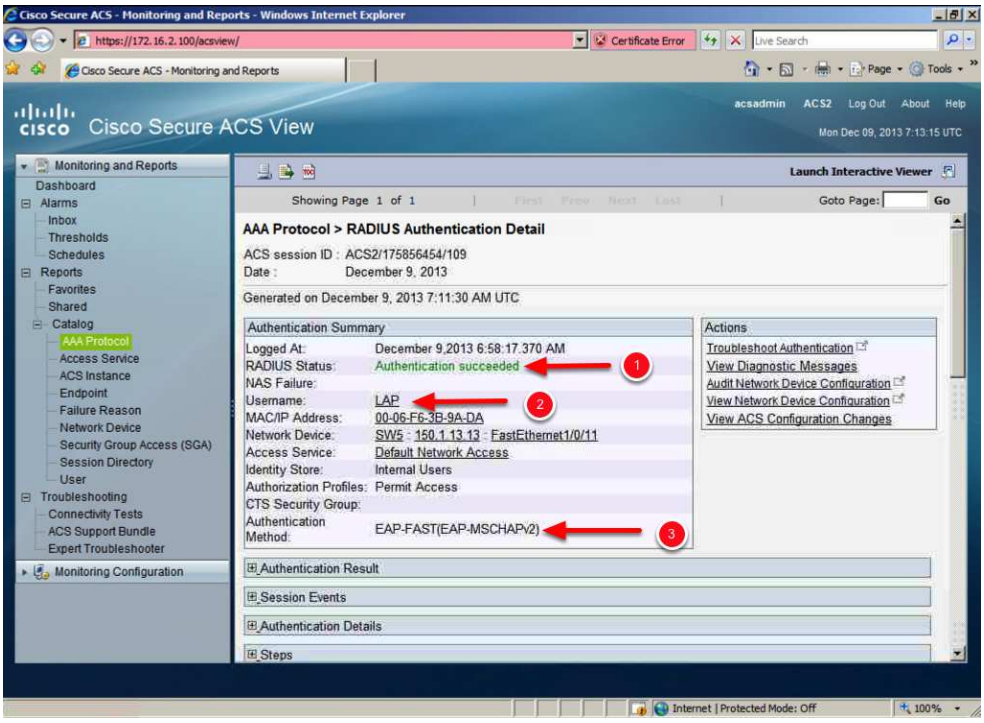
https://t.me/learningnets

```
Runnable methods list:
       Method   State dot1x     Authc Success


!
!AP0006.f63b.9ada#show dot1x  all details
Sysauthcontrol              Disabled
Dot1x Protocol Version            2

Dot1x Info for GigabitEthernet0
-----------------------------------
PAE                       = SUPPLICANT
StartPeriod               = 30
AuthPeriod                = 30
HeldPeriod                = 60
MaxStart                  = 3
Credentials profile       = lwapp_credentials
EAP profile               = lwapp_eap_profile
Dot1x Supplicant Client List
------------------------------
Authenticator             = 001a.e2b6.758d
        Supp SM State     = AUTHENTICATED
        Supp Bend SM Stat = IDLE Port Status                = AUTHORIZED
```

Verify the RADIUS authentication logs:



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

LAP Local Mode

Task

- Configure WLC to broadcast a SSID named **CCIE-LABXX**, where **X** equals the rack number you've been assigned.
- Configure authentication using preshared key.
- Protect management frames but accept non-capable clients.
- Configure WLC for additional filtering so that only WiFi TEST PC can connect to this SSID.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

Management Frame Protection (MFP) is used to increase the security of the 802.11 management frames that are otherwise unencrypted, thus protecting against DoS attacks. There are two types of MFP supported by the WLC:

- **Infrastructure MFP**: Protects management frames between access-points by adding message integrity check information elements (MIC-IE) to these frames; this is

https://t.me/learningnets

enabled globally on WLC.
- **Client MFP**: Protects management frames between CCXv5 (Cisco compatible extensions) compatible clients and access-points; this is only supported for SSIDs using WPA2 with TKIP or AES for encryption. This is enabled at the WLAN level; the difference between **optional** and **required** is that with required, only CCXv5-compatible clients will successfully associate to the SSID.

There are two main differences between WPA and WPA2:

- WPA2 support for AES encryption is mandatory; in WPA it was optional.
- WPA2 uses a four-way handshake between client and access-point, compared with a six-way handshake on WPA.

TKIP and AES are cipher suites used to protect/encrypt frames over the shared radio medium. Whereas TKIP appeared as an enhancement to WEP protocol, AES is an industry standard widely deployed. There is also a Cisco proprietary version of TKIP, known as CKIP.

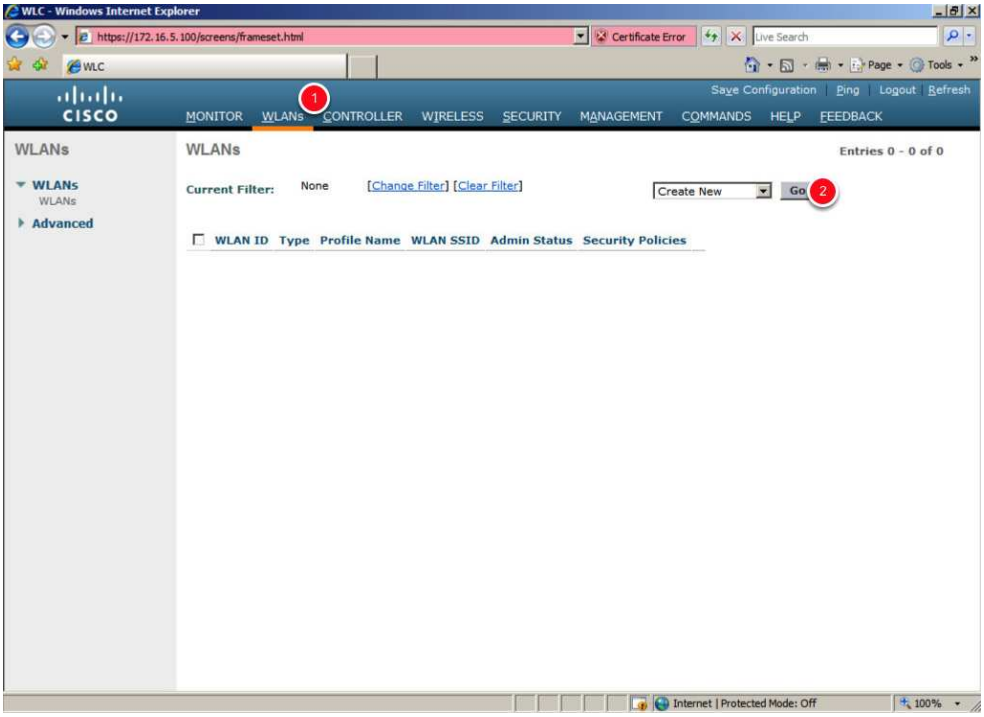
The following concepts must be clear enough in the process of a client joining a wireless infrastructure:

- **Authentication** is the first step and is optional. If mandatory, it can be one-way (client is authenticated) or two-way (mutual authentication between client and server); there is layer2 authentication and layer3 authentication.
- **Encryption** is optional; if layer2 authentication is enabled, encryption is automatically enabled and is used to protect data transfer over the shared medium.
- **Key distribution** is required if encryption is active; for the client and access-point to encrypt traffic, they both need a shared key that can be manually distributed (preshared key authentication), or dynamically distributed (through EAP methods); the dynamic method requires the existence of a backend RADIUS authentication server.
- **Key management** takes care of key refreshment and group key generation so that all clients associated to the same access-point can still talk to each other.
- **Key reuse** allows the client and access-point to reuse a previously negotiated key, which helps in roaming because restarting the key negotiation process is time consuming.

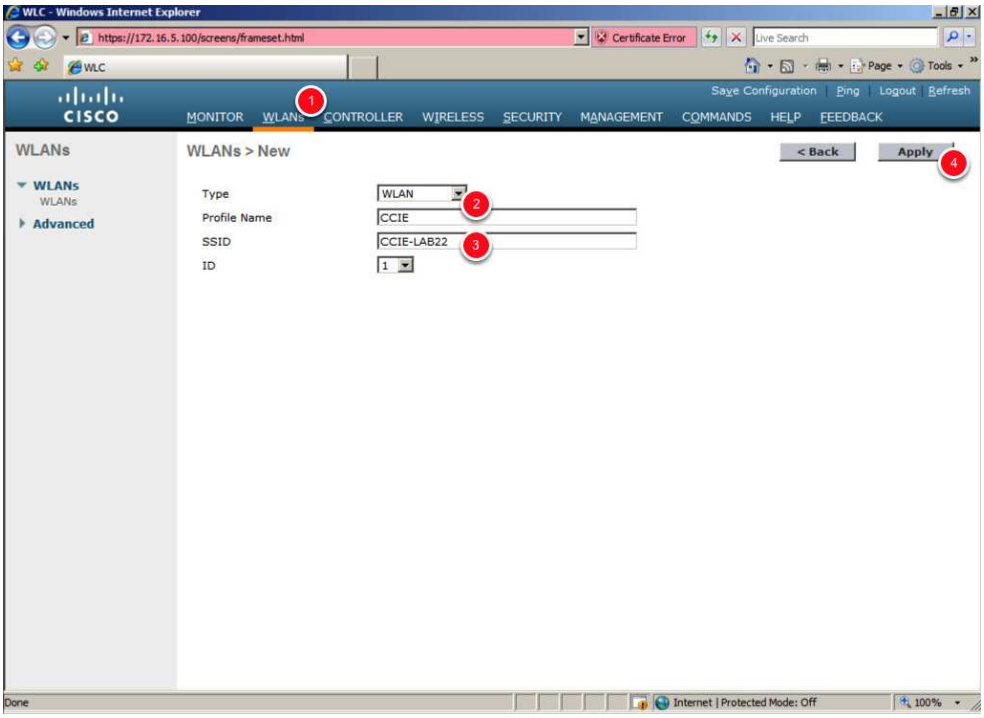
Configuration

Navigate through WLC GUI to create a new WLAN:

https://t.me/learningnets

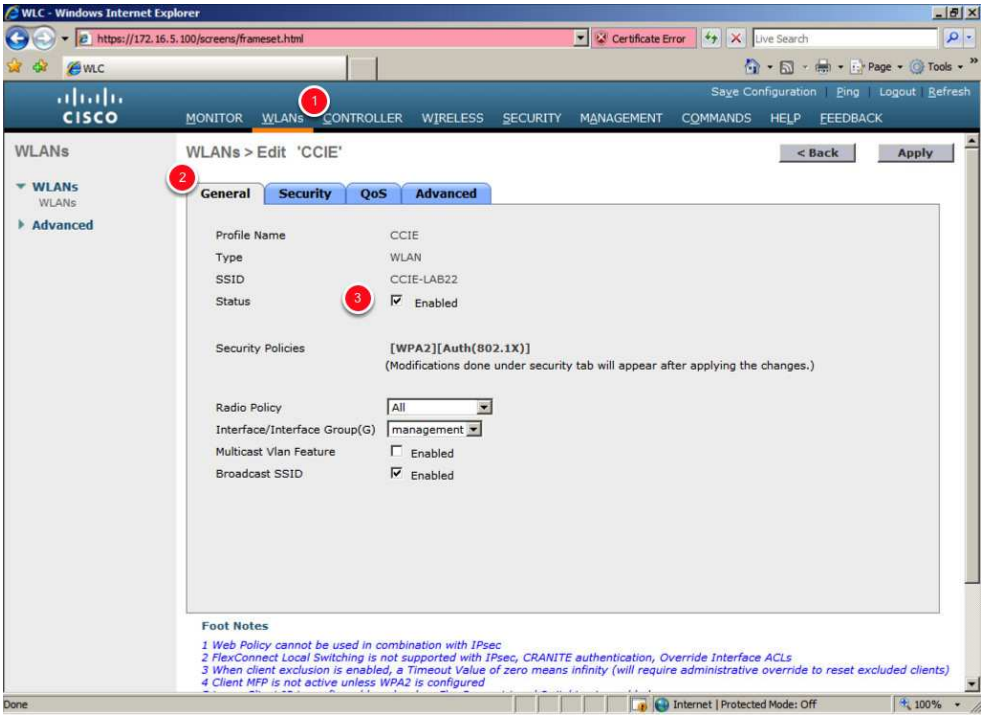


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

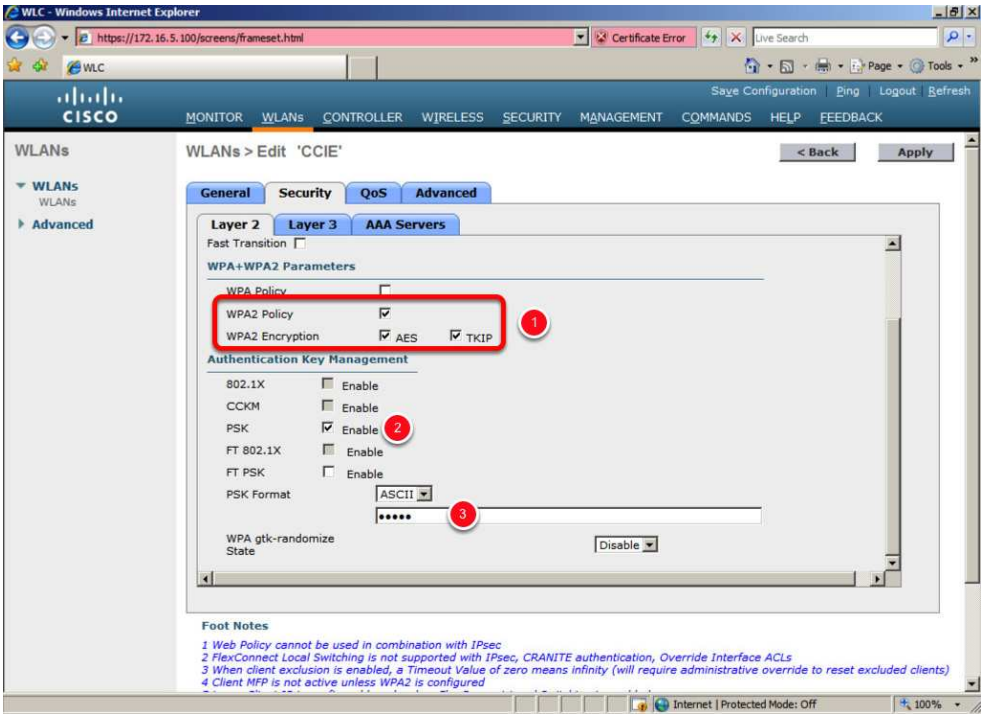


Enable the WLAN and leave WLC to broadcast the SSID:

https://t.me/learningnets

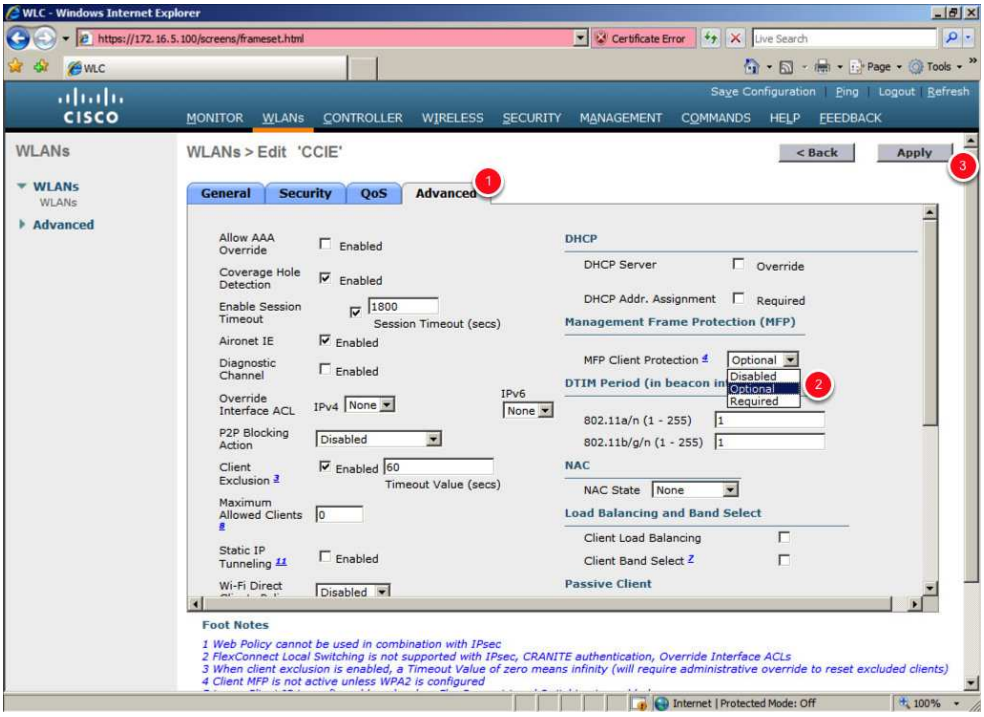


Use only WPA2 and specify the PSK; it must be between 8 and 63 characters:

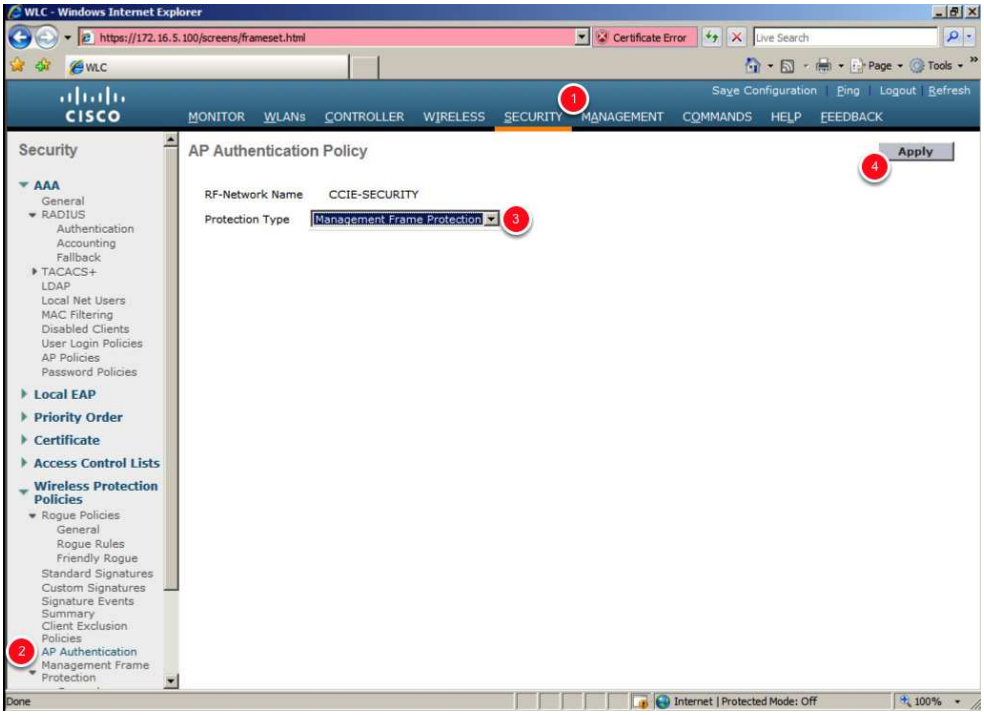


Leave client MFP to the default of optional (so that non-MFP-capable clients are accepted) and apply all configured settings for the WLAN:

https://t.me/learningnets

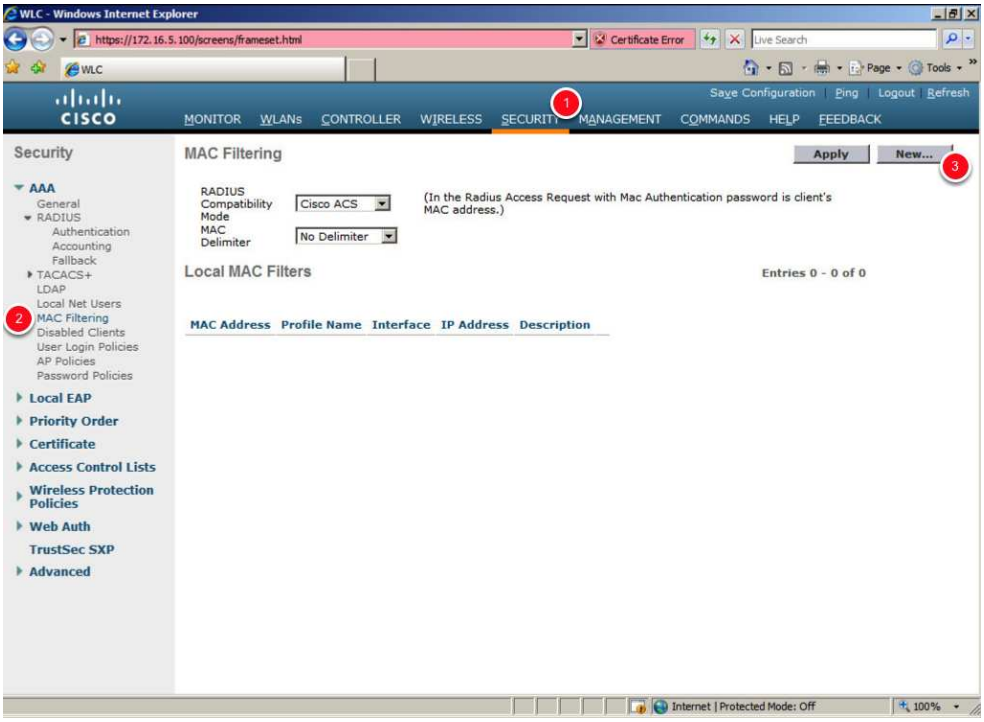


Globally enable infrastructure MFP:

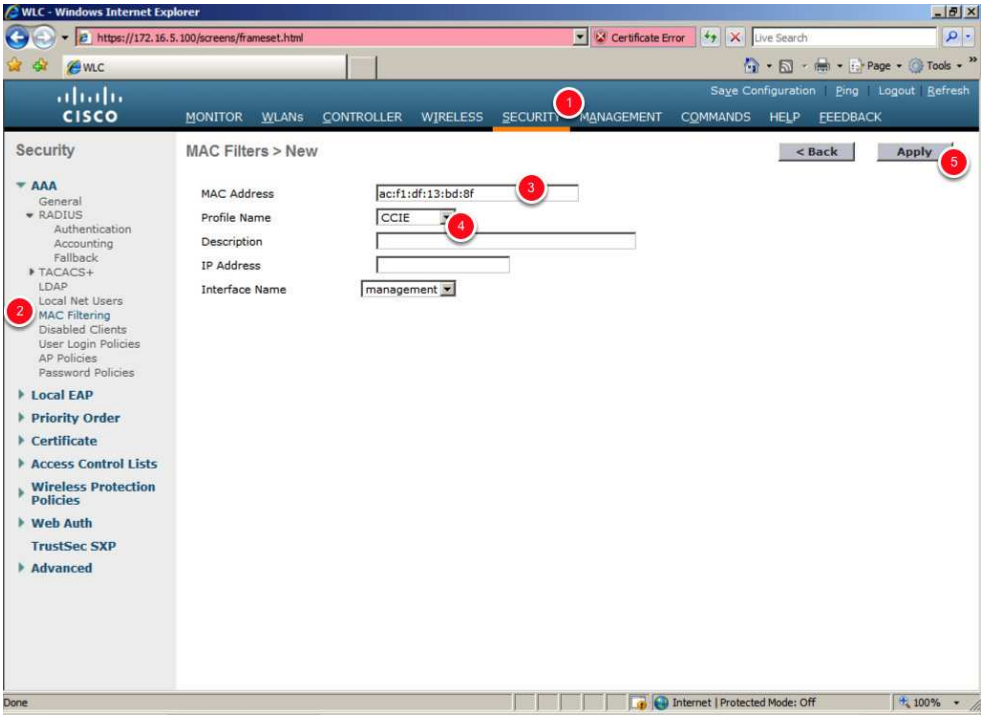


Navigate through WLC GUI to add your rack's WiFi Test-PC MAC address in the local database (based on the rack being assigned, this may differ):

https://t.me/learningnets



Add the MAC address:



Configure the WiFi client to connect; it should be successful, but because there are no DHCP servers configured for it, it will have no IP connectivity.

Verification

Verify that WLAN is enabled:

https://t.me/learningnets

```
(Cisco Controller) >show wlan summary

Number of WLANs.............................. 1

WLAN ID  WLAN Profile Name / SSID               Status    Interface Name
-------  -----------------------------------  --------  --------------------
1        CCIE / CCIE-LAB22                     Enabled   management
```

Verify that MFP is enabled:

```
(Cisco Controller) >show wps mfp summary
 Global Infrastructure MFP state.................. Enabled
Controller Time Source Valid..................... False

                                    WLAN       Client
WLAN ID  WLAN Name                  Status     Protection
-------  -------------------------  ---------  ----------
1        CCIE                       Enabled    Optional
```

Verify that MAC filtering is configured:

```
(Cisco Controller) >show macfilter summary

MAC Filter RADIUS Compatibility mode.............. Cisco ACS
MAC Filter Delimiter.............................. None

Local Mac Filter Table

MAC Address               WLAN Id          IP Addr           Description
-----------------------   --------------   ---------------   --------------------------------
ac:f1:df:13:bd:8f          1               unknown
```

Verify that the client is associated and authenticated:

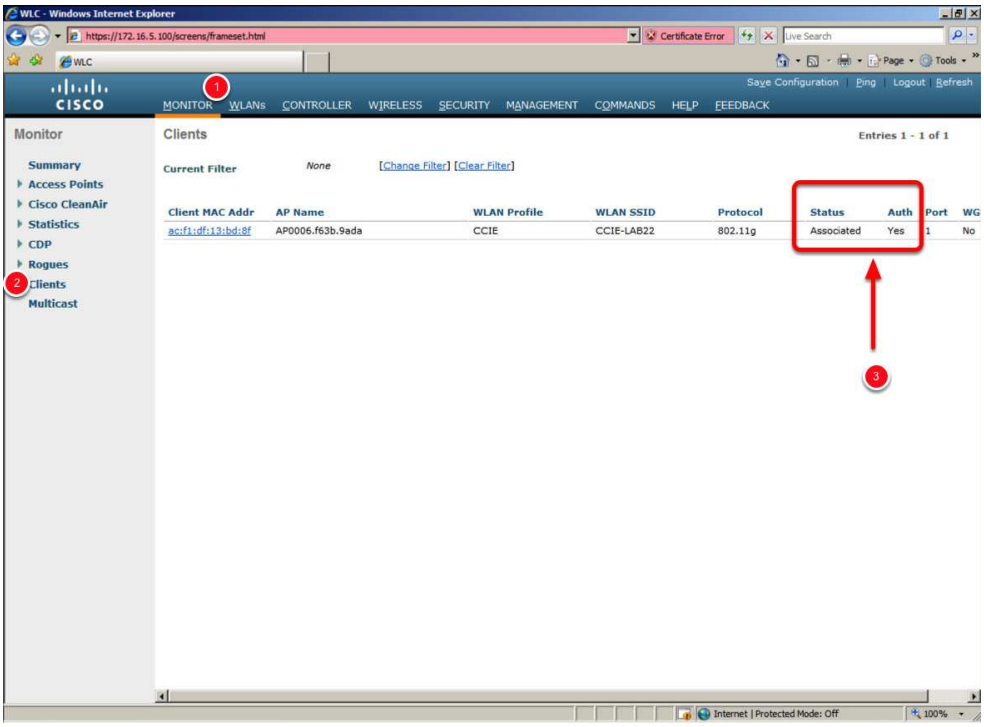
```
(Cisco Controller) >show client summary

Number of Clients............................. 1

MAC Address       AP Name           Status        WLAN/RLAN      Auth Protocol         Port Wired
----------------- ----------------- ------------- -------------- ---- ---------------- ---- -----
```

https://t.me/learningnets

```
ac:f1:df:13:bd:8f AP0006.f63b.9ada  Associated    1              Yes  802.11g          1    N/A
```



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

LAP FlexConnect Mode

Task

- Configure WLC with a hidden SSID named **CCIE-XX**, where **X** equals the rack number you've been assigned.
- Configure WPA authentication using a preshared key; WLC should exclude for 300 seconds only for clients failing layer2 authentication.
- User traffic for this SSID should not be tunneled to the controller and assigned to VLAN 147 (136.1.147.0/24).
- Configure SW5 to act as default gateway and DHCP server for WiFi clients.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

By default, access-points will regularly broadcast their active SSIDs in beacons so that clients can easily identity which networks are available. Optionally, this can be disabled, in which case you would have to know the network name and manually add it on the client side. To protect the control-plane from DoS attacks, WLC is by default configured with a set of thresholds to disallow WiFi access to clients in

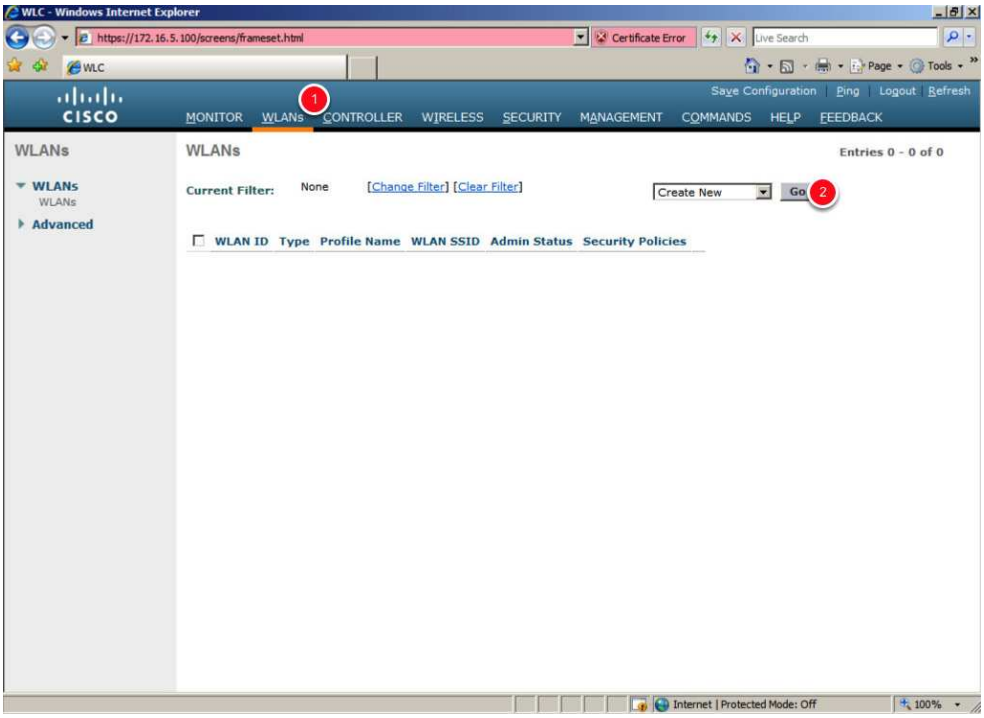
https://t.me/learningnets

certain cases:

- **Excessive 802.11 Association Failures**: Clients are excluded after five consecutive association failures.
- **Excessive 802.11 Authentication Failures**: Clients are excluded after five consecutive authentication failures.
- **Excessive 802.1X Authentication Failures**: Clients are excluded after three consecutive authentication failures.
- **IP Theft or IP Reuse**: Clients are excluded if the IP address is already assigned to another device.
- **Excessive Web Authentication Failures**: Clients are excluded after three consecutive authentication failures.

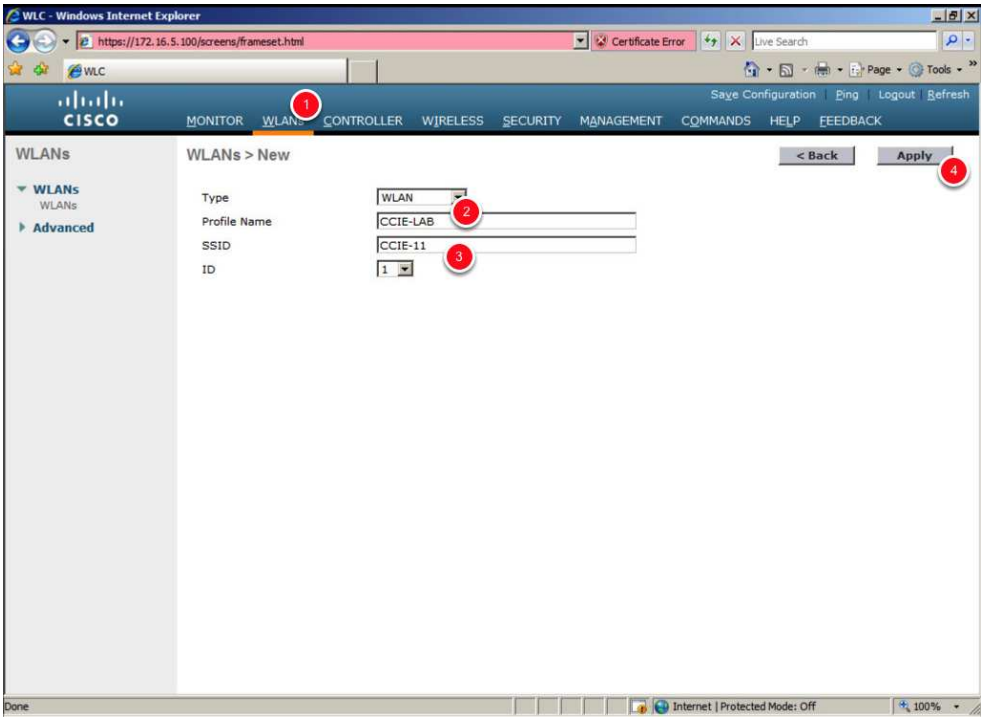
Configuration

Navigate through WLC GUI to create a new WLAN:

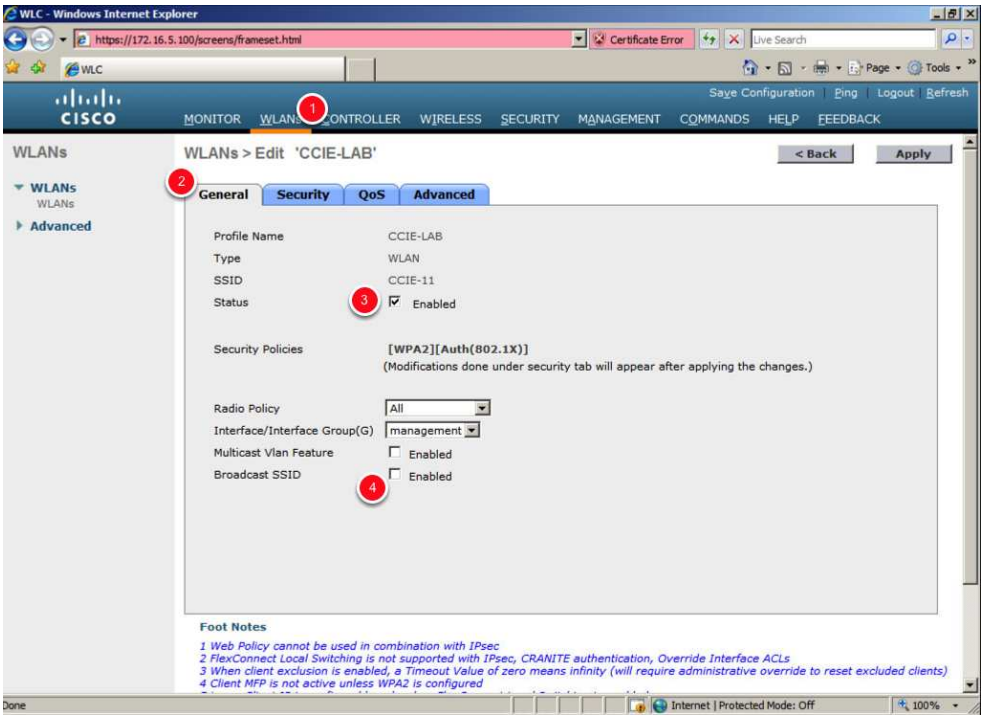


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

https://t.me/learningnets

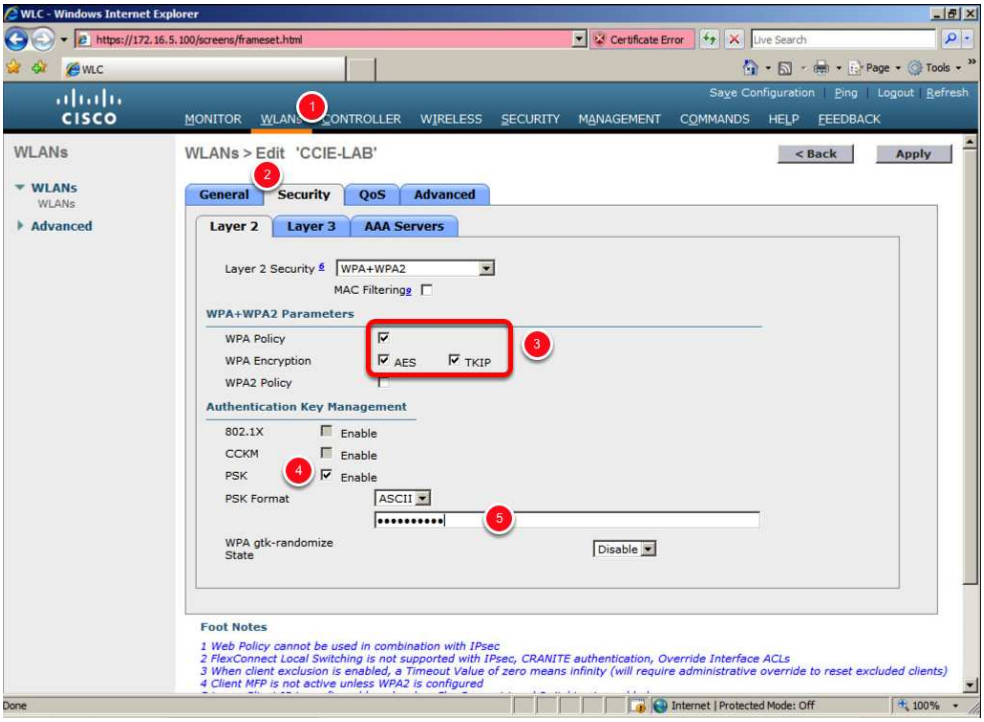


Enable the WLAN and disable the SSID broadcasting, thus hiding the network:

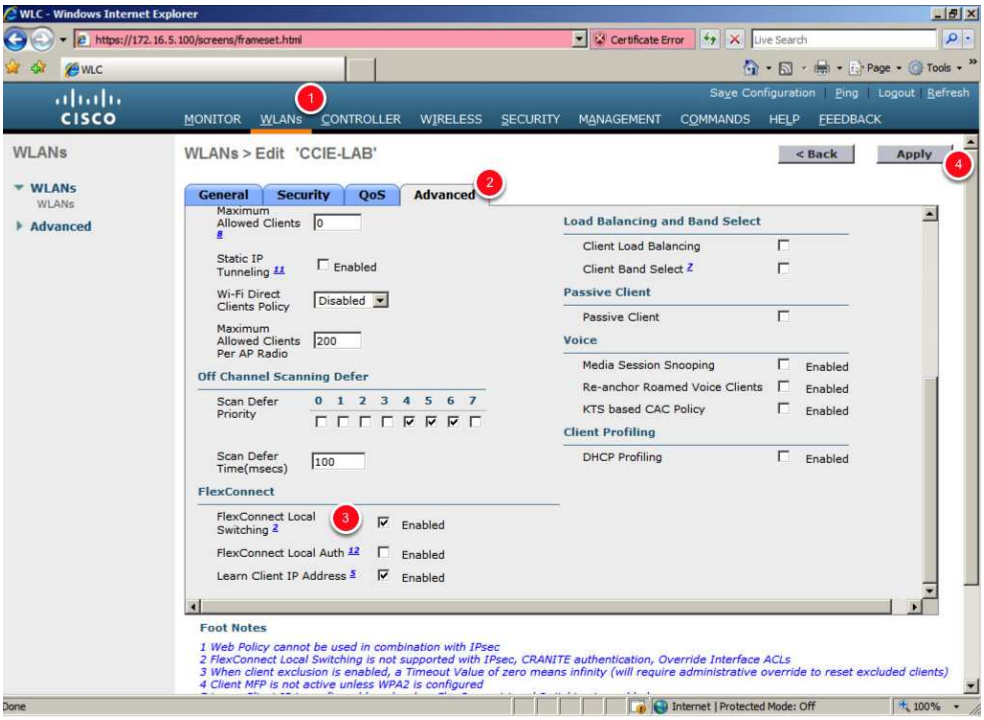


Configure the SSID for WPA with PSK authentication:

https://t.me/learningnets

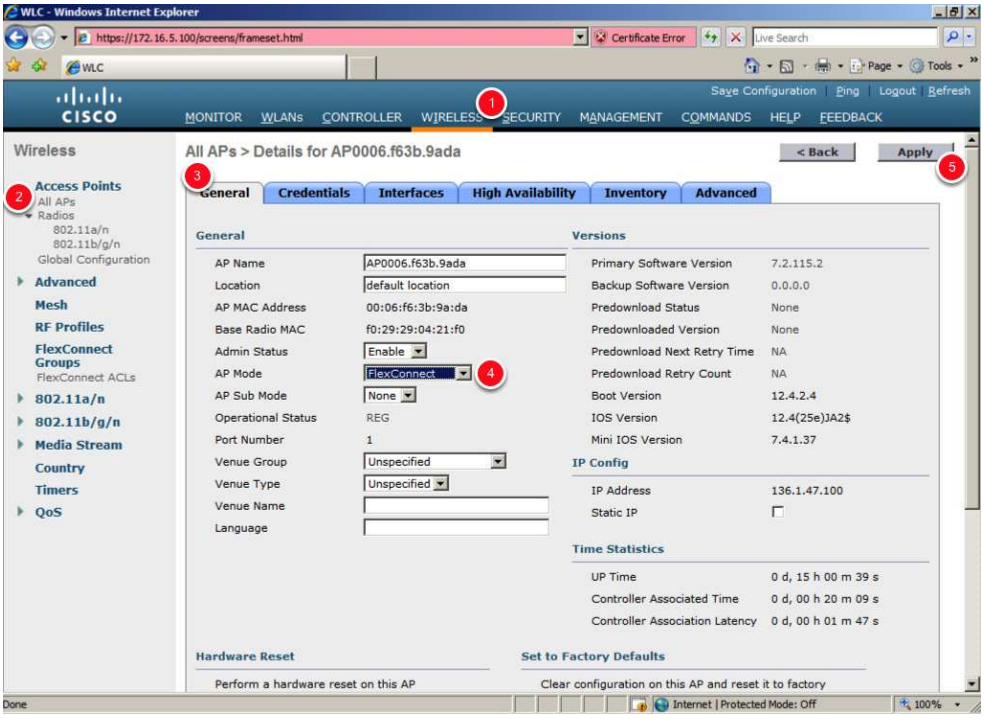


Enable the SSID for flexconnect switching:

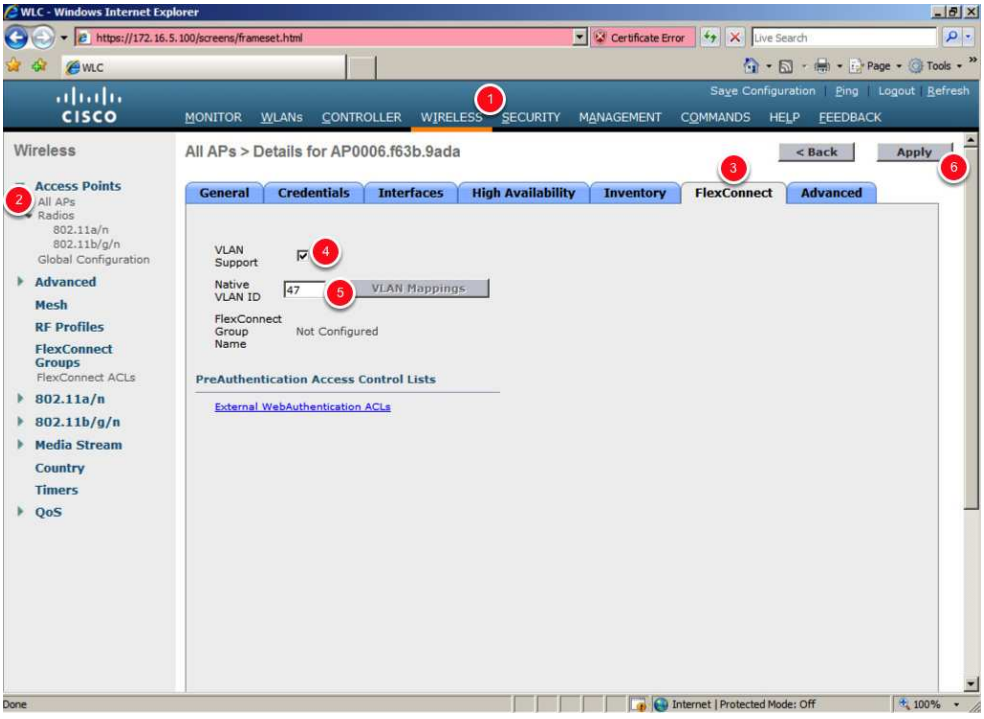


Configure the LAP to function in flexconnect mode; it will reload later:

https://t.me/learningnets

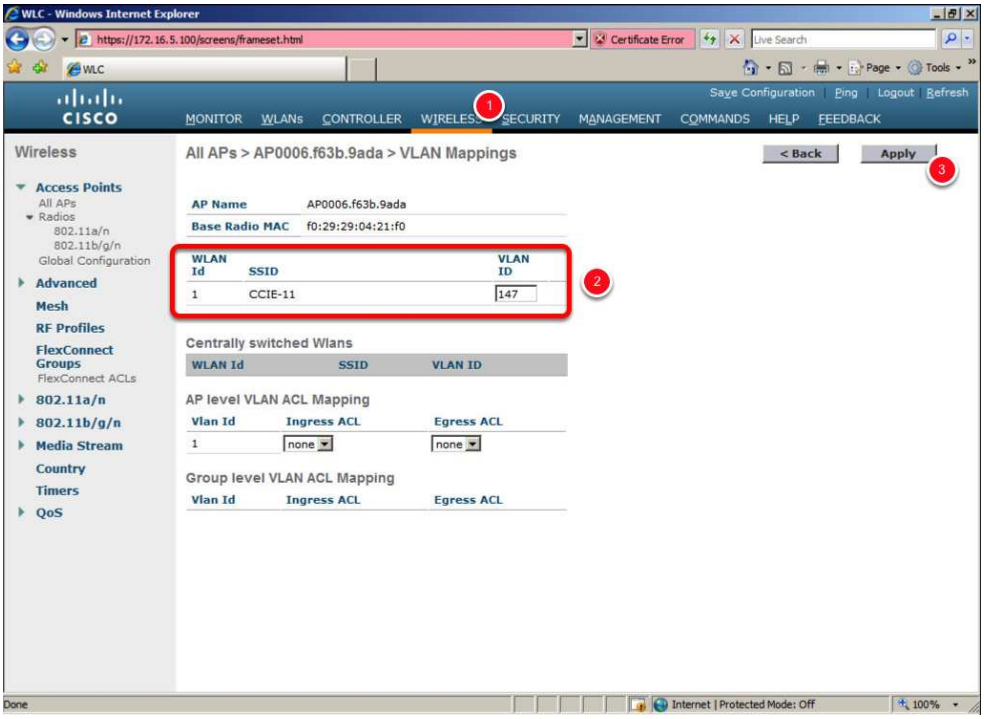


LAP will boot up in flexconnect mode, and a new tab will be visible for LAP configuration. Enable VLAN support and specify the native VLAN to be the management VLAN of the LAP; this is because LAP does not support VLAN tagging for the management VLAN.

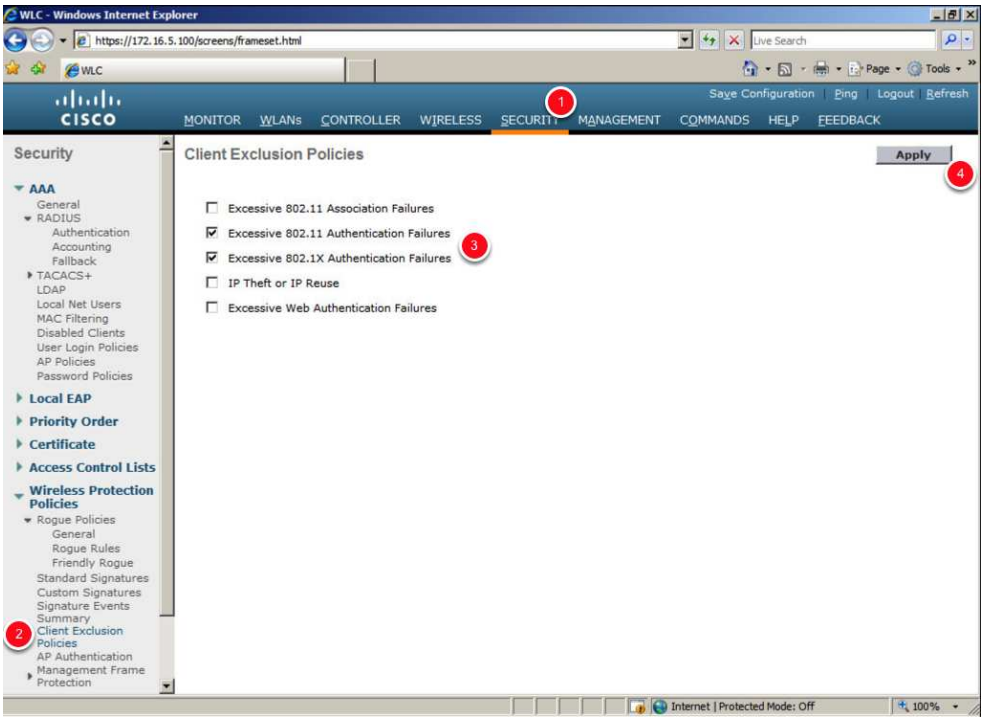


Navigate again to LAP settings, and configure the VLAN mapping for the defined SSID to be VLAN 147:

https://t.me/learningnets

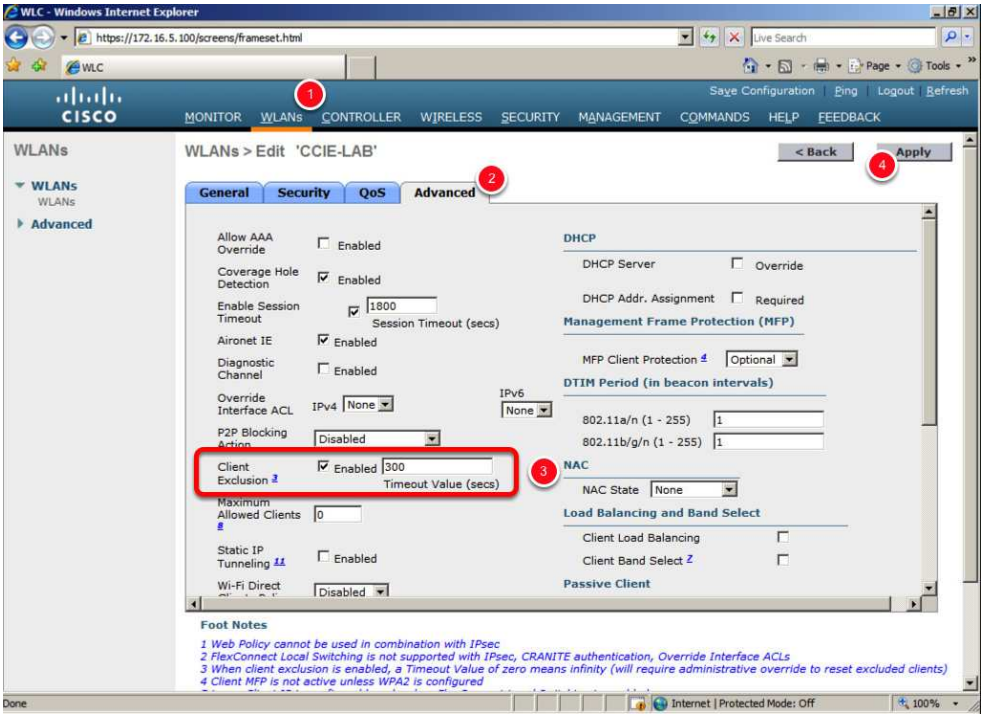


Configure client exclusion policies only for authentication failures; all options are enabled by default:



At the WLAN level, change the default interval for clients being excluded from 60 seconds to 300 seconds:

https://t.me/learningnets



SW5:

```
vlan 147
 name WIFI-PC
!
interface Vlan 147
 ip address 136.1.147.13 255.255.255.0
!
ip dhcp excluded-address 136.1.147.1 136.1.147.1
ip dhcp pool VLAN147
 network 136.1.147.0 255.255.255.0
 default-router 136.1.147.13
!
interface FastEthernet1/0/11
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 47
 switchport trunk allowed vlan 47,147
 switchport mode trunk
```

Associate the WiFi TEST-PC to the configured SSID. Because it is hidden, you must manually configure it.

Verification

Verify that WLAN is enabled:

https://t.me/learningnets

```
(Cisco Controller) >show wlan summary


Number of WLANs.................................. 1

WLAN ID  WLAN Profile Name / SSID                Status    Interface Name
-------  -----------------------------------  --------  --------------------
1        CCIE-LAB / CCIE-11                     Enabled   management
```

Verify advanced WLAN settings:

```
(Cisco Controller) >show wlan 1


WLAN Identifier.................................. 1
Profile Name..................................... CCIE-LAB
Network Name (SSID).............................. CCIE-11
Status........................................... Enabled
MAC Filtering.................................... Disabled
Broadcast SSID................................... Disabled
AAA Policy Override.............................. Disabled
Network Admission Control
  Client Profiling Status ....................... Disabled
  Radius-NAC State............................... Disabled
  SNMP-NAC State................................. Disabled
  Quarantine VLAN................................ 0
Maximum number of Associated Clients.............. 0
Maximum number of Clients per AP Radio........... 200
Number of Active Clients......................... 1
Exclusionlist Timeout............................ 300 seconds
Session Timeout.................................. 1800 seconds
CHD per WLAN..................................... Enabled
Webauth DHCP exclusion........................... Disabled
Interface........................................ management
--More-- or (q)uit
Wi-Fi Direct policy configured................. Disabled
EAP-Passthrough.............................. Disabled
CKIP ........................................ Disabled
Web Based Authentication..................... Disabled
Web-Passthrough.............................. Disabled
Conditional Web Redirect..................... Disabled
Splash-Page Web Redirect..................... Disabled
Auto Anchor.................................. Disabled
FlexConnect Local Switching.................. Enabled
```

https://t.me/learningnets

```
FlexConnect Local Authentication.............. Disabled
FlexConnect Learn IP Address.................. Enabled
Client MFP.................................... Optional but inactive (WPA2 not configured)
Tkip MIC Countermeasure Hold-down Timer....... 60
```

Verify that global client exclusion policies are enabled:

```
(Cisco Controller) >show wps summary

Auto-Immune
  Auto-Immune.................................. Disabled

Client Exclusion Policy
  Excessive 802.11-association failures.......... Disabled
Excessive 802.11-authentication failures....... Enabled
 Excessive 802.1x-authentication................ Enabled

  IP-theft..................................... Disabled
  Excessive Web authentication failure........... Disabled

Signature Policy
  Signature Processing......................... Enabled
```

Verify that the client is associated and authenticated:

```
(Cisco Controller) >show client summary

Number of Clients............................... 1

MAC Address       AP Name           Status        WLAN/RLAN      Auth Protocol         Port Wired
----------------- ----------------- ------------- -------------- ---- ---------------- ---- -----
 ac:f1:df:13:bd:8f AP0006.f63b.9ada  Associated    1              Yes  802.11n(2.4 GHz) 1    N/A
```

Verify that the client has been assigned an IP address and is reachable:

```
SW5#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
136.1.47.100        0100.06f6.3b9a.da       Mar 18 1993 01:26 AM    Automatic
136.1.147.2         01ac.f1df.13bd.8f       Mar 18 1993 09:51 AM    Automatic
```

https://t.me/learningnets

```
!
!SW5#ping 136.1.147.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.1.147.2, timeout is 2 seconds:!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/7/17 ms
!
!(Cisco Controller) >show client detail ac:f1:df:13:bd:8f
Client MAC Address.............................. ac:f1:df:13:bd:8f
Client Username ................................ N/A
AP MAC Address.................................. f0:29:29:04:21:f0
AP Name......................................... AP0006.f63b.9ada
Client State.................................... Associated
Client NAC OOB State............................ Access
Wireless LAN Id................................. 1
802.11u......................................... Not Supported
BSSID........................................... f0:29:29:04:21:f0
Connected For .................................. 550 secs
Channel......................................... 1
IP Address...................................... 136.1.147.2

Association Id.................................. 1
Authentication Algorithm........................ Open System
Reason Code..................................... 1
Status Code..................................... 0
Session Timeout................................. 1800
Client CCX version.............................. 5
Client E2E version.............................. No E2E support
Diagnostics Capability.......................... Supported
S69 Capability.................................. Not Supported
QoS Level....................................... Silver
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

PEAP (EAP-MSCHAPv2) Local Mode

Task

- Configure WLC with a SSID named **CCIE-LABXX**, where **X** equals the rack number you've been assigned:
  - Use PEAP (EAP-MSCHAPv2) authentication via the ACS RADIUS server.
  - Allow only WPA2 clients and enable Cisco's proprietary key management system, additionally to 802.1x.
  - ACS should authorize users to have only ICMP access after they are authenticated.
  - Use WLC as the DHCP server for clients of this WLAN, and assign the IP address from the range 172.16.5.200-172.16.5.254.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

In contrast to all the identity management tasks from the previous sections, in which the authenticator was the switch, in this case the authenticator is the WLC. On the client side, you must simply enable 802.1x settings for the Wireless NIC, not for the

https://t.me/learningnets

Wired NIC; on the authentication server (ACS/ISE), configuration is mostly identical. Because the WLC does not support dACL, to apply a per-user ACL through RADIUS authorization you must make use of the Cisco Airespace-ACL-Name attribute, which will only signal the ACL name, so ACL itself must be configured on the WLC. However, for the WLC to apply any authorization received from the RADIUS server, you must enable the **Allow AAA Override** option at the WLAN level.

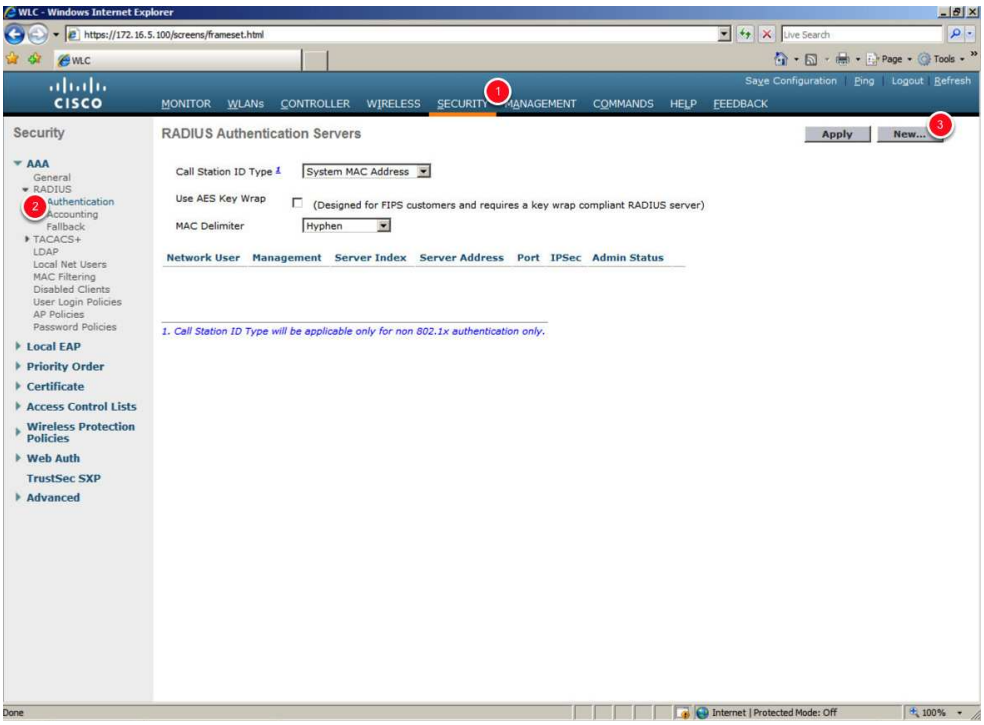
ACLs configured on the WLC can optionally have a direction specified, were the point of reference is the WLC itself:

- **Inbound** will apply to packets received from WiFi clients.
- **Outbound** will apply to packets sent toward WiFi clients.

Cisco Centralized Key Management (CCKM) is a Cisco proprietary feature that allows access-points to share client-negotiated keys through the WLC; this boosts the roaming performance because it allows clients to bypass a full reauthentication and bypass a key handshake. Although CCKM and 802.1x are both enabled, the access-point will try to negotiate CCKM first and fail over to regular 802.1x if the client does not support CCKM.

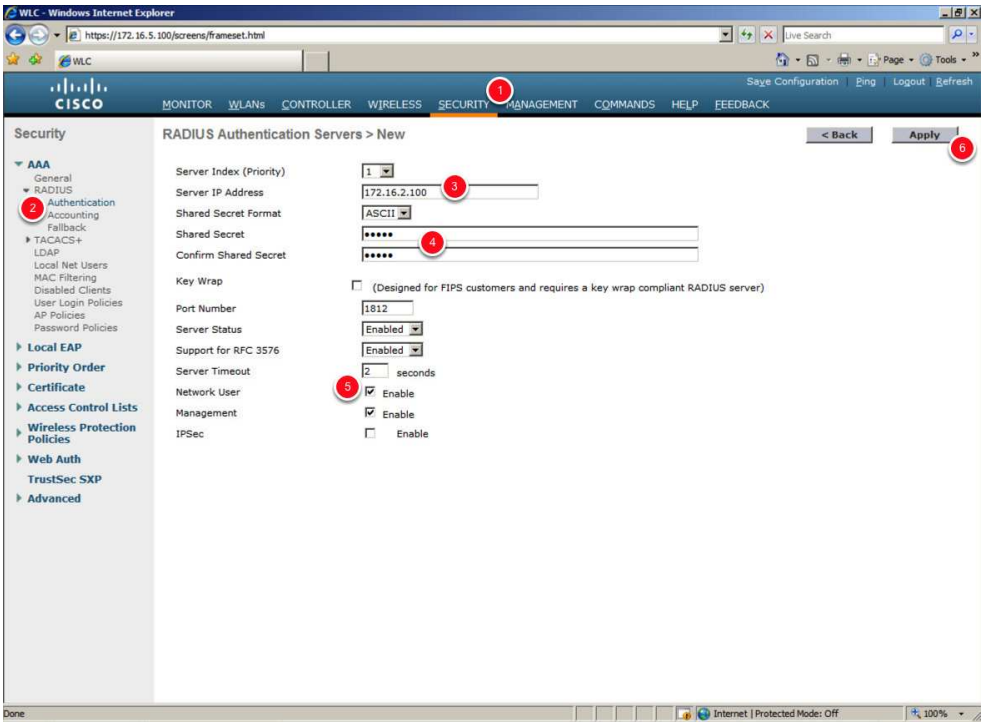
Configuration

Navigate through the WLC GUI to add the ACS RADIUS server:

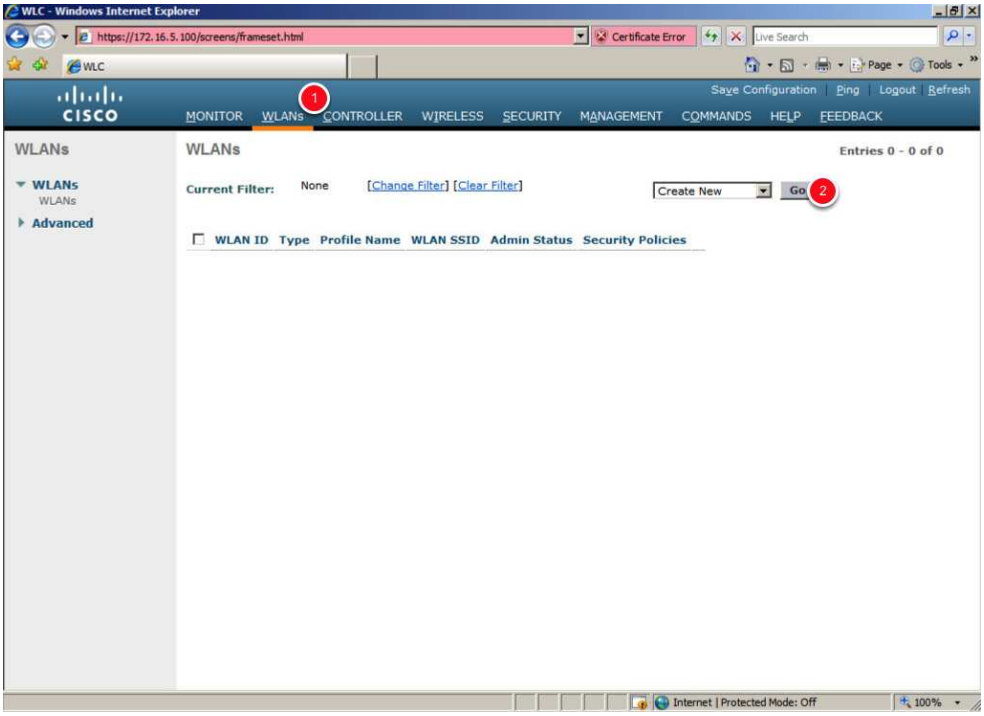


https://t.me/learningnets

Add the RADIUS server:

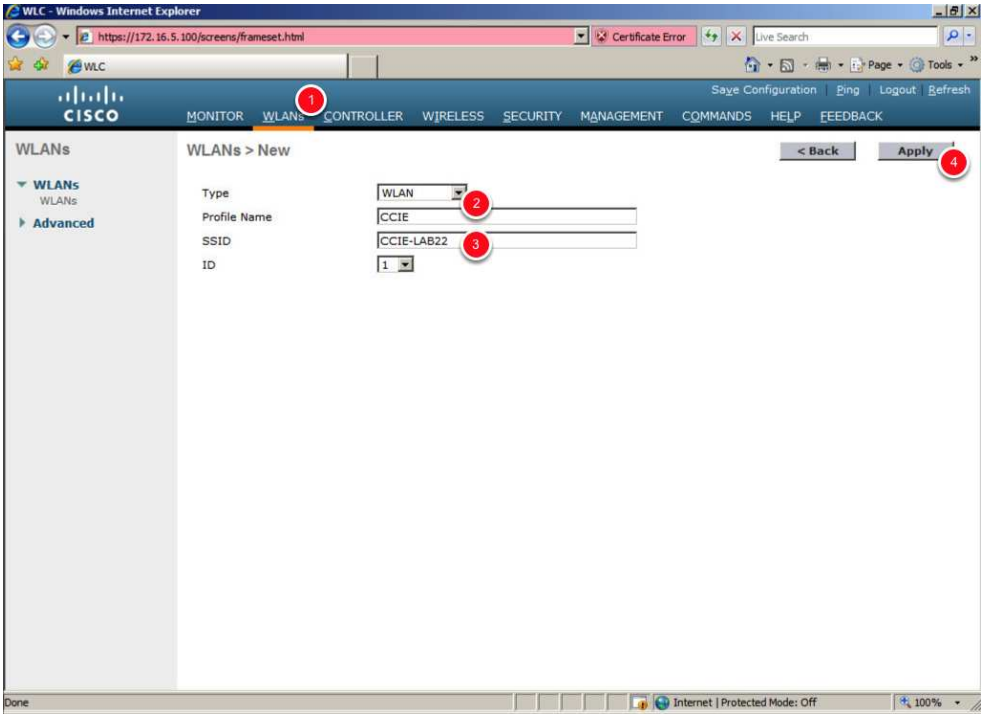


Navigate through WLC GUI to create a new WLAN:

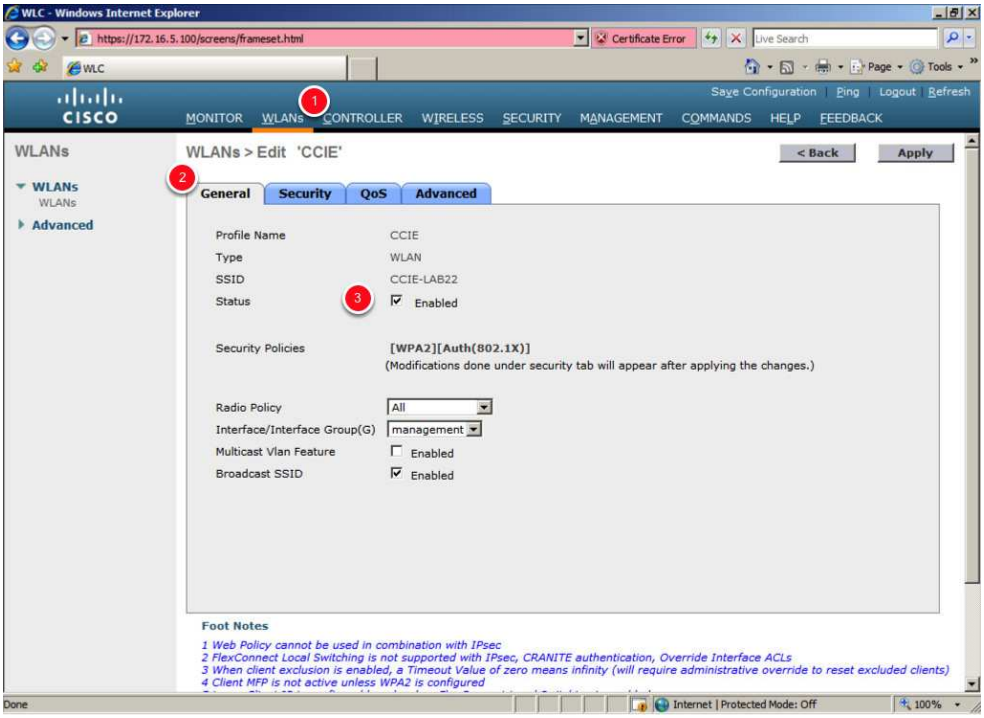


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

https://t.me/learningnets

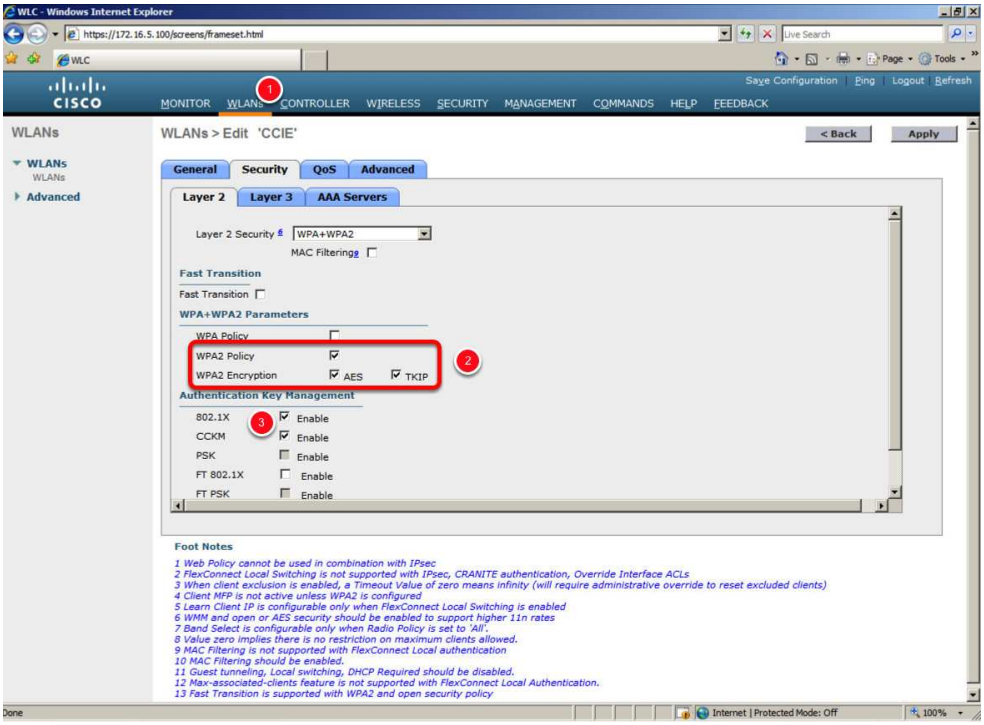


Enable the WLAN and leave WLC to broadcast the SSID:

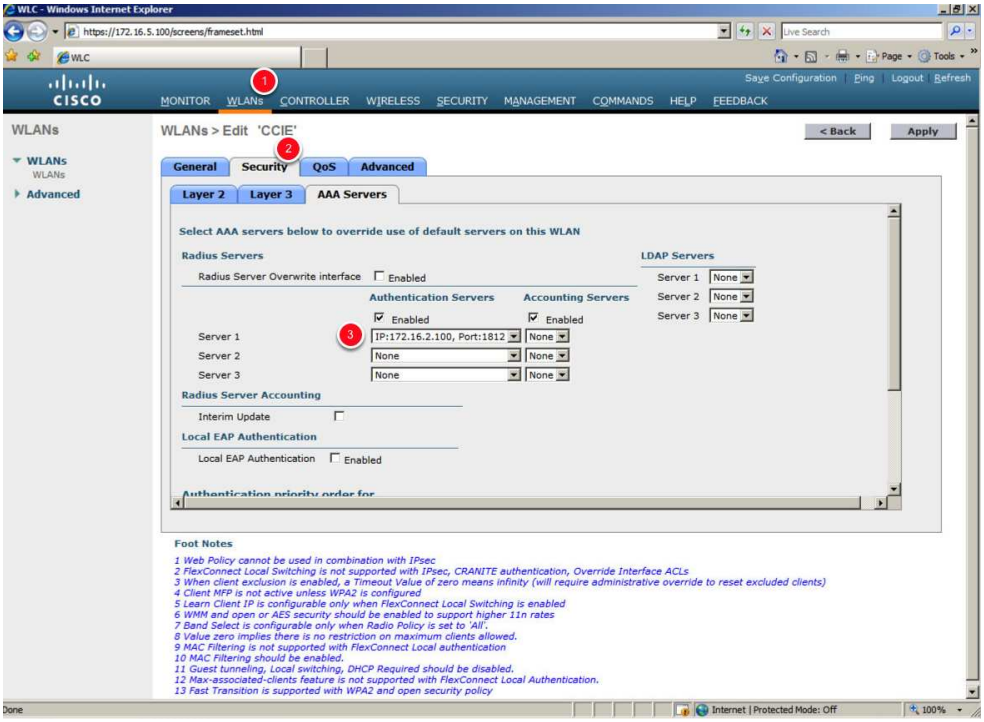


Configure WLAN security only for WPA2 and enable CCKM with 802.1x:

https://t.me/learningnets

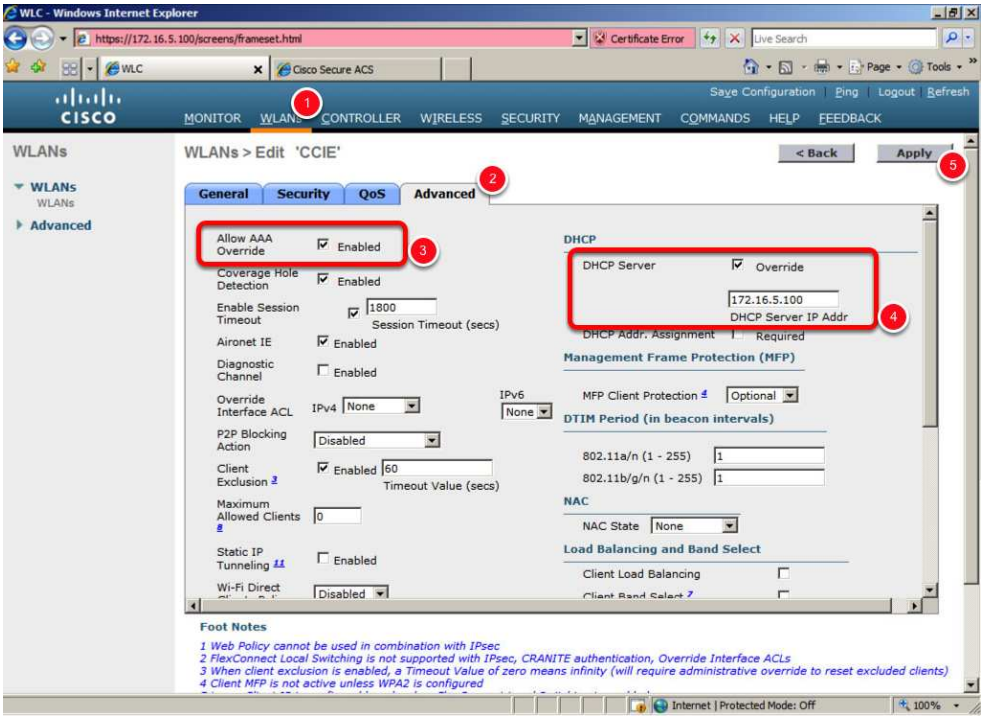


Activate the RADIUS server to be used for 802.1x WLAN authentication:

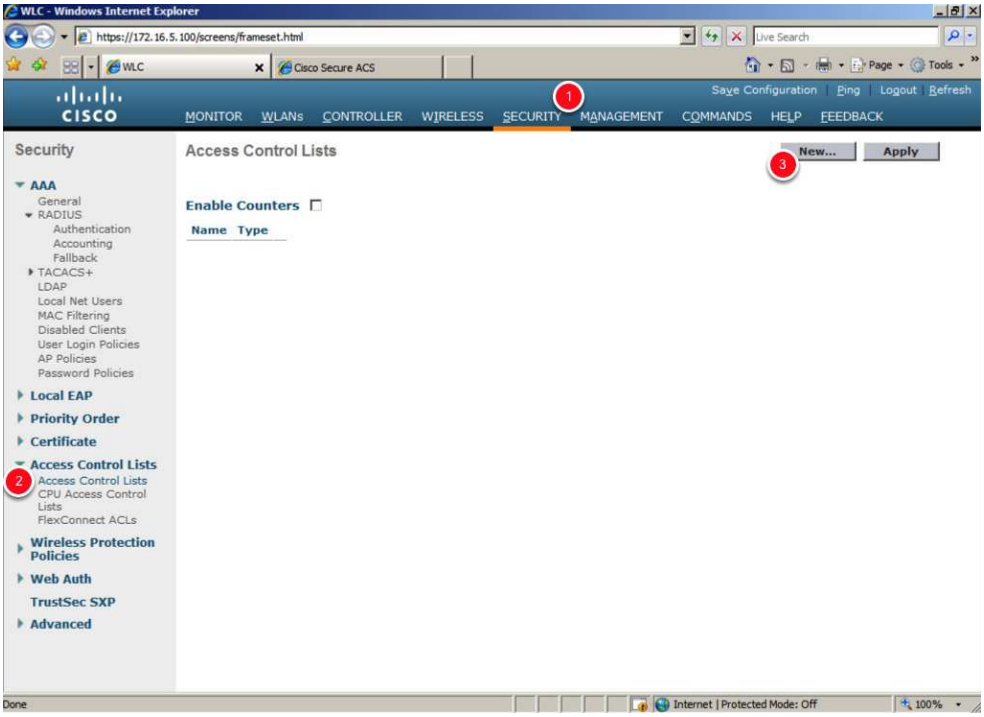


Enable AAA override at the WLAN level (required so that ACL received from RADIUS is applied) and DHCP override (to use the internal DHCP server):

https://t.me/learningnets

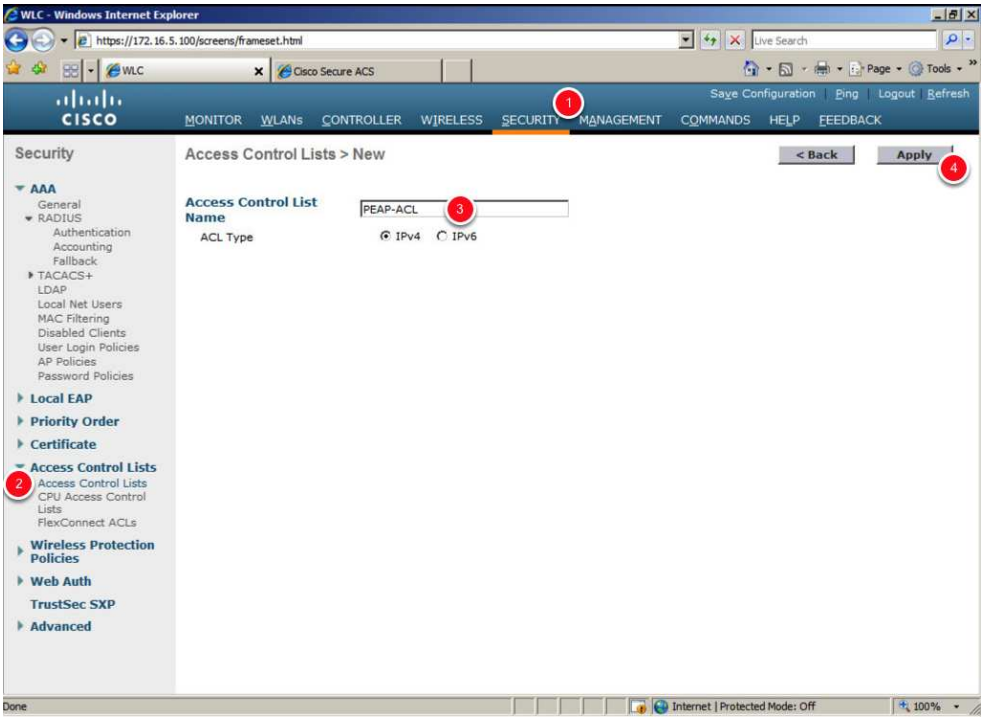


Navigate through the WLC GUI to configure an ACL:

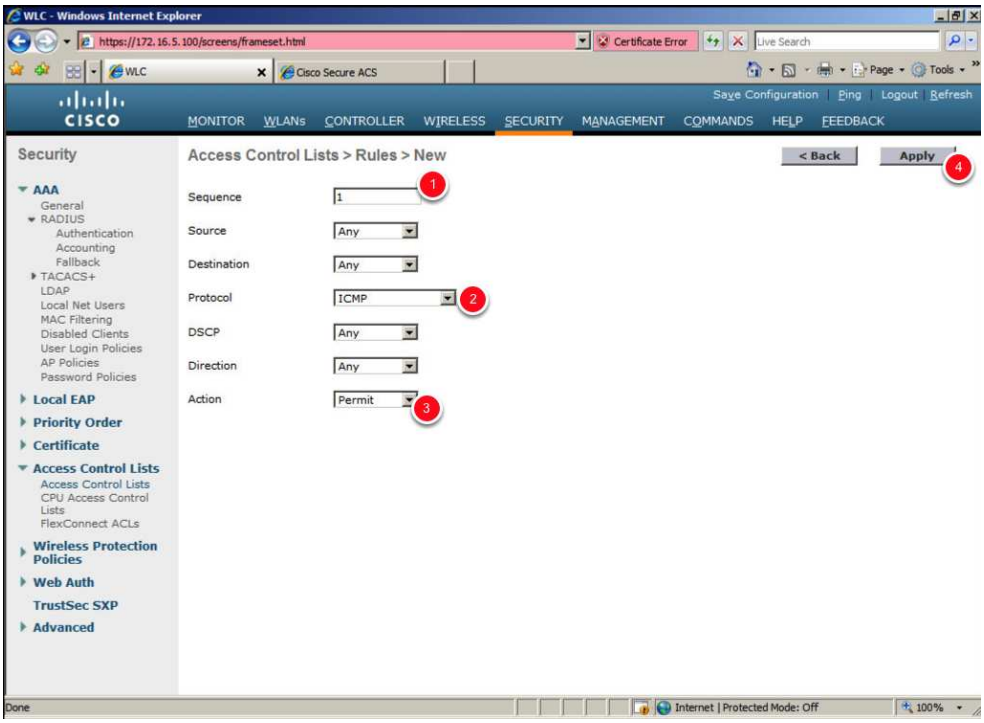


Select to configure an IPv4 ACL and give it a name that will be used later in the ACS RADIUS authorization profile:

https://t.me/learningnets

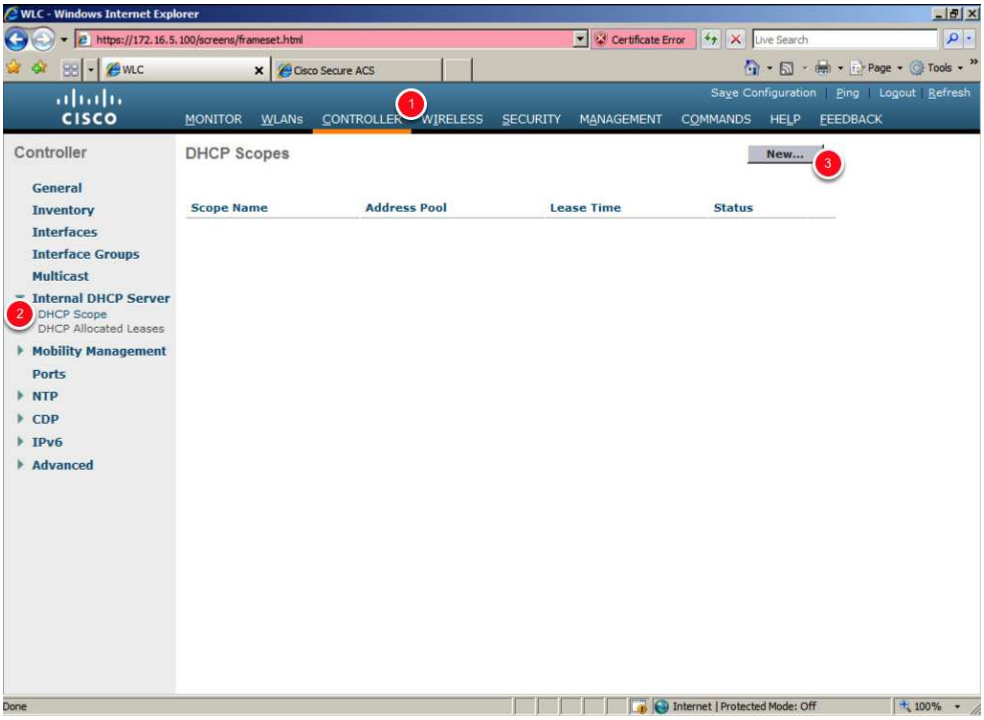


Edit the ACL, and select ICMP protocol and action of permit:

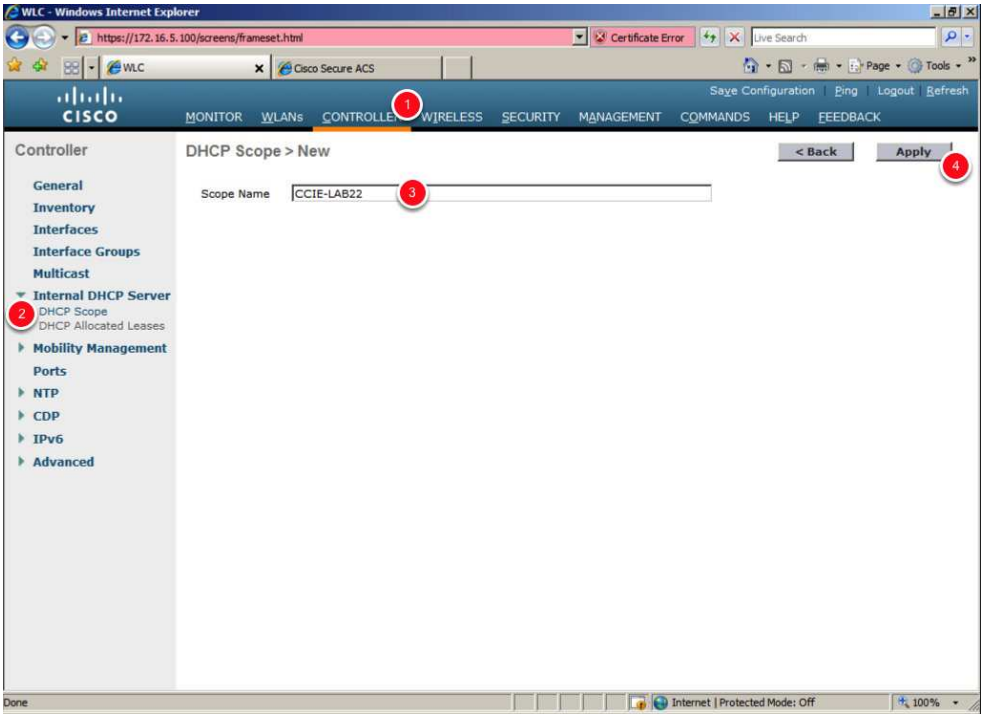


Navigate through the WLC GUI to create a new DHCP scope for the WLC internal server:

https://t.me/learningnets

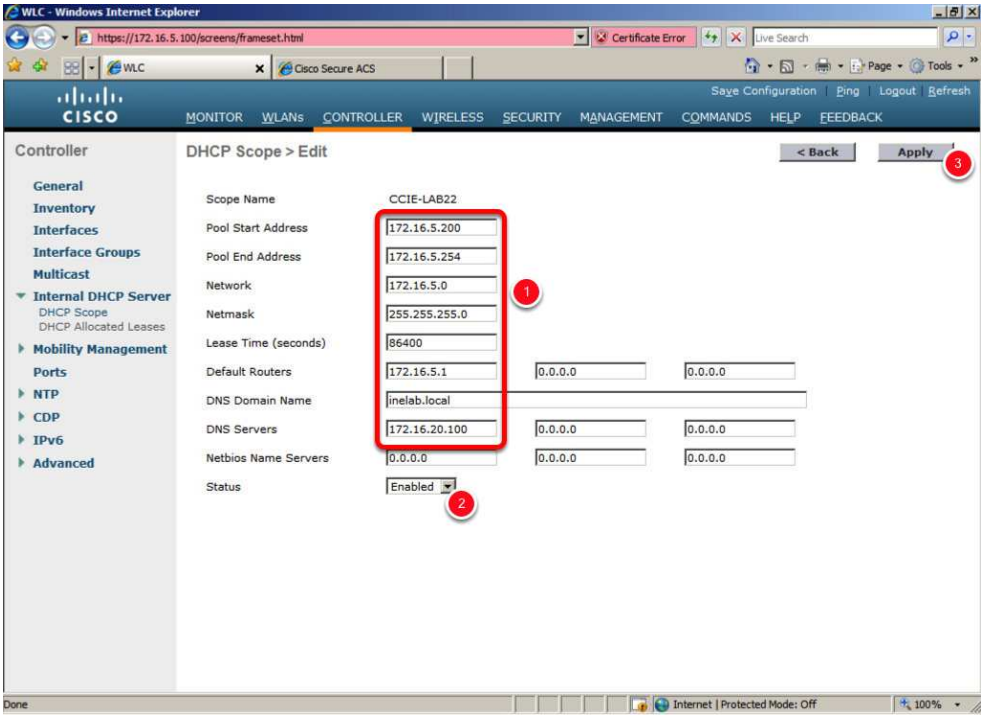


Give the DHCP scope a name:

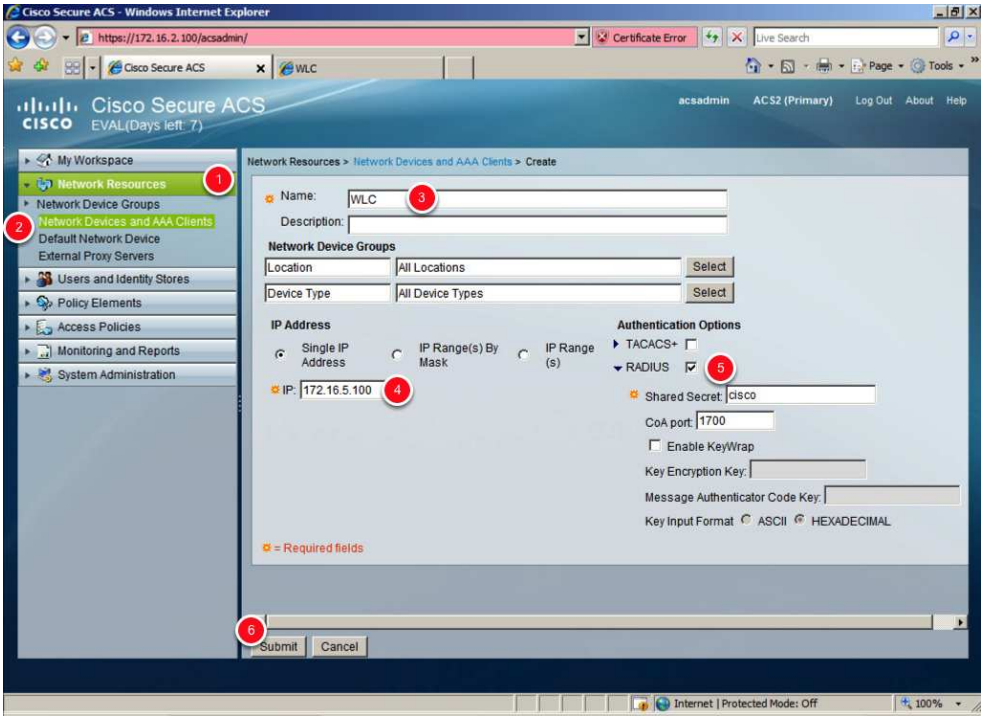


Configure the DHCP settings and enable it; it is disabled by default:

https://t.me/learningnets

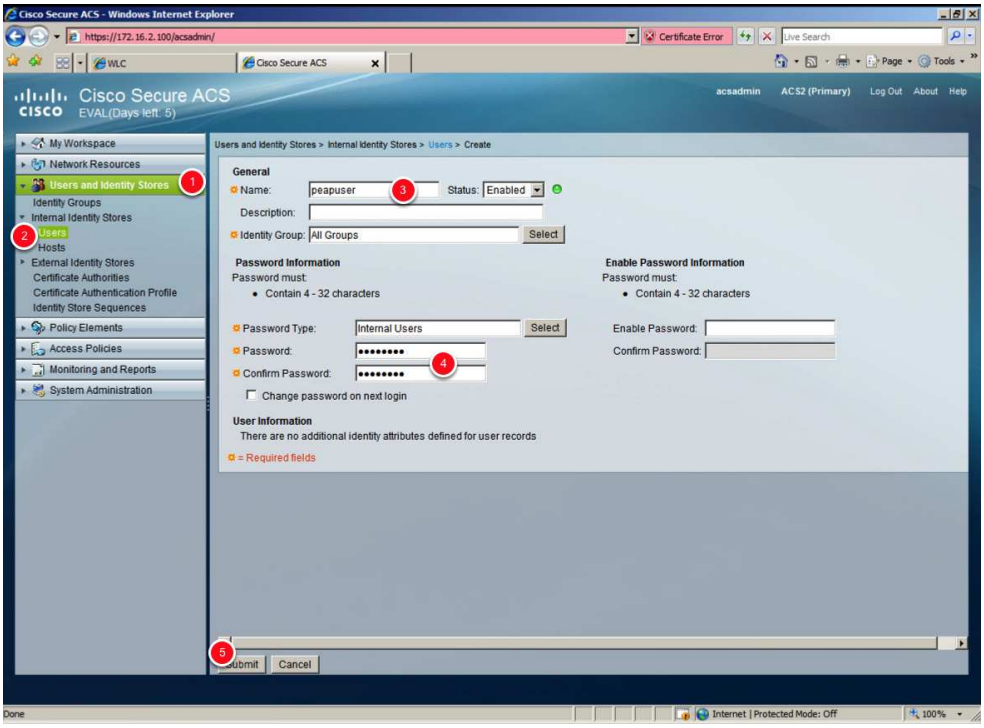


Add WLC as a NAD device into ACS, using RADIUS protocol:

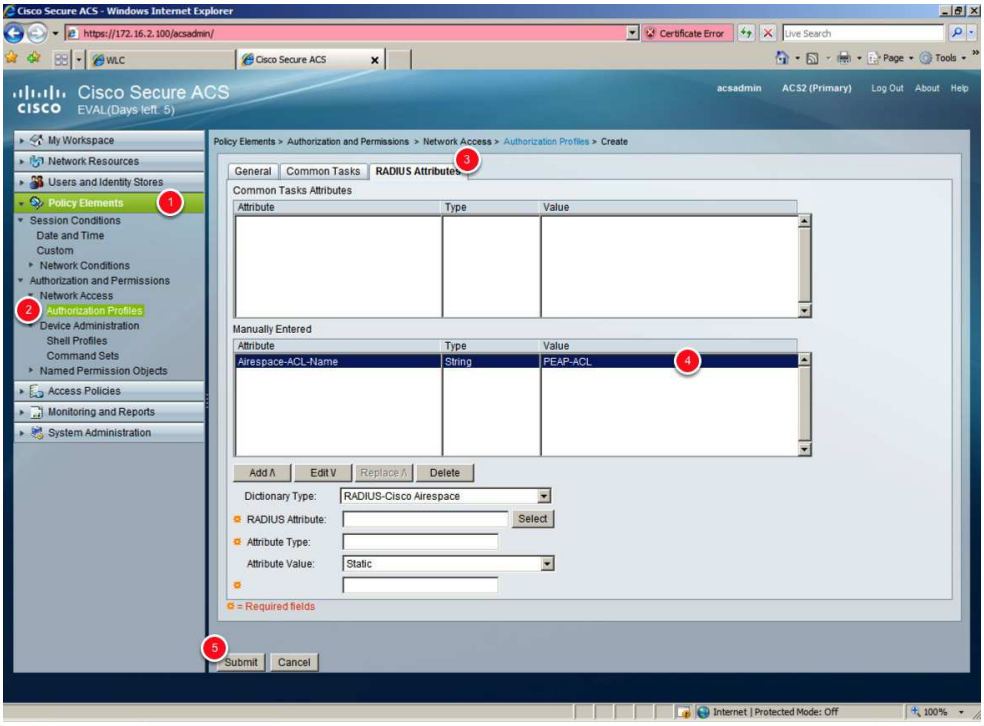


Add the user required for PEAP(EAP-MSCAHPv2) authentication:

https://t.me/learningnets

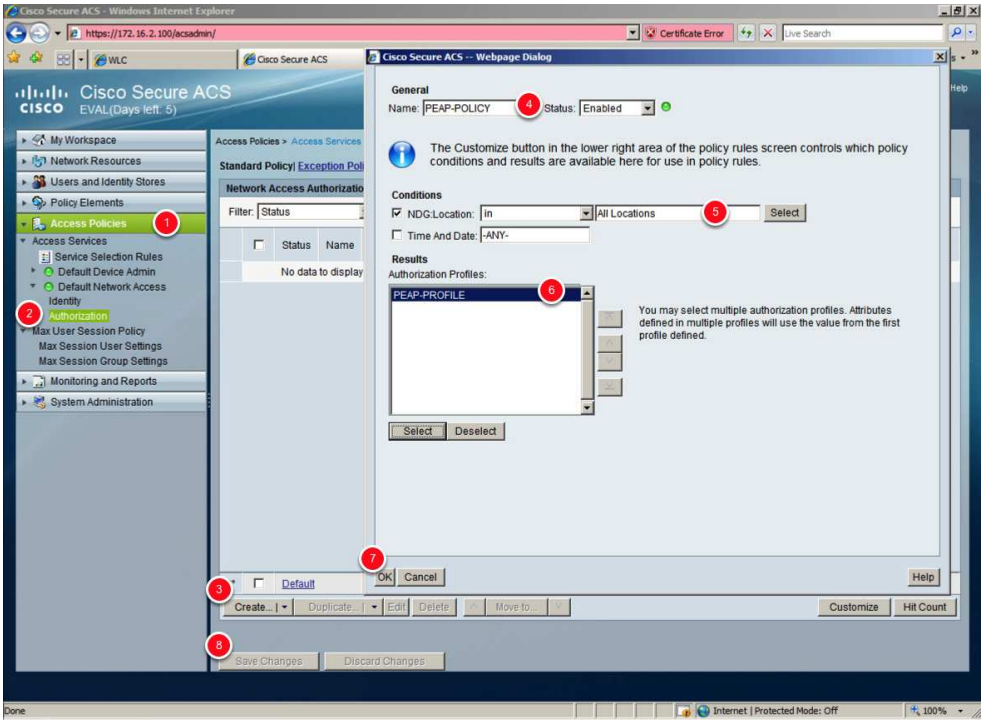


Configure a RADIUS authorization profile, and use the Cisco Airespace-ACL-Name attribute to apply the ACL defined on the WLC:

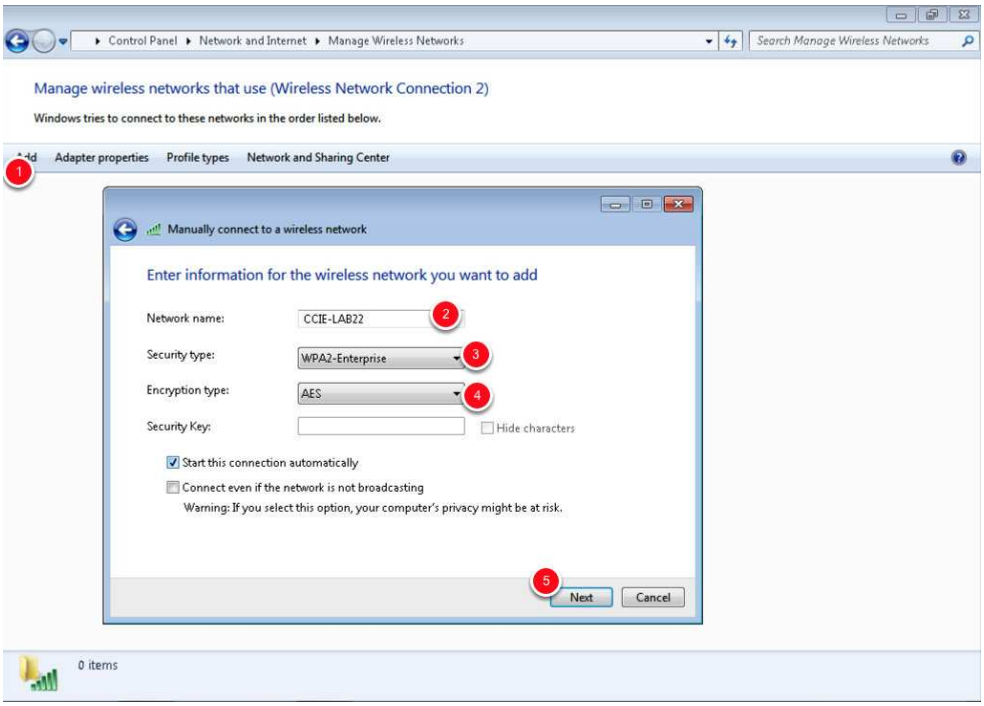


Configure a RADIUS access policy, and apply the previously configured profile:

https://t.me/learningnets

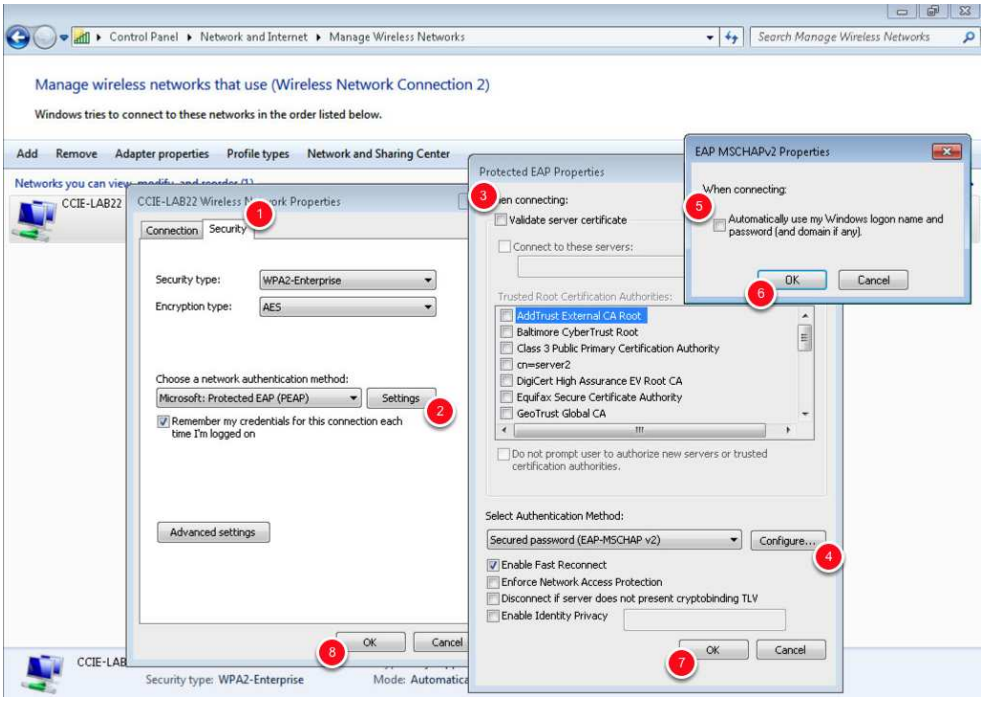


Go to the WiFi TEST-PC and manually add the WiFi network:



Edit the configured network, and make sure to disable the **Validate Server Certificate** option (because ACS is using a self-signed certificate); disable using the Windows credentials:

https://t.me/learningnets



You'll be asked for credentials in order to authenticate; use the ones configured on the ACS RADIUS server.

Verification

Verify the RADIUS server configuration:

```
(Cisco Controller) >show radius summary

Vendor Id Backward Compatibility..................... Disabled
Call Station Id Case................................. lower
Call Station Id Type................................. Mac Address
Aggressive Failover.................................. Enabled
Keywrap.............................................. Disabled
Fallback Test:
    Test Mode........................................ Off
    Probe User Name.................................. cisco-probe
    Interval (in seconds)............................ 300
MAC Delimiter for Authentication Messages........ hyphen
MAC Delimiter for Accounting Messages............ hyphen

Authentication Servers

Idx  Type  Server Address    Port    State     Tout  MgmtTout  RFC3576  IPSec - AuthMode/Phase1/Group/Lifetime/Auth/
---  ----  ---------------  ------  --------  ----  --------  -------  ----------------------------------------
1    NM    172.16.2.100     1812    Enabled   2     2         Enabled   Disabled - none/unknown/group-0/0 none/none
```

https://t.me/learningnets

```
Accounting Servers

Idx  Type  Server Address    Port    State     Tout  MgmtTout  RFC3576  IPSec - AuthMode/Phase1/Group/Lifetime/Auth/

--More-- or (q)uit
---  ----  ---------------  ------  --------  ----  --------  -------  ----------------------------------------
```

Verify that the client is associated and authenticated:

```
(Cisco Controller) >show client summary

Number of Clients.............................. 1

MAC Address       AP Name           Status        WLAN/RLAN      Auth Protocol         Port Wired
----------------- ----------------- ------------ -------------- ---- ---------------- ---- -----
ac:f1:df:13:bd:8f AP0006.f63b.9ada  Associated    1              Yes  802.11n(2.4 GHz) 1    N/A
```

Verify that DHCP scope on WLC is enabled and WiFi TEST-PC has been assigned an IP address:

```
(Cisco Controller) >show dhcp summary

  Scope Name                    Enabled          Address Range
CCIE-LAB22                        Yes      172.16.5.200 -> 172.16.5.254
!
!(Cisco Controller) >show dhcp leases

       MAC               IP          Lease Time Remaining
ac:f1:df:13:bd:8f    172.16.5.200    23 hours 53 minutes 58 seconds
```

Verify that the client has been assigned the IP address and the ACL has also been applied:

```
(Cisco Controller) >show client detail ac:f1:df:13:bd:8f
Client MAC Address............................. ac:f1:df:13:bd:8f
Client Username ..............................  peapuser
AP MAC Address................................. f0:29:29:04:21:f0
AP Name........................................ AP0006.f63b.9ada
Client State................................... Associated
Client NAC OOB State........................... Access
Wireless LAN Id................................ 1
802.11u........................................ Not Supported
```

https://t.me/learningnets

```
BSSID..................................... f0:29:29:04:21:f0
Connected For ............................ 610 secs
Channel................................... 11
IP Address................................ 172.16.5.200
IPv6 Address.............................. fe80::3c16:4073:f3a4:c679
Association Id............................ 1
Authentication Algorithm.................. Open System
Reason Code............................... 1
Status Code............................... 0
Client CCX version........................ 5
Client E2E version........................ No E2E support
Diagnostics Capability.................... Supported
S69 Capability............................ Not Supported
Re-Authentication Timeout................. 1193
QoS Level................................. Silver
802.1P Priority Tag....................... disabled
CTS Security Group Tag.................... Not Applicable
KTS CAC Capability........................ No
WMM Support............................... Enabled
  APSD ACs.................................  BK  BE  VI  VO
Power Save................................ OFF
Current Rate.............................. 5.5
Supported Rates........................... 1.0,2.0,5.5,11.0,6.0,9.0,
    ..................................... 12.0,18.0,24.0,36.0,48.0,
    ..................................... 54.0
Mobility State............................ Local
Mobility Move Count....................... 0
Security Policy Completed................. Yes
Policy Manager State...................... RUN
Policy Manager Rule Created............... Yes
Audit Session ID.......................... ac1005640000000552a71787
IPv4 ACL Name............................. PEAP-ACL
IPv4 ACL Applied Status................... Yes
IPv6 ACL Name............................. none
IPv6 ACL Applied Status................... Unavailable
Policy Type............................... WPA2
Authentication Key Management............. 802.1x
```

Verify the ACL settings:

```
(Cisco Controller) >show acl summary

ACL Counter Status              Disabled
-----------------------------------IPv4 ACL Name                     Applied
```

https://t.me/learningnets

```
-------------------------------- -------PEAP-ACL                          Yes

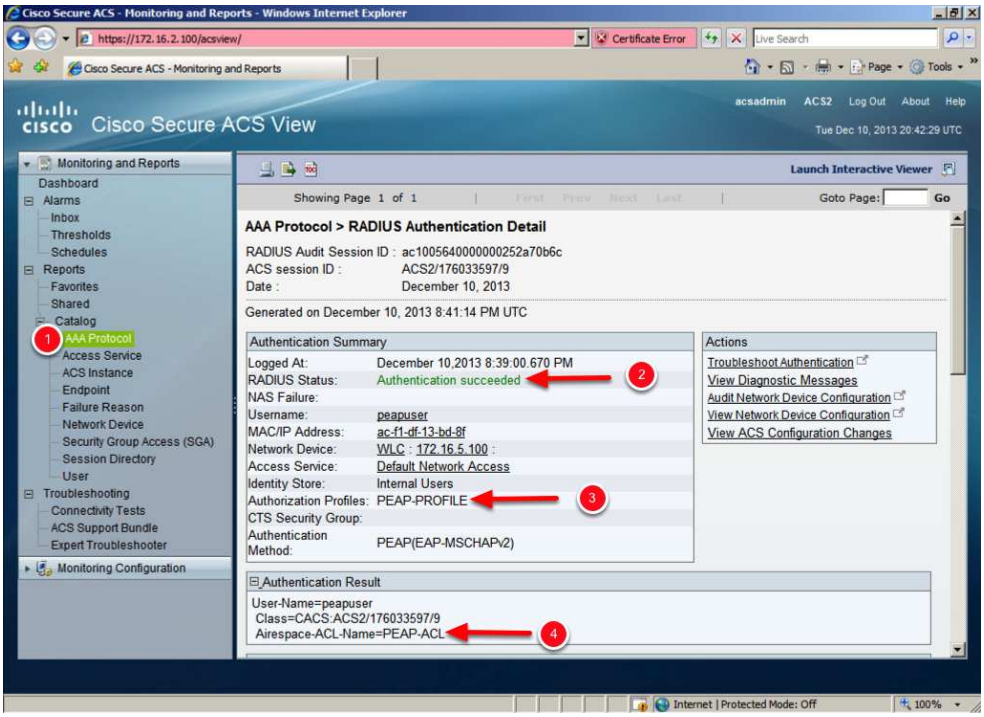

---------------------------------------
IPv6 ACL Name                    Applied
-------------------------------- -------
!
!(Cisco Controller) >show acl detailed PEAP-ACL

                        Source                          Destination                Source Port  Dest Port
Index  Dir       IP Address/Netmask               IP Address/Netmask       Prot    Range       Range    DSCP  Action
------ --- ------------------------------- ------------------------------- ---- ----------- ----------- ----- ------

1 Any         0.0.0.0/0.0.0.0                 0.0.0.0/0.0.0.0            1     0-65535     0-65535  Any Permit

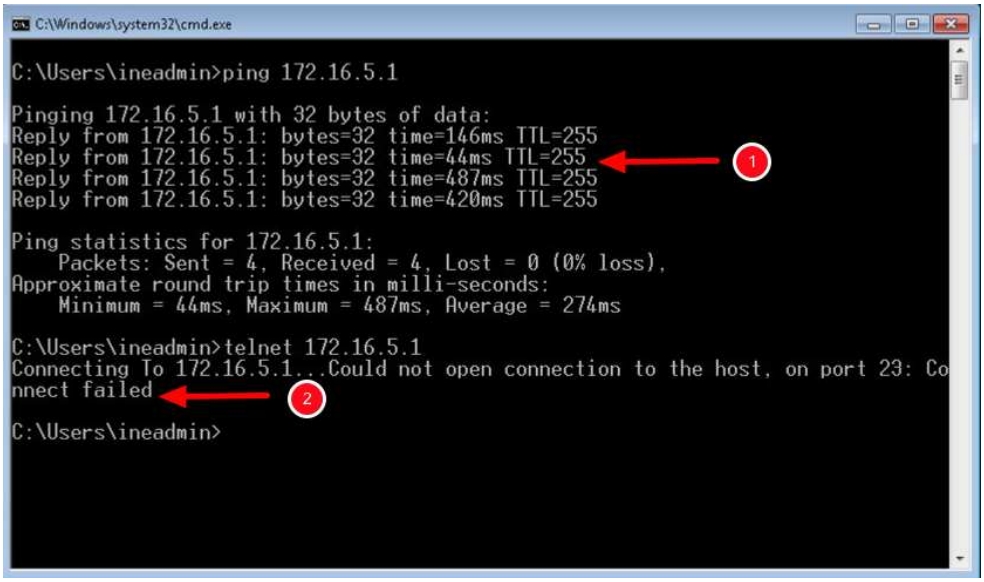

 DenyCounter : 0
```

Verify the RADIUS authentication logs, and check that ACL name has been sent as part of the authorization process:



Verify from the client that ICMP works, but for example TCP traffic is filtered:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

PEAP (EAP-MSCHAPv2) FlexConnect Mode

Task

- Configure WLC with a SSID named **CCIE-XX**, where **X** equals the rack number you've been assigned.
  - Use PEAP (EAP-MSCHAPv2) authentication via the ISE RADIUS server.
  - Enable WPA2 authentication with 802.1x.
  - Authenticated users should have only ICMP access.
  - User traffic for this SSID should not be tunneled to the controller and assigned to VLAN 147 (136.1.147.0/24).
- Activate SW5 as the DHCP server and default gateway for VLAN 147.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands -->Download File**).

Overview

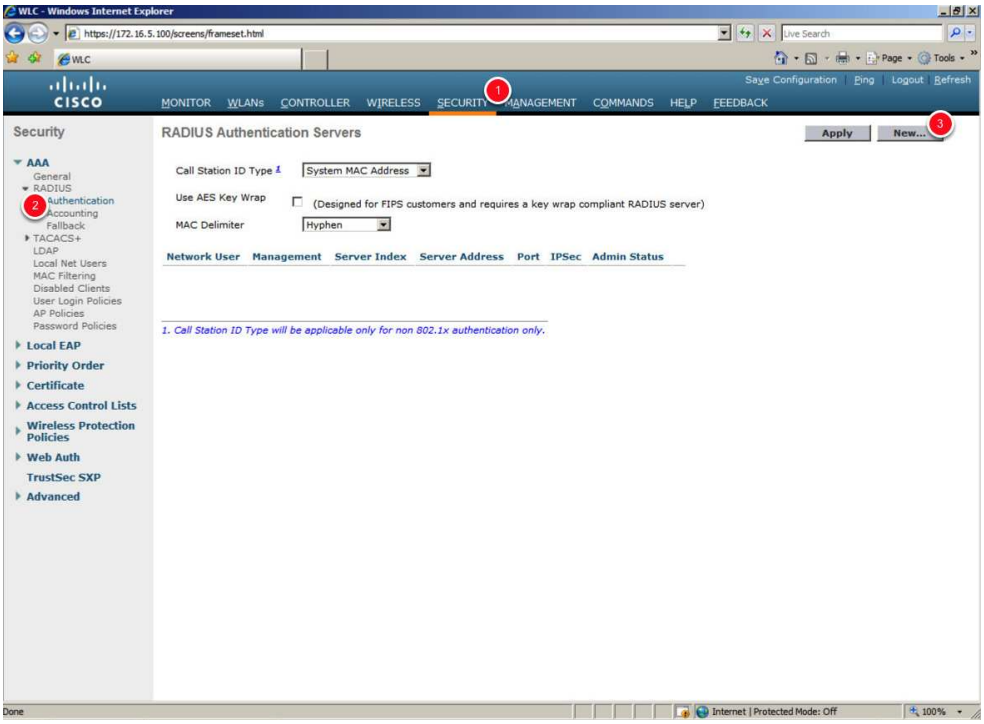
PEAP authentication and flexconnect have already been outlined in previous tasks. The difference here is the use of flexconnect ACLs, which do not have a direction option when being configured. Flexconnect ACLs can be applied at the AP or

https://t.me/learningnets

flexconnect group level, and direction is specified at the point of insertion. Additionally, WLC does not support **AAA Override** for flexconnect ACL in this code; it is supported starting with 7.5 version. **AAA Override** for flexconnect supports dynamic VLAN assignment for code 7.2.

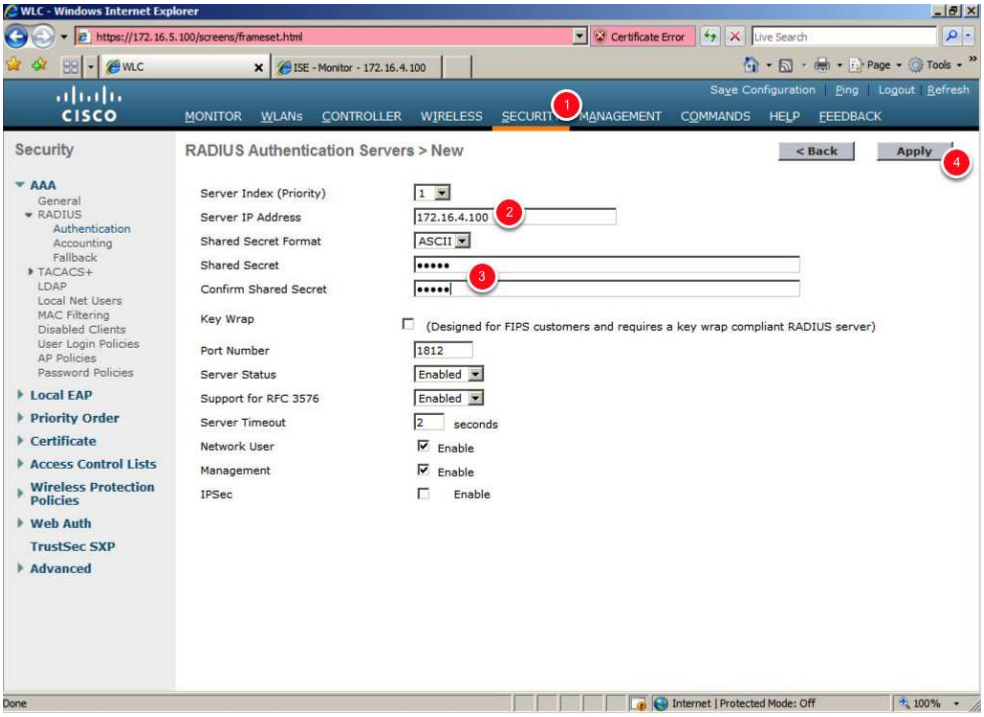
Configuration

Navigate through the WLC GUI to add the ISE RADIUS server:

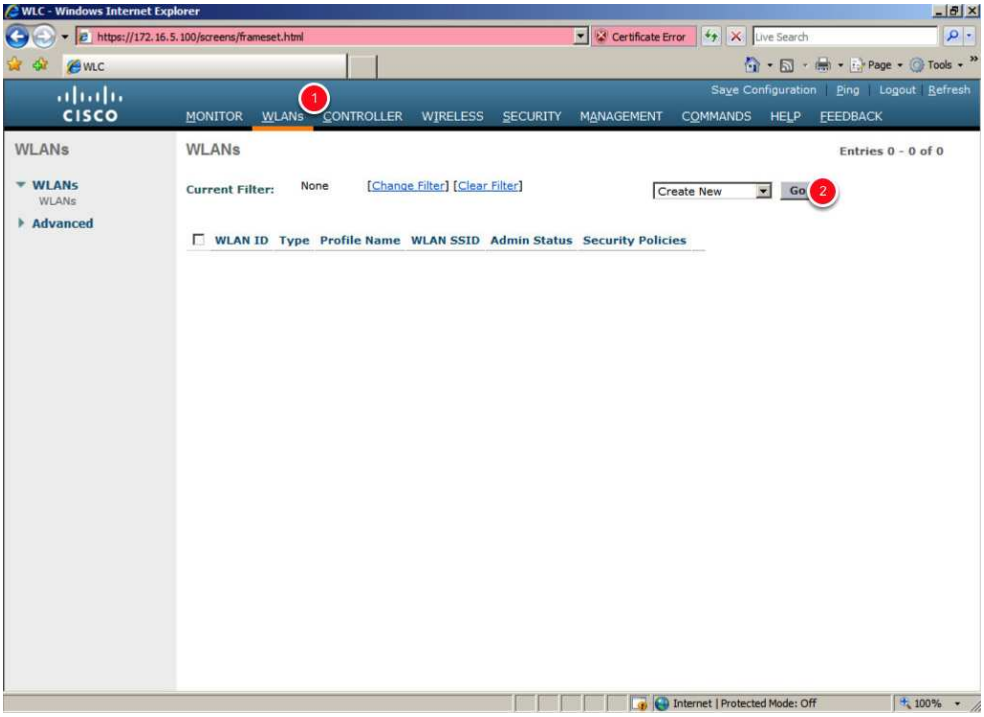


Add the RADIUS server:

https://t.me/learningnets

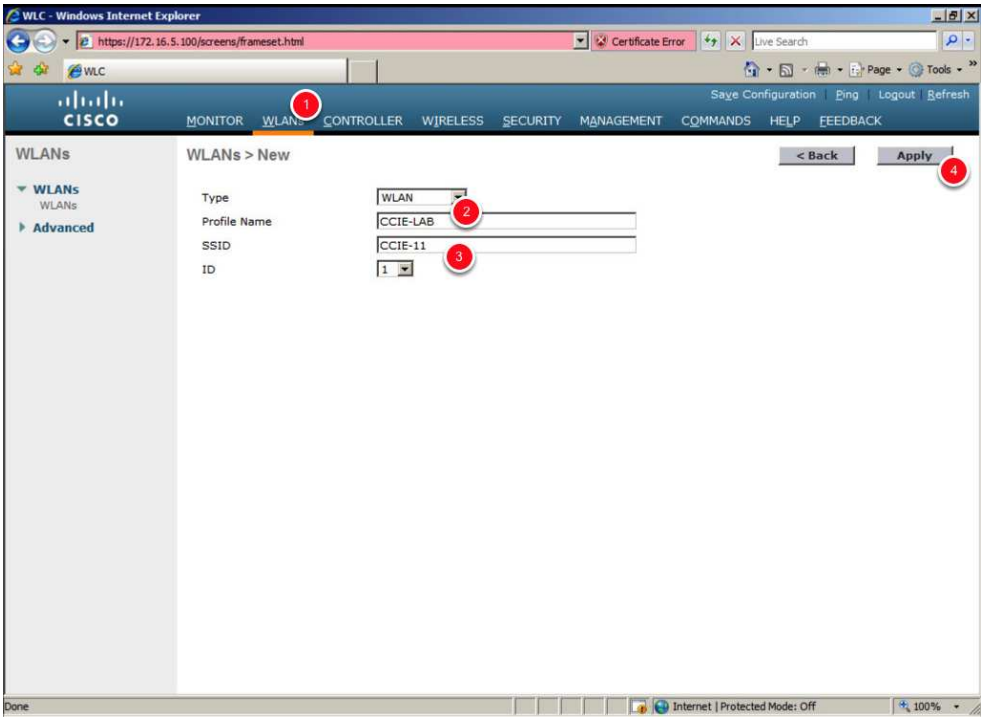


Navigate through WLC GUI to create a new WLAN:

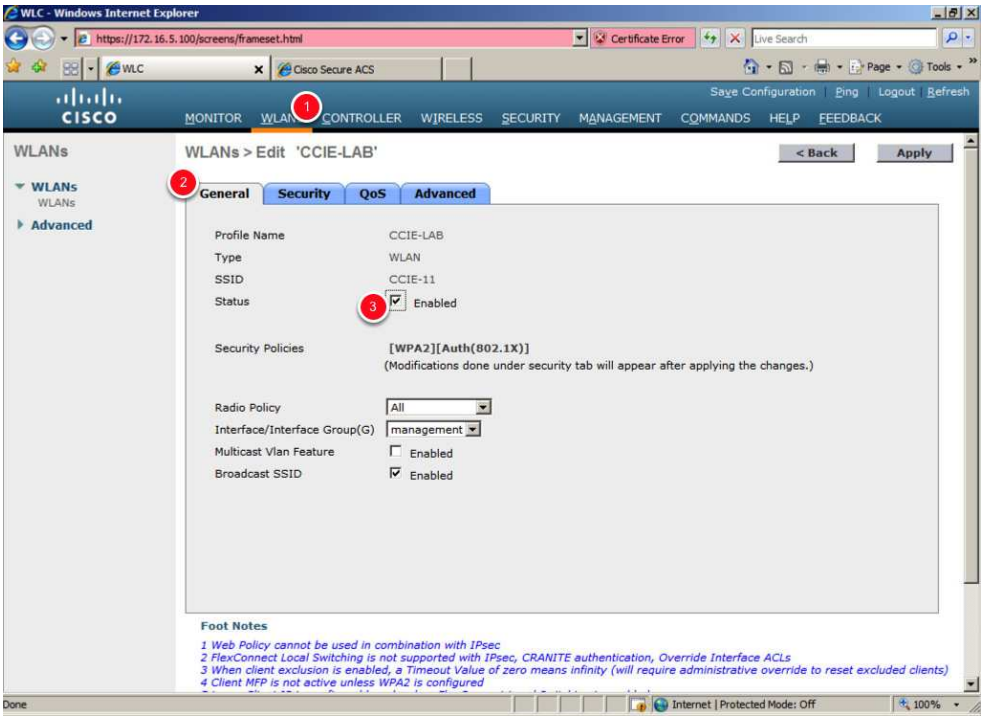


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

https://t.me/learningnets

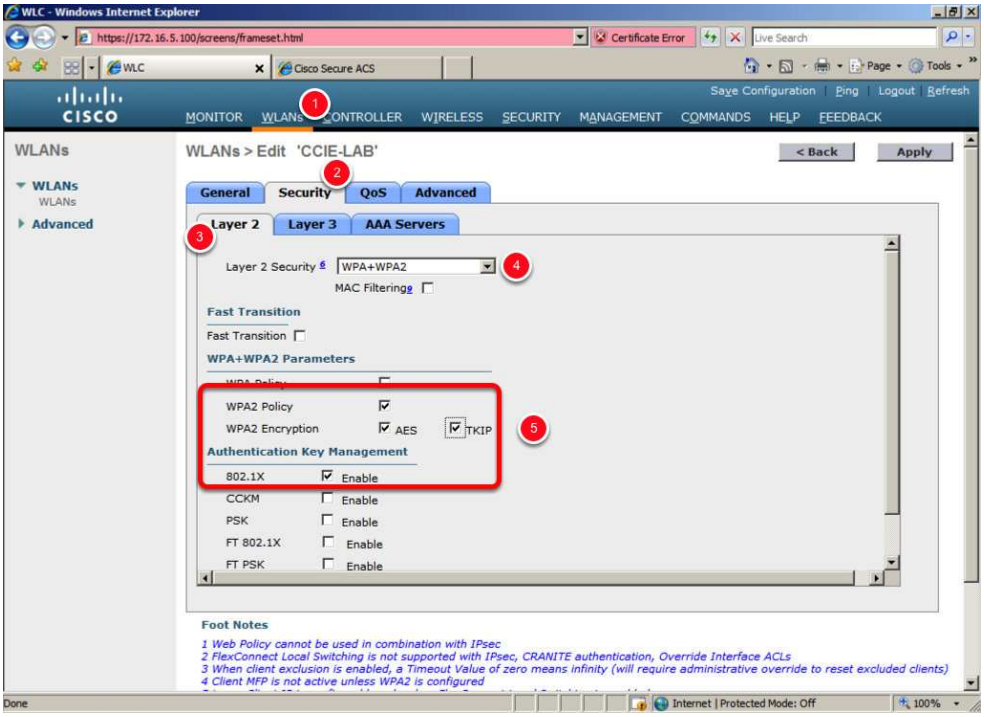


Enable the WLAN and leave it at default to be broadcasted:

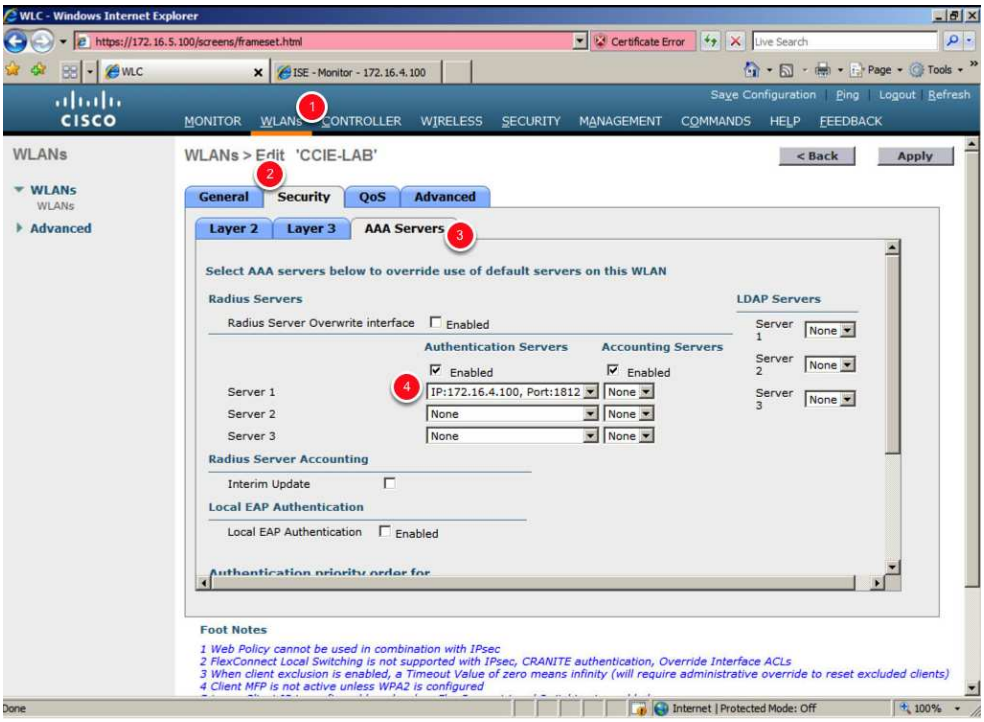


Configure the WLAN for 802.1x authentication with WPA2:

https://t.me/learningnets

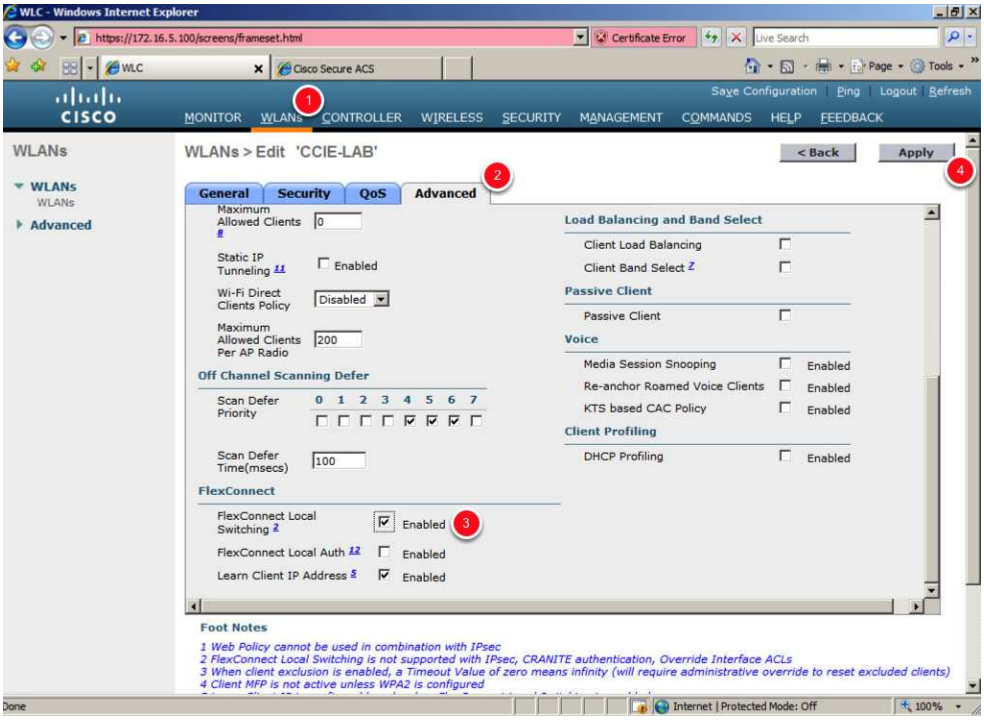


Specify the RADIUS server to be used for 802.1x authentication:

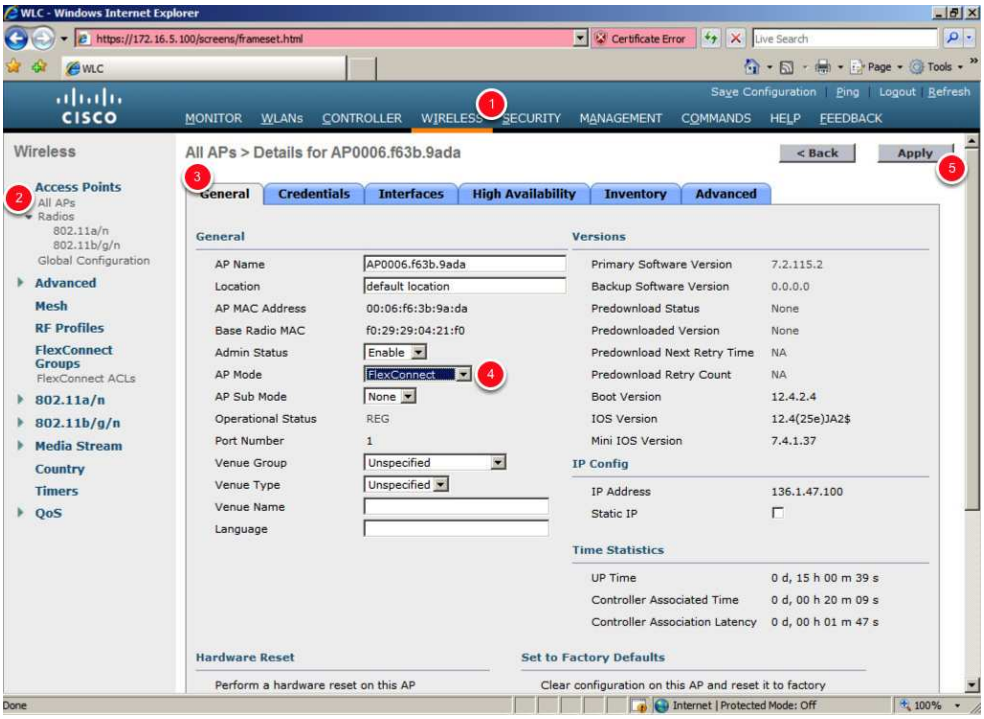


Configure the WLAN for local switching (flexconnect):

https://t.me/learningnets

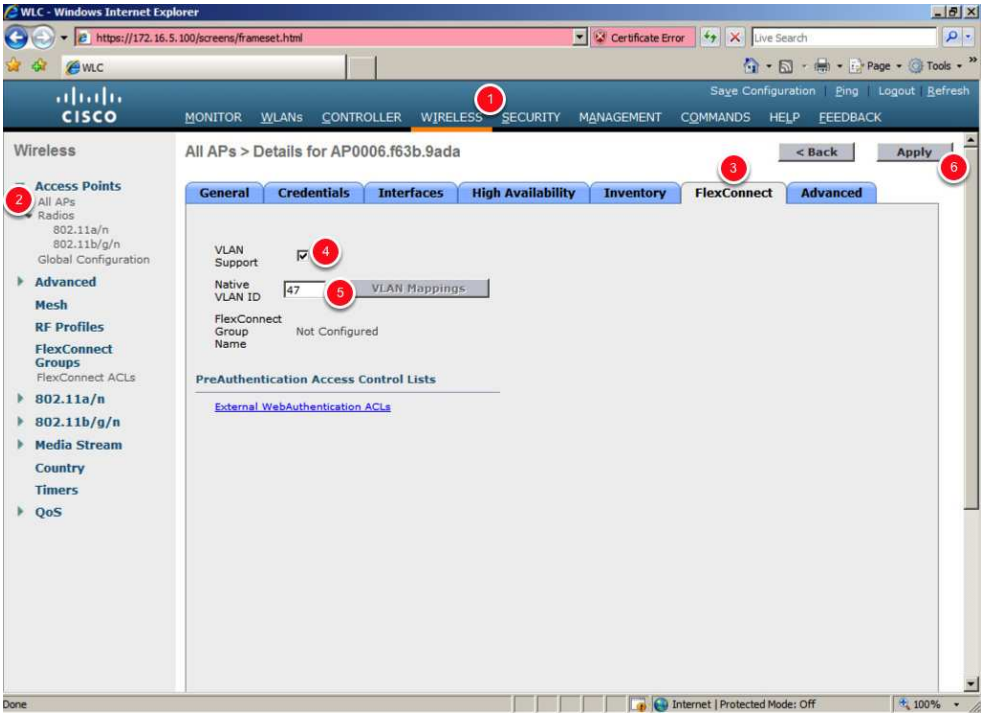


Configure the LAP to function in flexconnect mode; it will reload later:

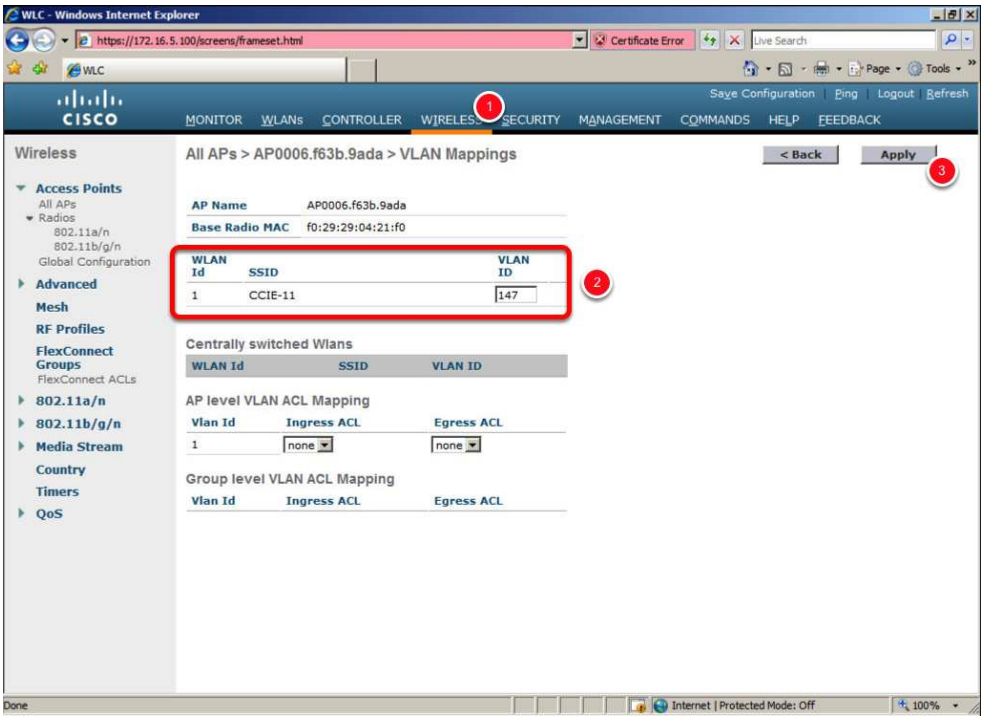


LAP will boot up in flexconnect mode and a new tab will be visible for LAP configuration; enable VLAN support and specify the native VLAN to be the management VLAN of the LAP. This is because LAP does not support VLAN tagging for the management VLAN:

https://t.me/learningnets

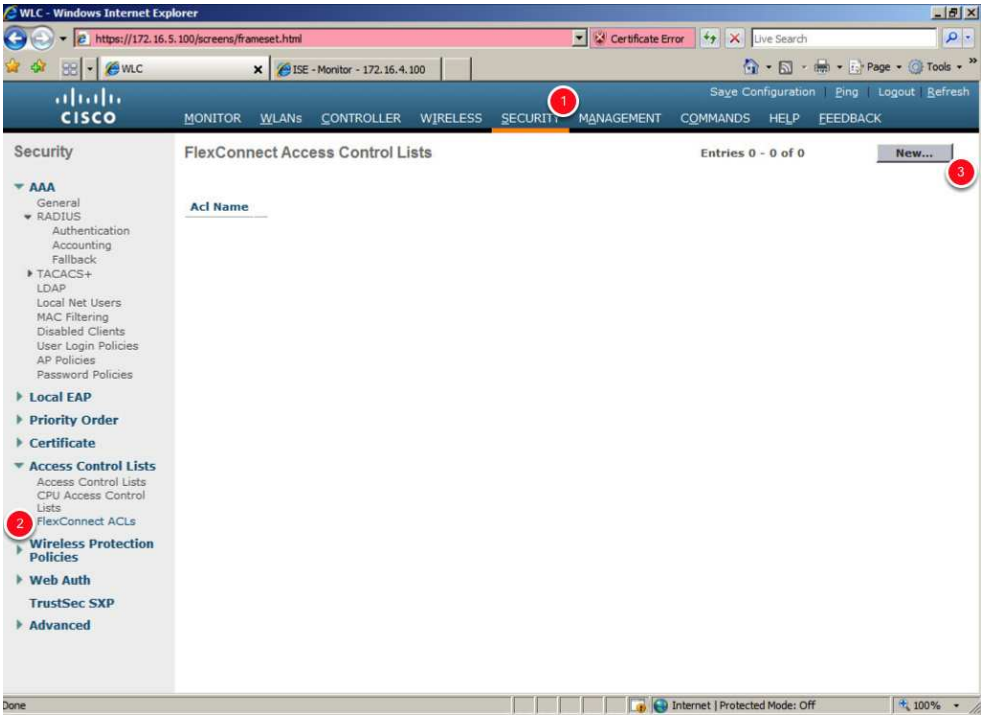


Navigate again to LAP settings and configure the VLAN mapping for the defined SSID to be VLAN 147:

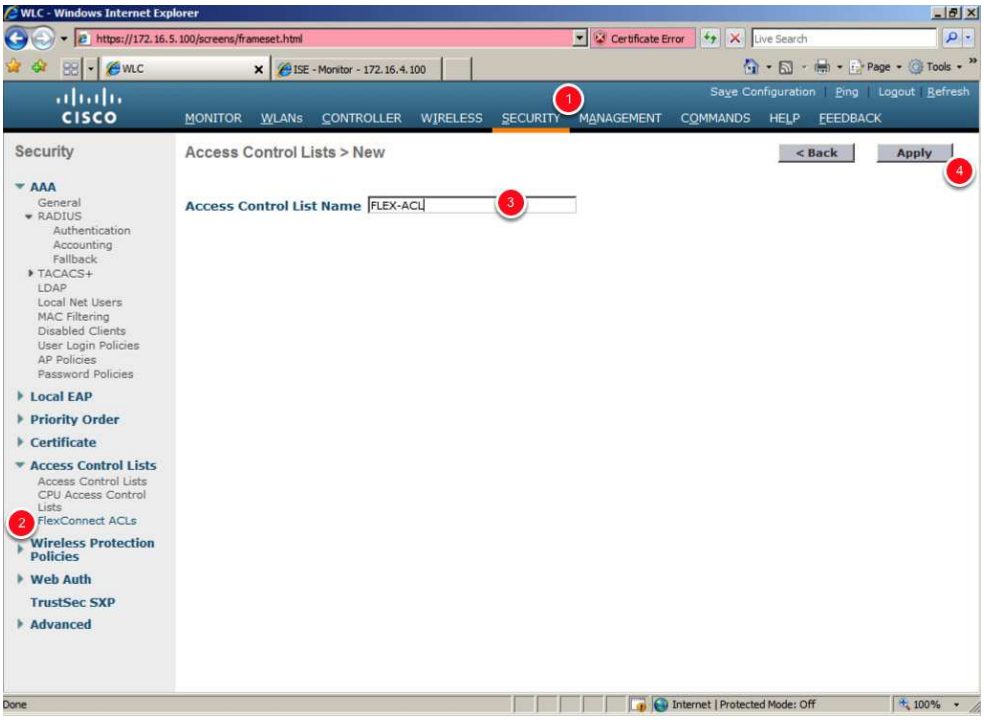


Navigate through the WLC GUI to configure a flexconnect ACL:

https://t.me/learningnets

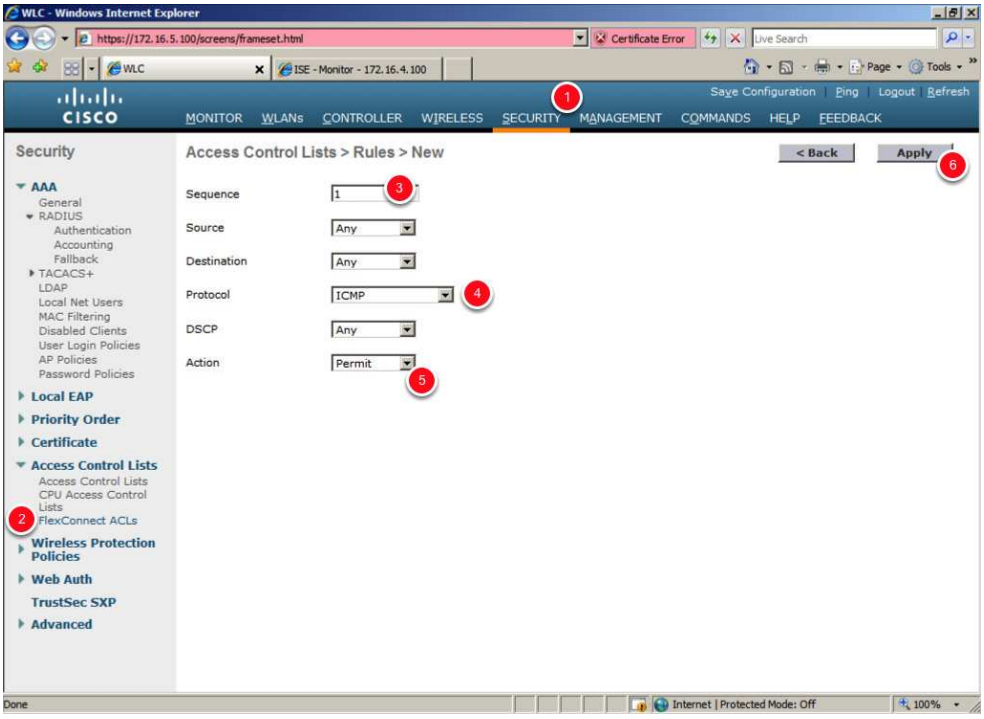


Give the ACL a relevant name:

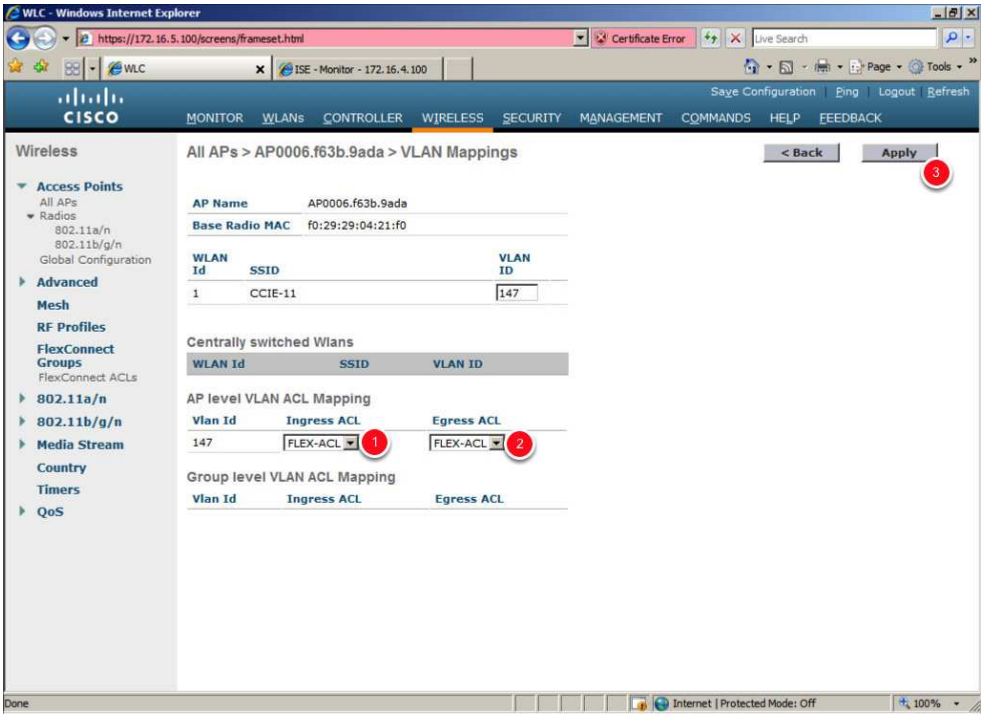


Configure required rules for the ACL:

https://t.me/learningnets

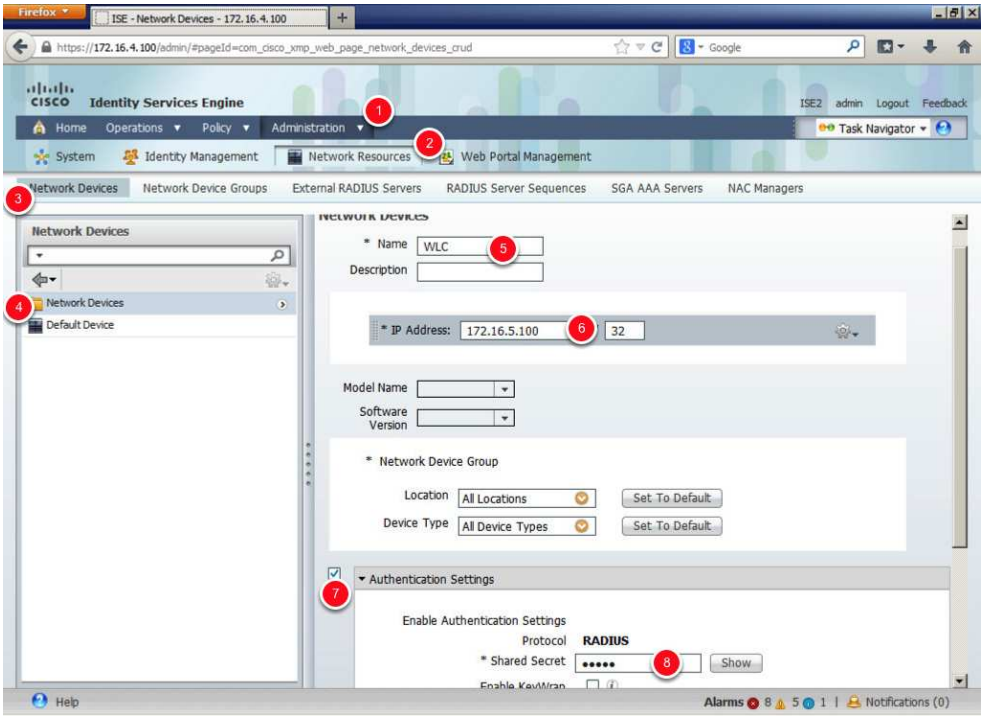


Apply the ACL at the LAP level; the other option is to configure a flexconnect group, add the LAP to it, and apply the ACL at the group level:

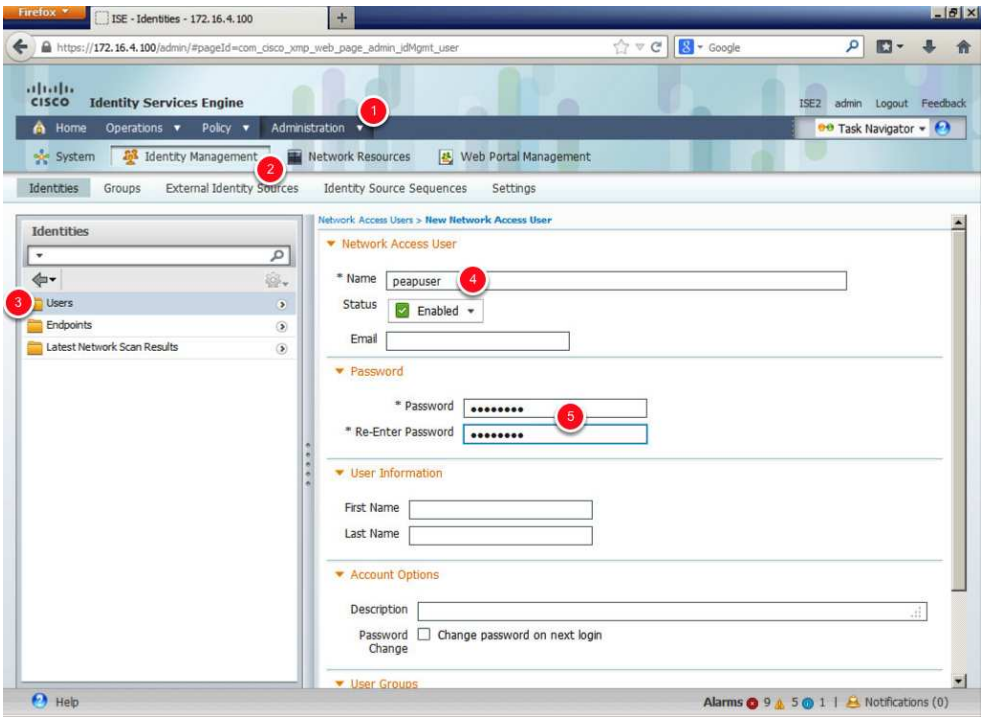


Add WLC as a NAD device into ISE:

https://t.me/learningnets

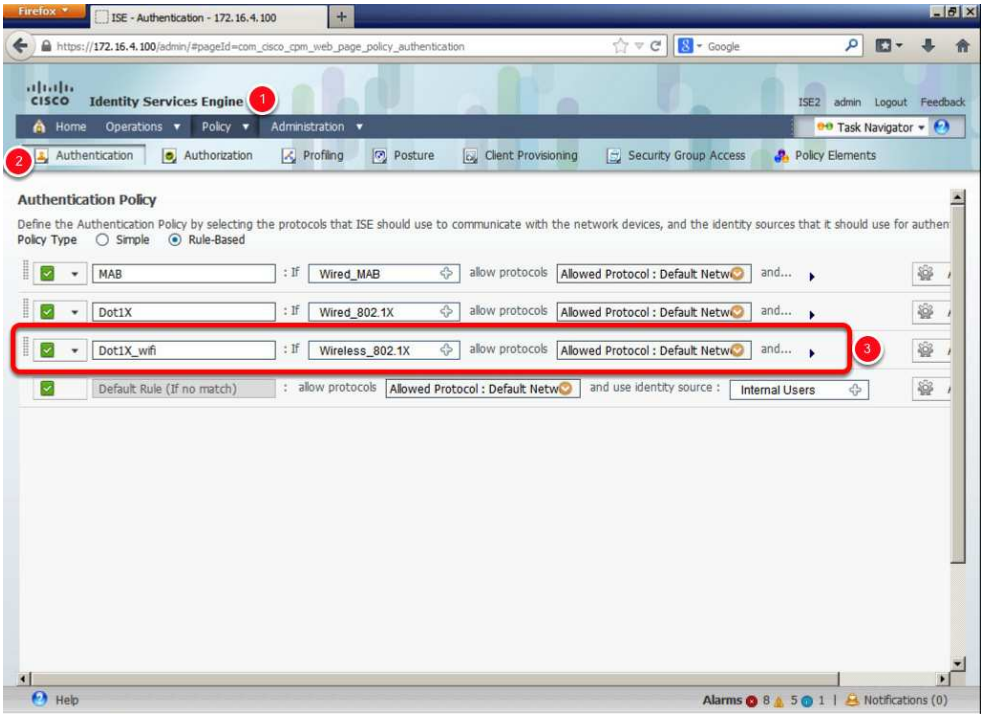


Configure a user in the ISE internal store, which will be used for PEAP authentication:

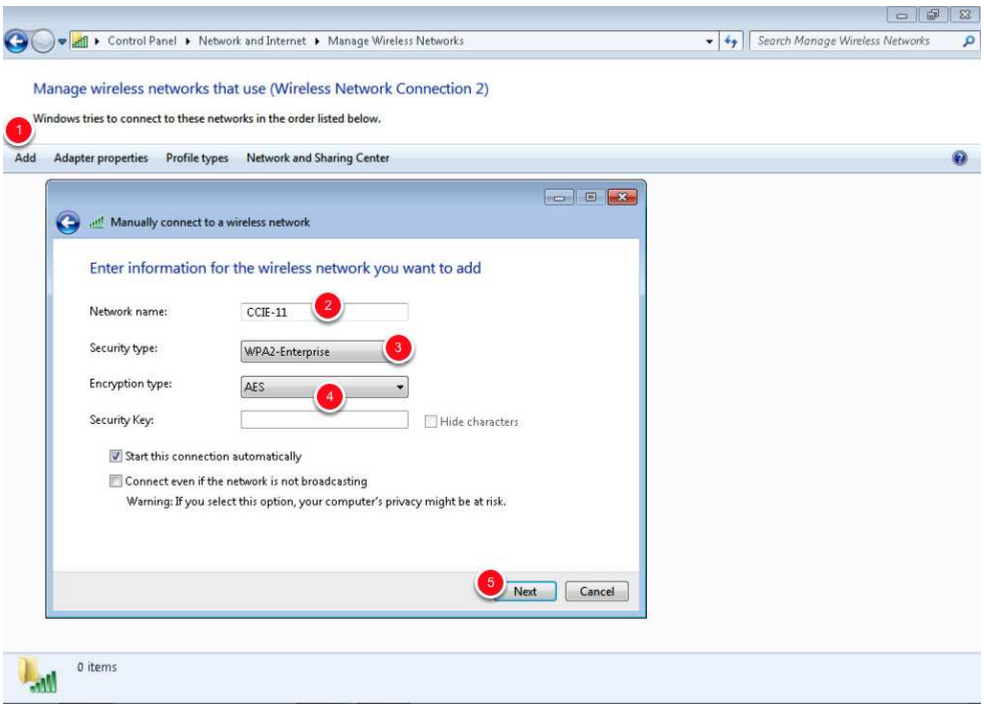


The default 802.1x authentication rule allows PEAP only for wired, so we must configure a similar rule for wireless:

https://t.me/learningnets

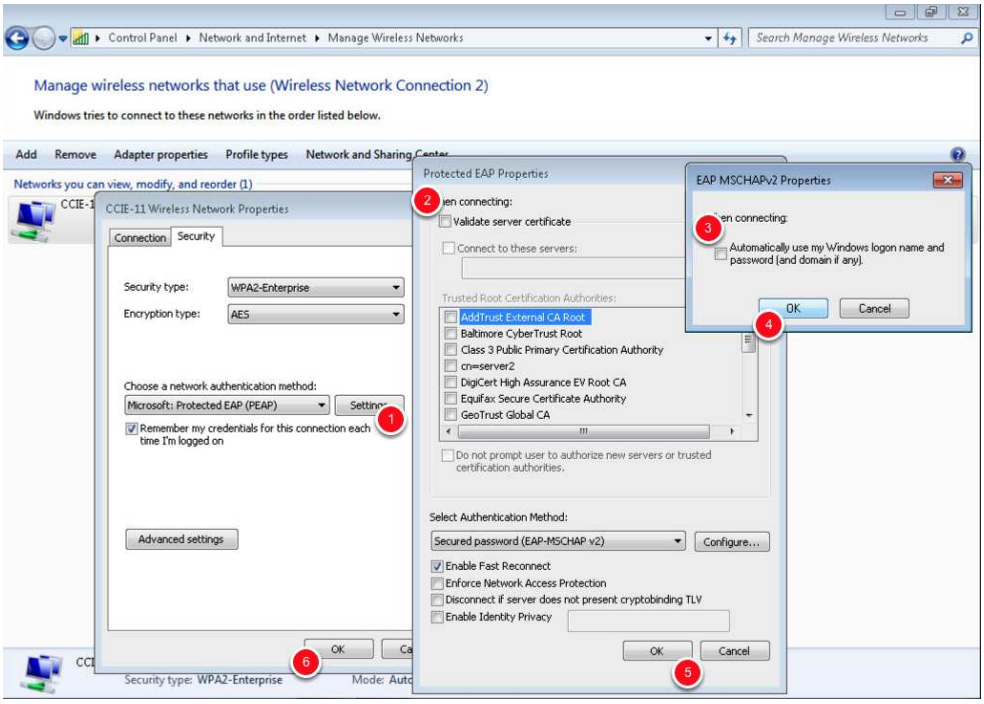


Configure an authorization policy specific for WiFi clients, or just use the default policy; here we make use of the default one. Go to the WiFi TEST-PC and manually add the WiFi network:



Edit the configured network; make sure to disable the **Validate Server Certificate** option (because ISE is using a self-signed certificate), and disable using the Windows credentials:

https://t.me/learningnets



You'll be asked for credentials to authenticate; use the ones configured on the ISE RADIUS server.

SW5:

```
vlan 147
 name WIFI-PC
!
interface Vlan 147
 ip address 136.1.147.13 255.255.255.0
!
ip dhcp excluded-address 136.1.147.1 136.1.147.1
ip dhcp pool VLAN147
 network 136.1.147.0 255.255.255.0
 default-router 136.1.147.13
!
interface FastEthernet1/0/11
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 47
 switchport trunk allowed vlan 47,147
 switchport mode trunk
```

Verification

Verify the RADIUS server configuration:

https://t.me/learningnets

```
(Cisco Controller) >show radius summary


Vendor Id Backward Compatibility..... Disabled
Call Station Id Case..... lower
Call Station Id Type..... Mac Address
Aggressive Failover..... Enabled
Keywrap..... Disabled
Fallback Test:
    Test Mode..... Off
    Probe User Name..... cisco-probe
    Interval (in seconds)..... 300
MAC Delimiter for Authentication Messages..... hyphen
MAC Delimiter for Accounting Messages..... hyphen

Authentication Servers

Idx  Type  Server Address    Port    State     Tout  MgmtTout  RFC3576  IPSec - AuthMode/Phase1/Group/Lifetime/Auth/
---  ----  ---------------  ------  --------  ----  --------  -------  ----------------------------------------
1    NM    172.16.4.100     1812    Enabled   2     2         Enabled   Disabled - none/unknown/group-0/0 none/none


Accounting Servers

Idx  Type  Server Address    Port    State     Tout  MgmtTout  RFC3576  IPSec - AuthMode/Phase1/Group/Lifetime/Auth/

--More-- or (q)uit
---  ----  ---------------  ------  --------  ----  --------  -------  ----------------------------------------
```

Verify that the client is associated and authenticated:

```
(Cisco Controller) >show client summary

Number of Clients..... 1

MAC Address       AP Name           Status        WLAN/RLAN      Auth Protocol         Port Wired
----------------- ----------------- ------------- -------------- ---- ---------------- ---- -----
ac:f1:df:13:bd:8f AP0006.f63b.9ada  Associated    1               Yes  802.11n(2.4 GHz) 1    N/A
```

Verify the ACL settings:

```
(Cisco Controller) >show flexconnect acl summary

ACL Name                          Status
```

https://t.me/learningnets

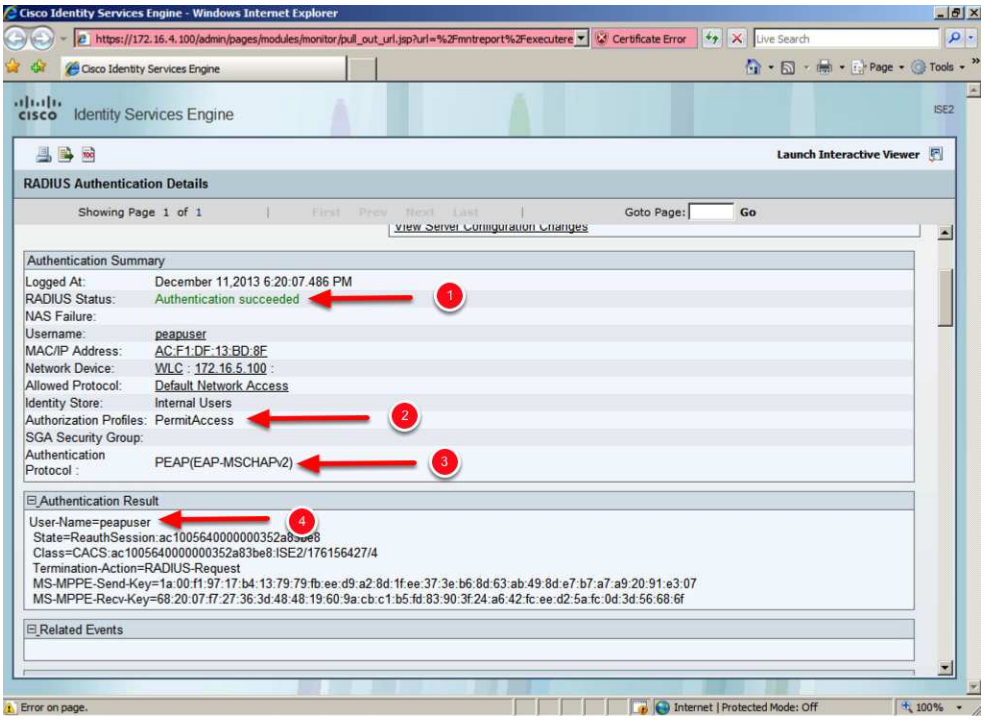
```
--------------------------------- -------FLEX-ACL                          Applied


!
!(Cisco Controller) >show flexconnect acl detailed FLEX-ACL

                    Source                          Destination                Source Port  Dest Port
Index        IP Address/Netmask              IP Address/Netmask        Prot     Range        Range     DSCP  Action
------ ------------------------------ ------------------------------ ---- ----------- ----------- ----- -------

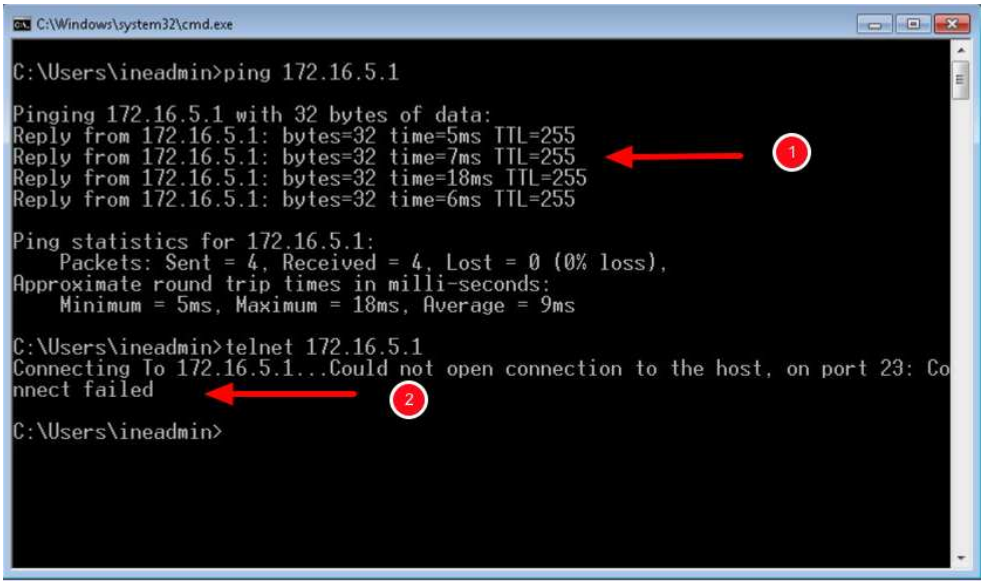
1          0.0.0.0/0.0.0.0                 0.0.0.0/0.0.0.0             1     0-65535     0-65535  Any Permit
```

Verify the RADIUS authentication logs on ISE, and check that correct authorization was applied:



Verify that WiFI TEST-PC is allowed ICMP traffic, but for example telnet is restricted:

https://t.me/learningnets



https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC Local Web Authentication

Task

- Configure WLC with a SSID named **CCIE-XX**, where **X** equals the rack number you've been assigned.
  - Authenticate users through the internal WLC portal, using the LDAP server at 172.16.20.100, and fallback to the local username database.
  - Query the LDAP server by username **ldapuser** with a password of **1qaz@WSX**, member of **Users** organization unit in the **inelab.local** domain.
  - Test client authentication by LDAP username of **USER1** with a password of **1qaz@WSX**, or by local username of **localwebuser** with a password of **localwebpass**.
  - Clients associated to the same SSID should not be allowed to communicate.
- Enable the WLC internal DHCP server, assign IP addresses from range 172.16.5.200-172.16.5.254, and assign domain **inelab.local** and DNS server of 172.16.20.100.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

https://t.me/learningnets

Overview

Web authentication is a layer3 method that cannot be combined with 802.1x layer2 security (unless running WLC version 7.4), but it can be combined with all other layer2 security methods, such as WPA/WPA2 with PSK. Although most implementations have either layer2 or layer3 security configured but not both, combining these together offers the advantage of using a web-portal for user authentication and user traffic encryption at the same time.

WLC will only allow DHCP and DNS; all other traffic is dropped before the client is authenticated in the captive portal (optionally, a preauthentication ACL can be applied to give access to additional resources before authentication). Depending on the operating system your client runs, it may be required to manually start an HTTP/HTTPS session that will be intercepted by WLC; if you access URLs by FQDN, name resolution must be functional first. The captive portal can be local (the internal WLC one), or external (such as the built-in one from ISE). For the built-in portal, WLC can authenticate users through the local user database, through RADIUS or LDAP.

The layer3 web authentication functionality can be implemented in several additional ways, depending on what needs to be achieved:

- **Passthrough**: Reinforces the captive portal where the user does not need to authenticate, but rather commit for an AUP (acceptable usage policy).
- **Conditional web redirect**: Conditionally redirects the user to a specific portal only after it has been authenticated through 802.1x or WPA/WPA2 layer2 security; the user does not have any other access but within the presented portal.
- **Splash page web redirect**: Redirects the user to a specific portal only after it has been authenticated through 802.1x or WPA/WPA2 layer2 security; the user does not need to perform any further authentication and is allowed access to the network.

WLC has a specific setting, called **P2P Blocking Action**, that controls whether two authenticated users of the same SSID are allowed to communicate:

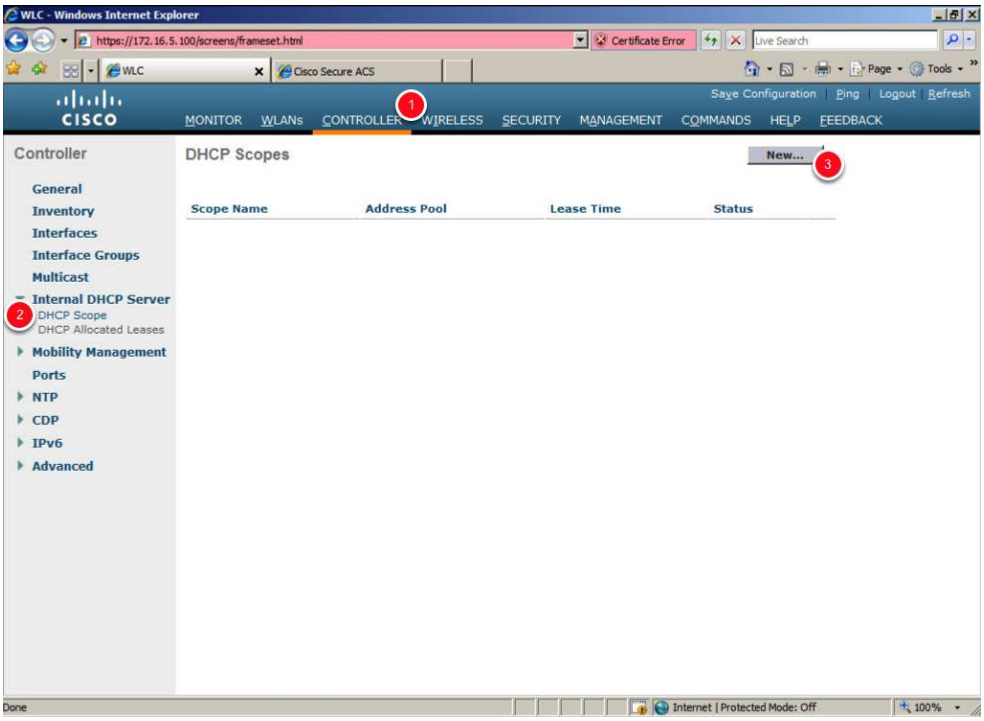
- **drop** option will restrict communication.
- **forward-upstream** will divert the traffic out of the wireless infrastructure to the directly connected switch of the WLC, where traffic can be controlled through a VACL.

Configuration

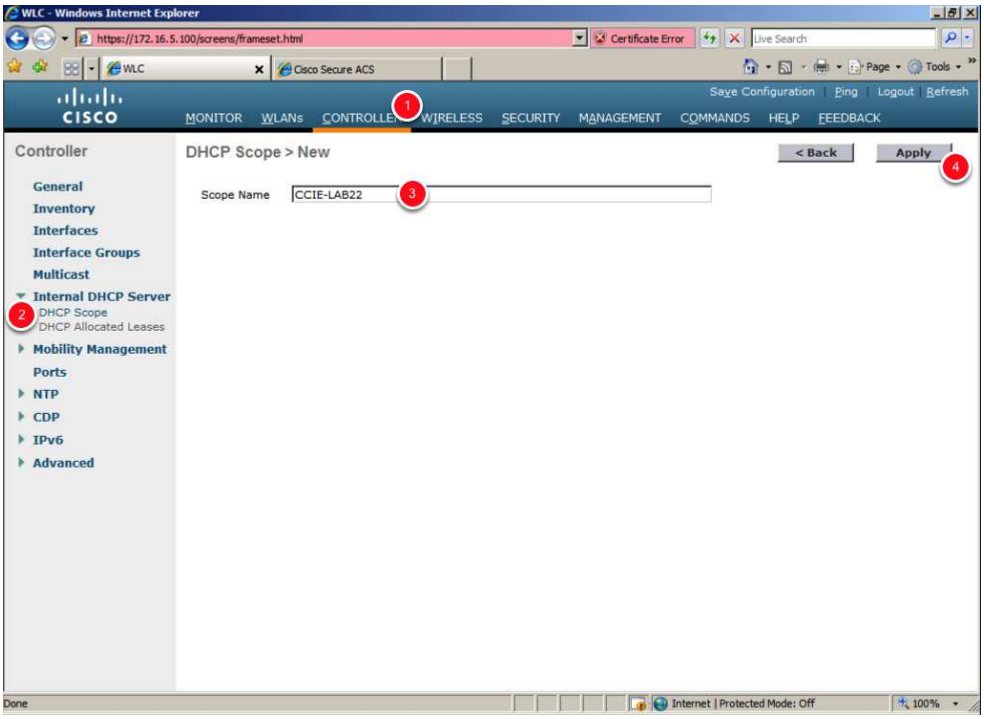
Navigate through the WLC GUI to create a new DHCP scope for the WLC internal

https://t.me/learningnets

server:

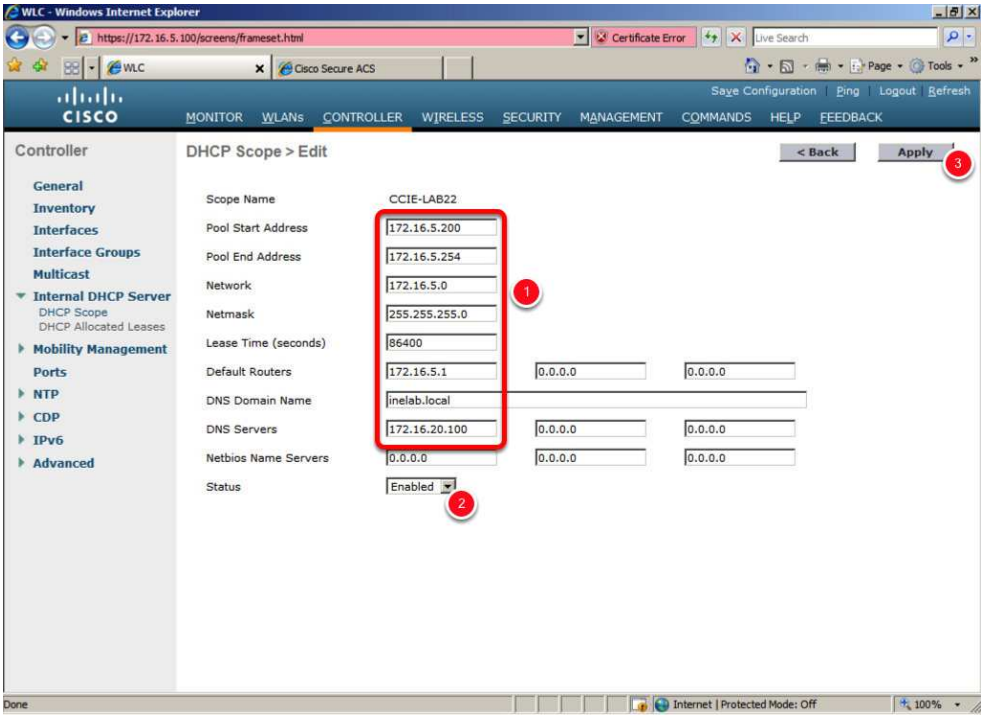


Give the DHCP scope a name:

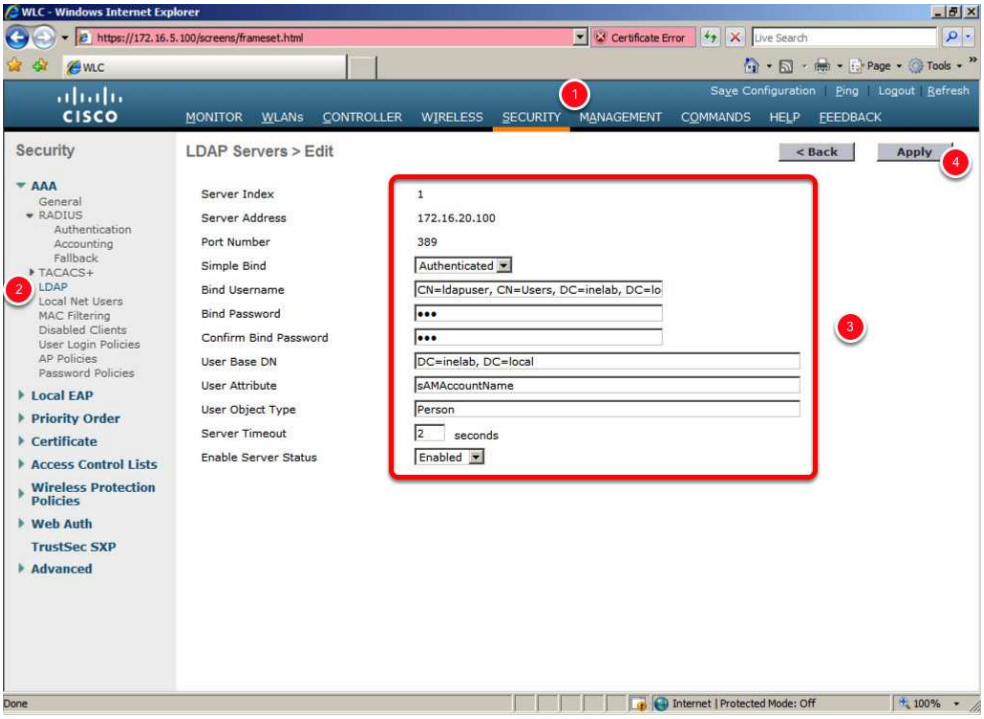


Configure the DHCP settings and enable it; by default it is disabled:

https://t.me/learningnets

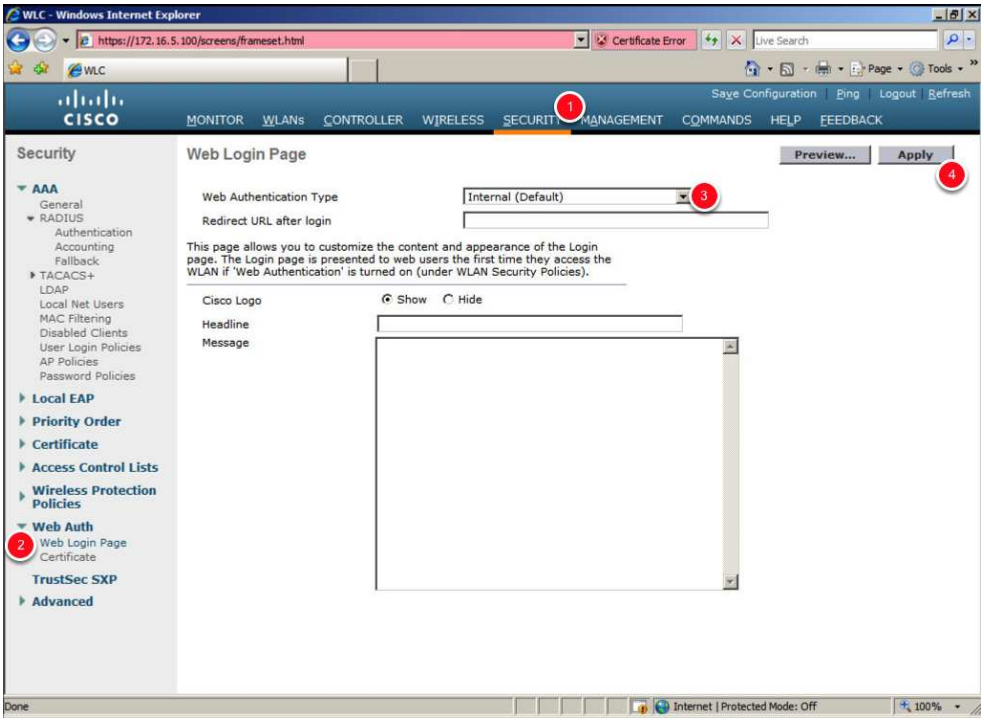


Navigate through the WLC GUI to configure an LDAP server:

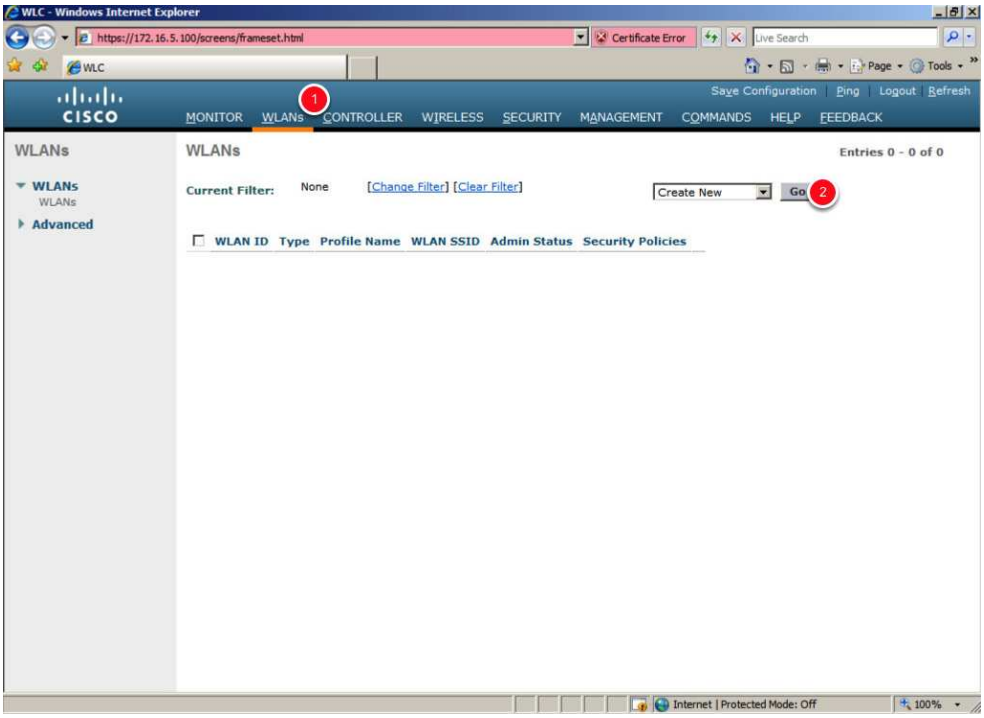


Ensure that the WLC web authentication uses the default settings, displaying the internal webpage:

https://t.me/learningnets

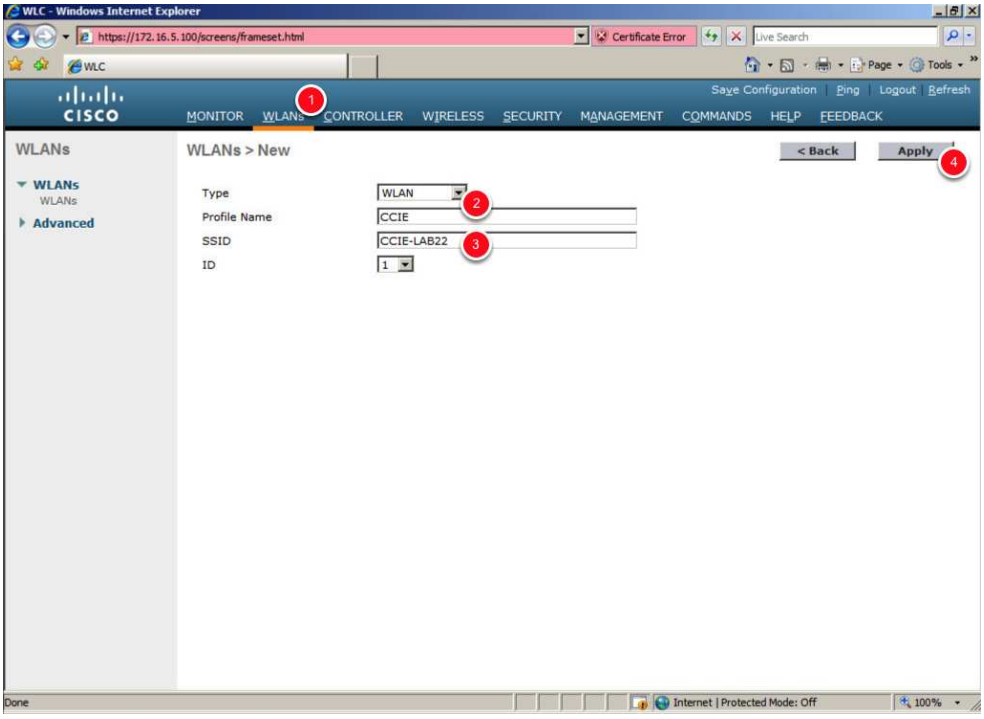


Navigate through WLC GUI to create a new WLAN:

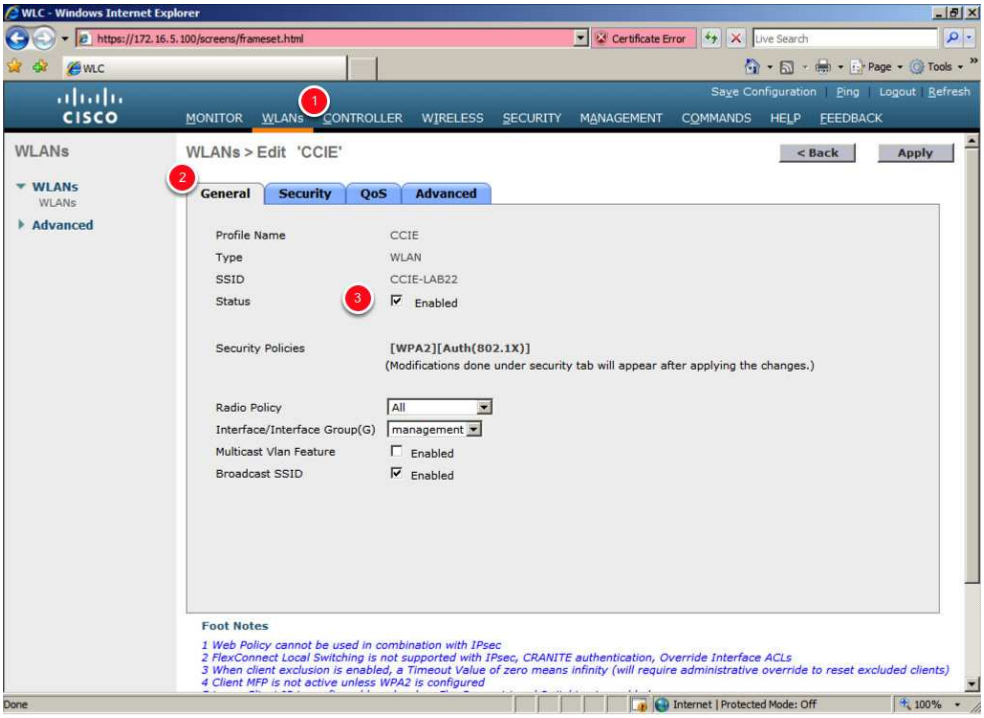


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

https://t.me/learningnets

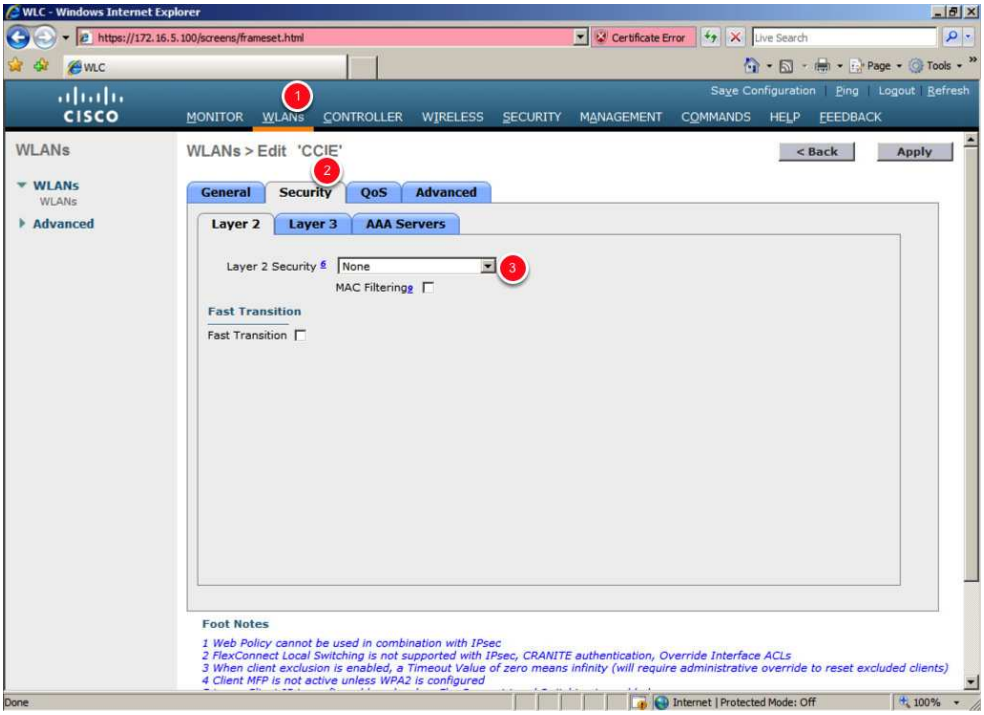


Enable the WLAN and leave WLC to broadcast the SSID:

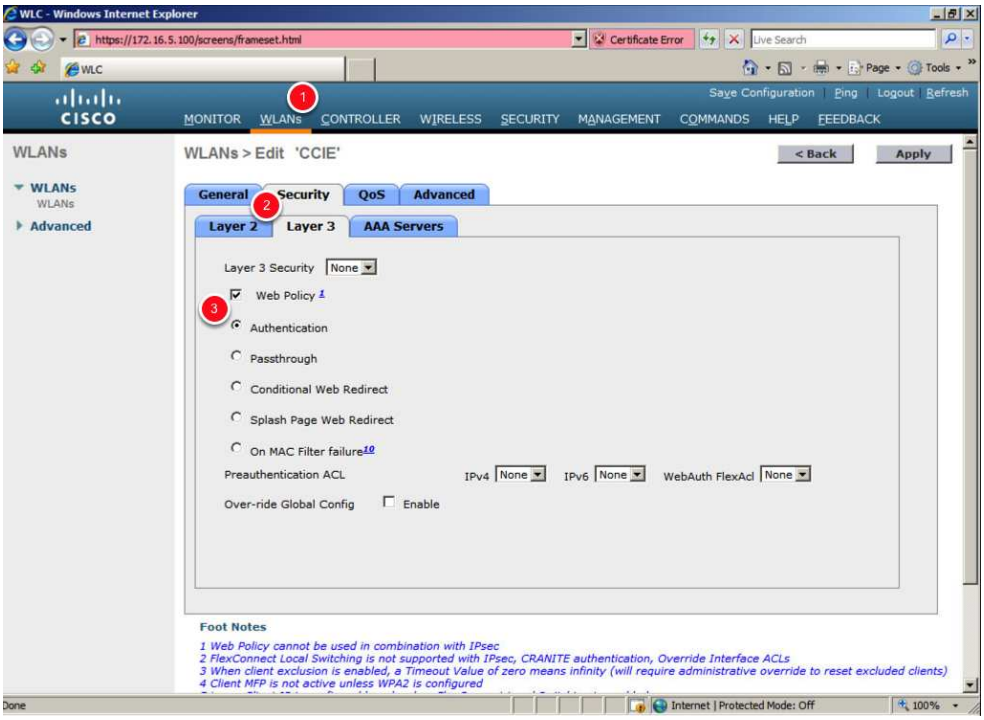


Disable layer2 authentication (set it to none) for the WLAN:

https://t.me/learningnets

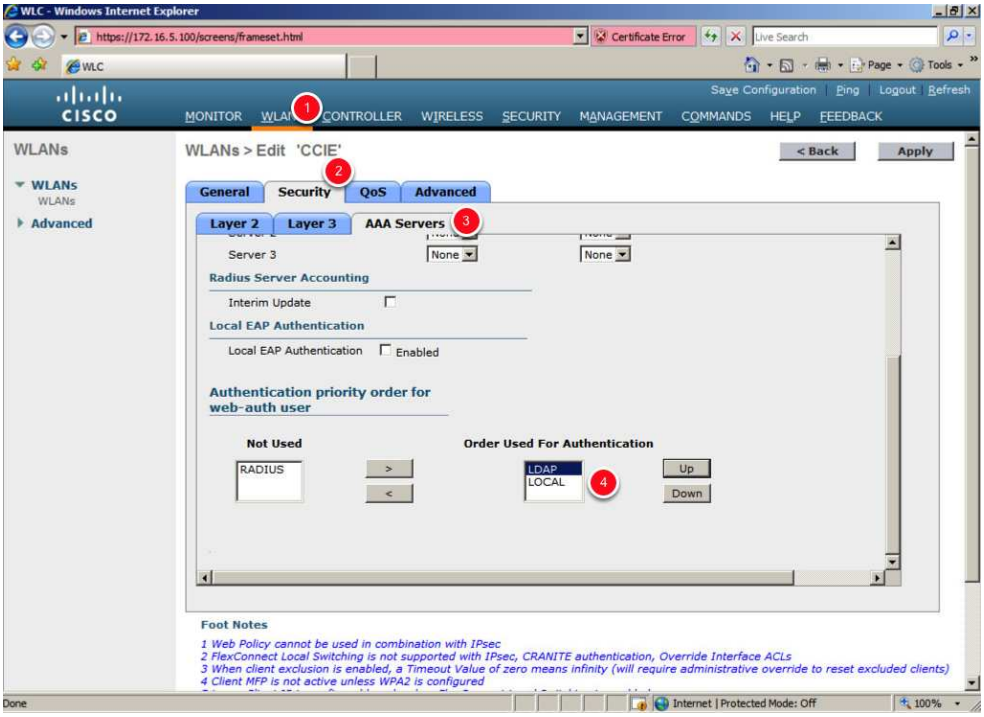


Enable layer3 authentication for the WLAN, using the **authentication** method:

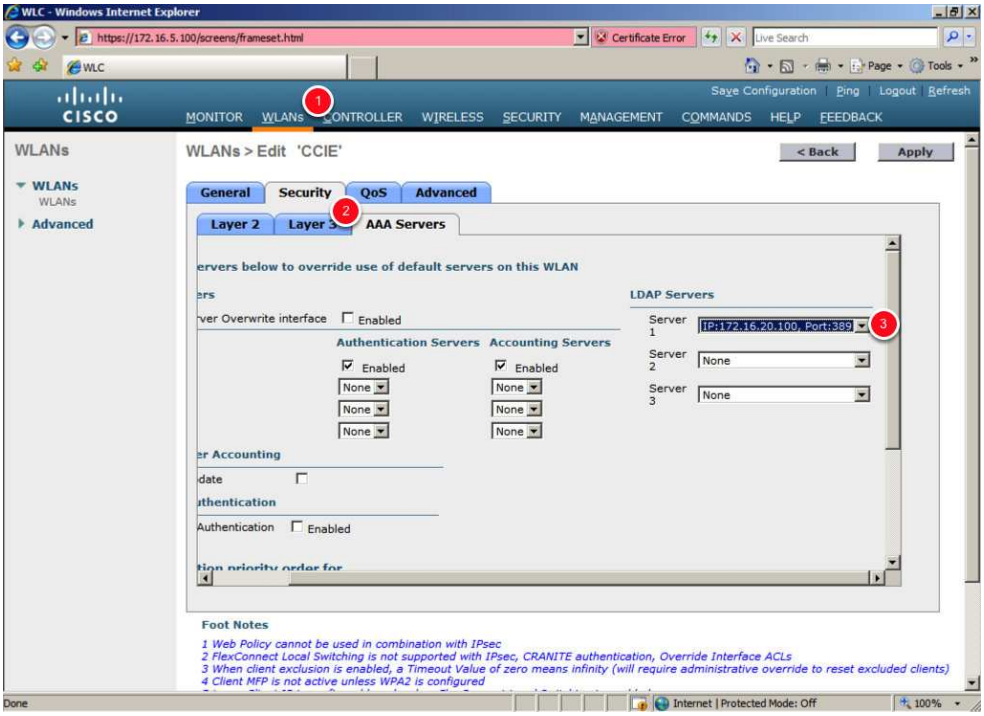


Configure the WLAN so that for web authentication it first uses LDAP and then uses the local user database:

https://t.me/learningnets

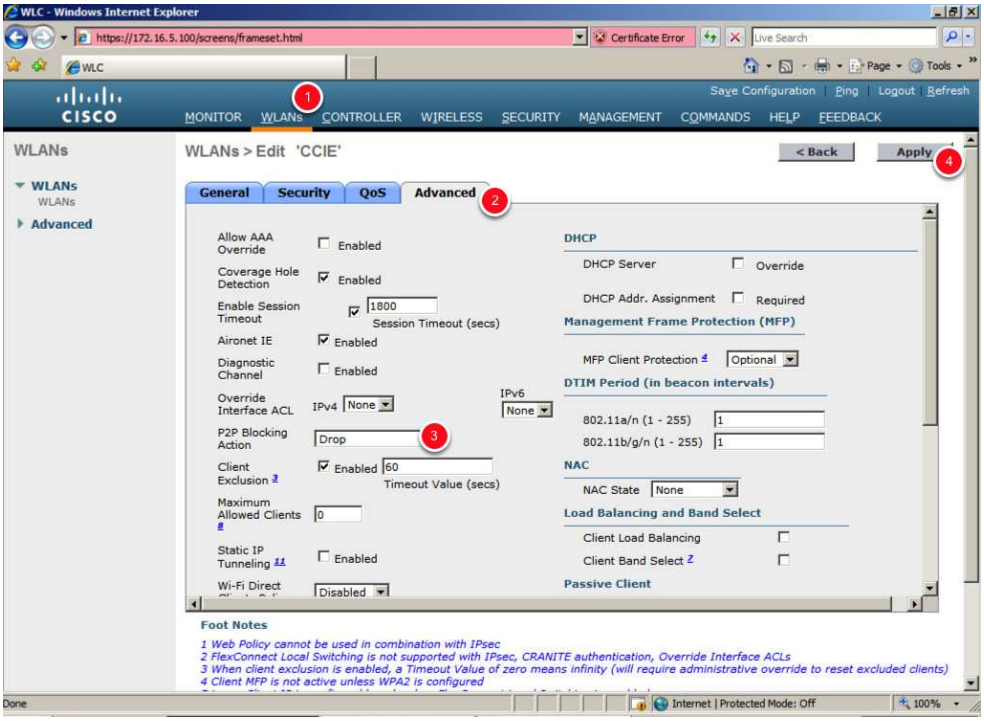


Specify which LDAP servers are used for the WLAN:

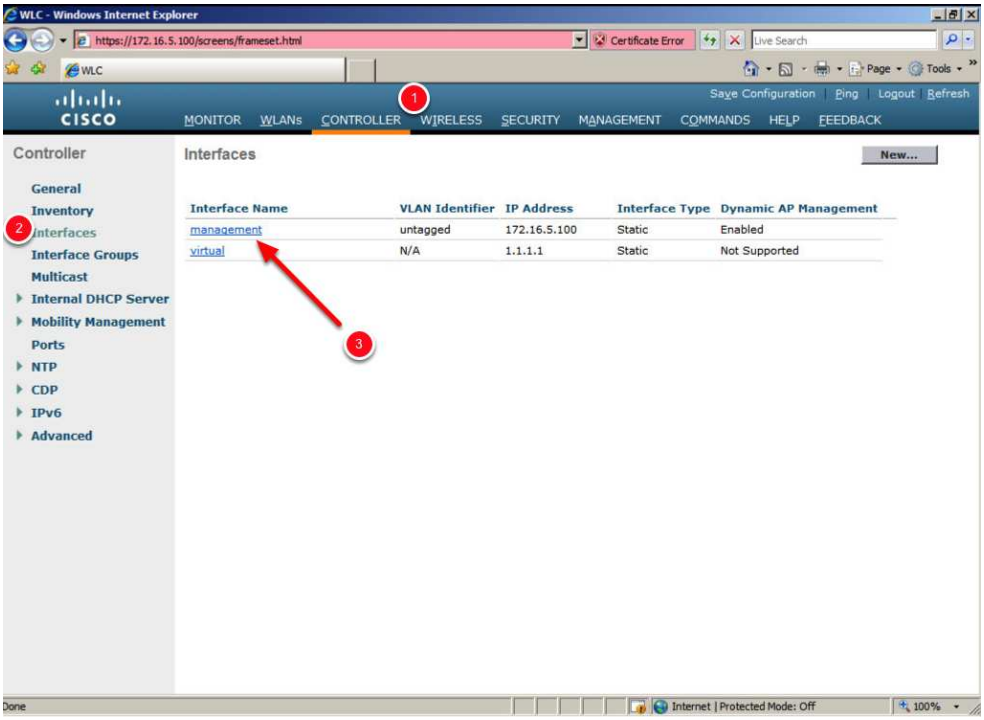


Configure the WLAN so that two clients associated to same access-point are not allowed to communicate:

https://t.me/learningnets

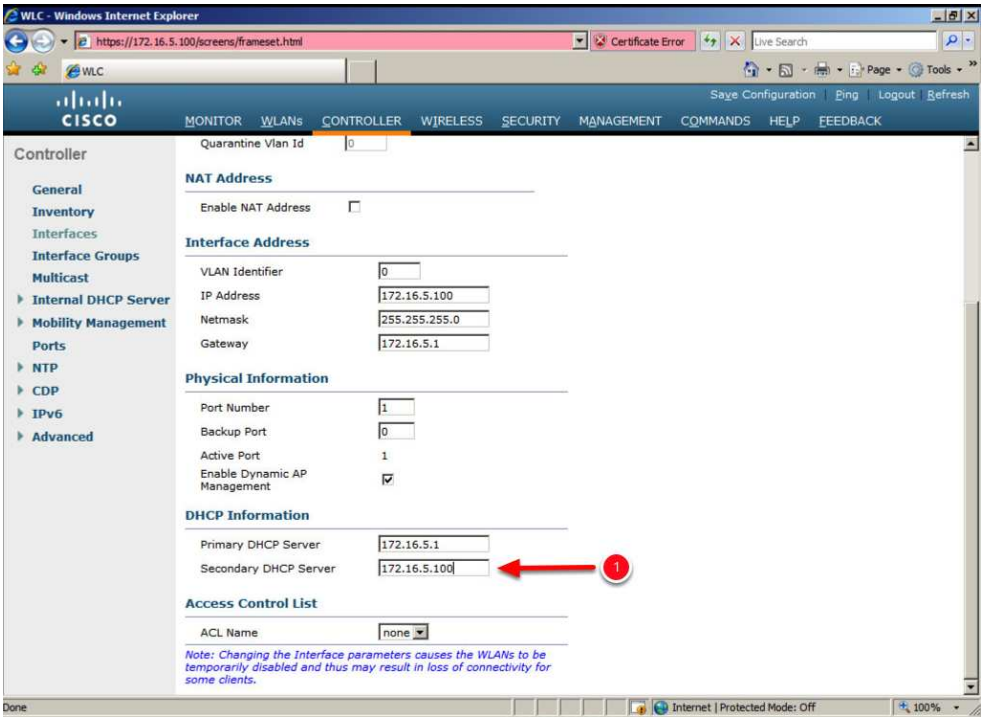


Navigate through the GUI WLC to configure WLC as DHCP server, at the interface level:

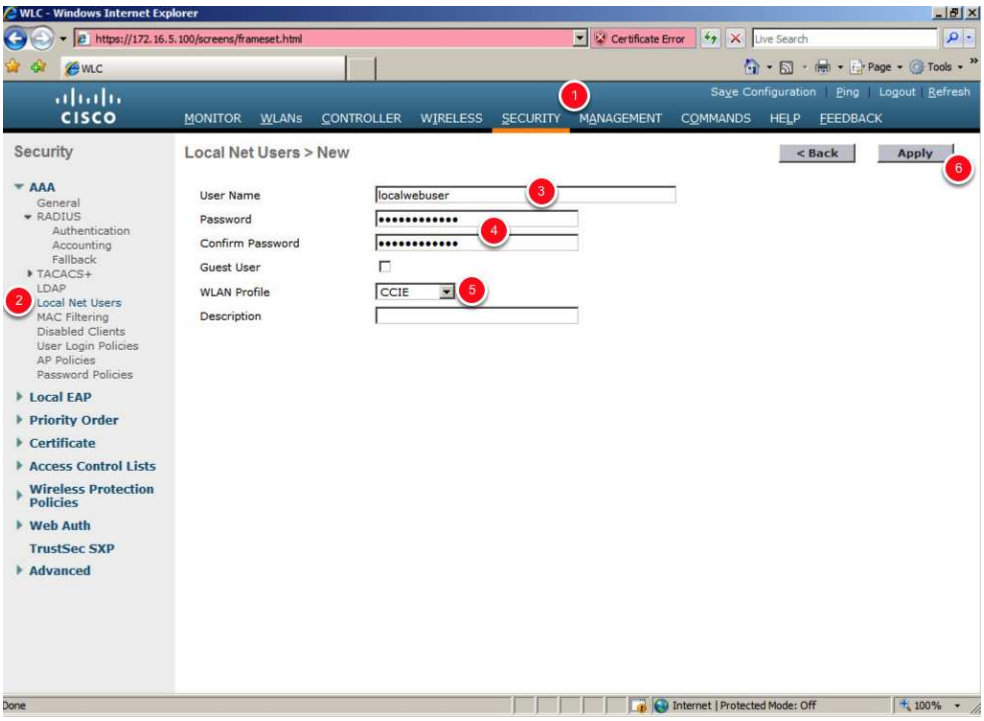


Configure the internal WLC DHCP server as the secondary DHCP server for WLANs, using the management interface:

https://t.me/learningnets



Add a user in the WLC local database, which will be used for only for **CCIE-LABXX** WLAN authentication purposes:



Verification

Verify the LDAP configuration:

```
(Cisco Controller) >show ldap 1
```

https://t.me/learningnets

```
Server Index..................................... 1
Address.......................................... 172.16.20.100
Port............................................. 389
Enabled.......................................... Yes
User DN.......................................... DC=inelab, DC=local
User Attribute................................... sAMAccountName
User Type........................................ Person
Retransmit Timeout............................... 2 seconds
Bind Method ..................................... Authenticated
Bind Username.................................... CN=ldapuser, CN=Users, DC=inelab, DC=local
```

Verify the local username database used only for WLAN authentication:

```
(Cisco Controller) >show netuser summary

Maximum logins allowed for a given user name..... Unlimited


User Name                 WLAN Id  User Type  Lifetime                          Description
------------------------  -------- ---------  ------------------------------    --------------------------------
localwebuser              WLAN   1 Permanent  N/A
```

After WiFi TEST-PC is associated, check its state before performing web authentication (the client requires an IP address to perform layer3 authentication):

```
(Cisco Controller) >show client summary

Number of Clients.............................. 1

MAC Address       AP Name           Status        WLAN/RLAN      Auth Protocol         Port Wired
----------------- ----------------- ------------- -------------- ---- ---------------- ---- -----
ac:f1:df:13:bd:8f AP0006.f63b.9ada  Associated    1              No   802.11n(2.4 GHz) 1    N/A
!
!(Cisco Controller) >show client detail ac:f1:df:13:bd:8f
Client MAC Address.............................. ac:f1:df:13:bd:8f
Client Username ................................ N/A
AP MAC Address.................................. f0:29:29:04:21:f0
AP Name......................................... AP0006.f63b.9ada
Client State.................................... Associated
Client NAC OOB State............................ Access
Wireless LAN Id................................. 1
802.11u......................................... Not Supported
BSSID........................................... f0:29:29:04:21:f0
Connected For .................................. 23 secs
```

https://t.me/learningnets

```
Channel..................................... 11
IP Address.................................. 172.16.5.200

IPv6 Address................................ fe80::3c16:4073:f3a4:c679
Association Id.............................. 1
Authentication Algorithm.................... Open System
Reason Code................................. 1
Status Code................................. 0
Session Timeout............................. 1800
Client CCX version.......................... 5
Client E2E version.......................... No E2E support
Diagnostics Capability...................... Supported
S69 Capability.............................. Not Supported
QoS Level................................... Silver
802.1P Priority Tag......................... disabled
CTS Security Group Tag...................... Not Applicable
KTS CAC Capability.......................... No
WMM Support................................. Enabled
  APSD ACs.................................  BK  BE  VI  VO
Power Save.................................. OFF
Supported Rates............................. 1.0,2.0,5.5,11.0,6.0,9.0,
    ........................................ 12.0,18.0,24.0,36.0,48.0,
    ........................................ 54.0
Mobility State.............................. Local
Mobility Move Count......................... 0
Security Policy Completed................... No
Policy Manager State........................ WEBAUTH_REQD

Policy Manager Rule Created................. Yes
AAA Override ACL Name....................... none
AAA Override ACL Applied Status............. Unavailable
AAA Override Flex ACL Name.................. none
AAA Override Flex ACL Applied Status........ Unavailable
AAA URL redirect............................ none
Audit Session ID............................ none
IPv4 ACL Name............................... none
IPv4 ACL Applied Status..................... Unavailable
```

Verify the DHCP server configuration and leases:

```
(Cisco Controller) >show dhcp detailed CCIE-LAB22
Scope: CCIE-LAB22

Enabled..................................... Yes
Lease Time.................................. 86400 (1 day )
```

https://t.me/learningnets

```
Pool Start..................................... 172.16.5.200
Pool End....................................... 172.16.5.254
Network........................................ 172.16.5.0
Netmask........................................ 255.255.255.0
Default Routers................................ 172.16.5.1  0.0.0.0  0.0.0.0
DNS Domain..................................... inelab.local
DNS............................................ 172.16.20.100  0.0.0.0  0.0.0.0
Netbios Name Servers........................... 0.0.0.0  0.0.0.0  0.0.0.0
!
!(Cisco Controller) >show dhcp leases

       MAC                IP          Lease Time Remaining
ac:f1:df:13:bd:8f    172.16.5.200     23 hours 51 minutes 13 seconds
```

Before authenticating the client, verify that you have no IP connectivity:

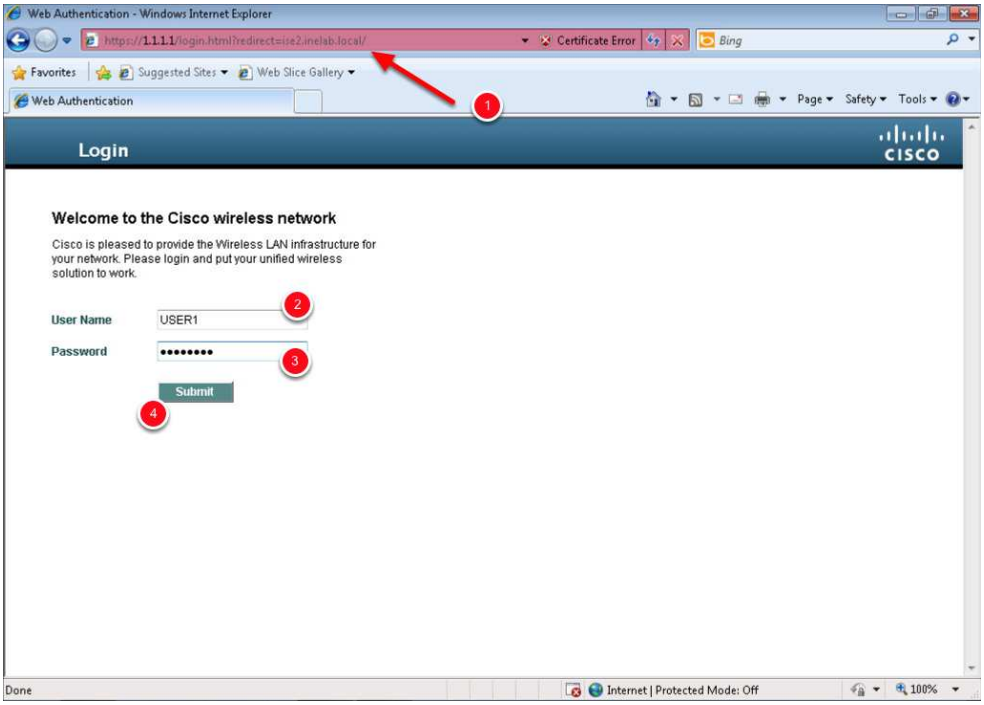
```
SW5#ping 172.16.5.200

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.5.200, timeout is 2 seconds:.....

Success rate is 0 percent (0/5)
```

From the WiFi TEST-PC, browse to an IP address or to an FQDN resolvable by the DNS server (TEST-SRV-B), such as ise2.inelab.local. WLC will intercept the HTTP request and ask for authentication; type the LDAP username/password:

https://t.me/learningnets



Authentication should be successful, and you'll be redirected the requested URL. Check that the client state changed after successful authentication:

```
(Cisco Controller) >show client detail ac:f1:df:13:bd:8f
Client MAC Address.............................. ac:f1:df:13:bd:8f
Client Username ................................ USER1
AP MAC Address.................................. f0:29:29:04:21:f0
AP Name......................................... AP0006.f63b.9ada
Client State.................................... Associated
Client NAC OOB State............................ Access
Wireless LAN Id................................. 1
802.11u......................................... Not Supported
BSSID........................................... f0:29:29:04:21:f0
Connected For .................................. 298 secs
Channel......................................... 11
IP Address...................................... 172.16.5.200
IPv6 Address.................................... fe80::3c16:4073:f3a4:c679
Association Id.................................. 1
Authentication Algorithm........................ Open System
Reason Code..................................... 1
Status Code..................................... 0
Session Timeout................................. 1800
Client CCX version.............................. 5
Client E2E version.............................. No E2E support
Diagnostics Capability.......................... Supported
S69 Capability.................................. Not Supported
QoS Level....................................... Silver
802.1P Priority Tag............................. disabled
CTS Security Group Tag.......................... Not Applicable
```

https://t.me/learningnets

```
KTS CAC Capability..............................  No
WMM Support.....................................  Enabled
  APSD ACs.....................................   BK  BE  VI  VO
Power Save......................................  OFF
Current Rate....................................  5.5
Supported Rates.................................  1.0,2.0,5.5,11.0,6.0,9.0,
    ..............................................  12.0,18.0,24.0,36.0,48.0,
    ..............................................  54.0
Mobility State..................................  Local
Mobility Move Count.............................  0
Security Policy Completed.......................  Yes
Policy Manager State............................  RUN

Policy Manager Rule Created.....................  Yes
AAA Override ACL Name...........................  none
AAA Override ACL Applied Status.................  Unavailable
AAA Override Flex ACL Name......................  none
AAA Override Flex ACL Applied Status............  Unavailable
AAA URL redirect................................  none
Audit Session ID................................  none
IPv4 ACL Name...................................  none
```

When the client is authenticated, verify that you have IP connectivity:

```
SW5#ping 172.16.5.200

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.5.200, timeout is 2 seconds:!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 50/176/596 ms
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC Central Web Authentication

Task

- Configure WLC with an SSID named **CCIE-XX**, where **X** equals the rack number you've been assigned.
  - Authenticate users through the built-in ISE portal, but do not manually configure any users in the ISE database.
  - Enable the DHCP server on WLC for this SSID, and assign IP addresses from range 172.16.5.200-172.16.5.254, domain **inelab.local**, and DNS server of 172.16.20.100.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If you have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, and for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (**Commands** -->**Download File**).

Overview

This is the same principle as outlined in the Identity Management section for CWA with ISE. The only difference is that now we allow the guest user to self-enroll, so we need to enable this functionality. One other difference, compared with the wired version of CWA, is on the redirect-ACL:

https://t.me/learningnets

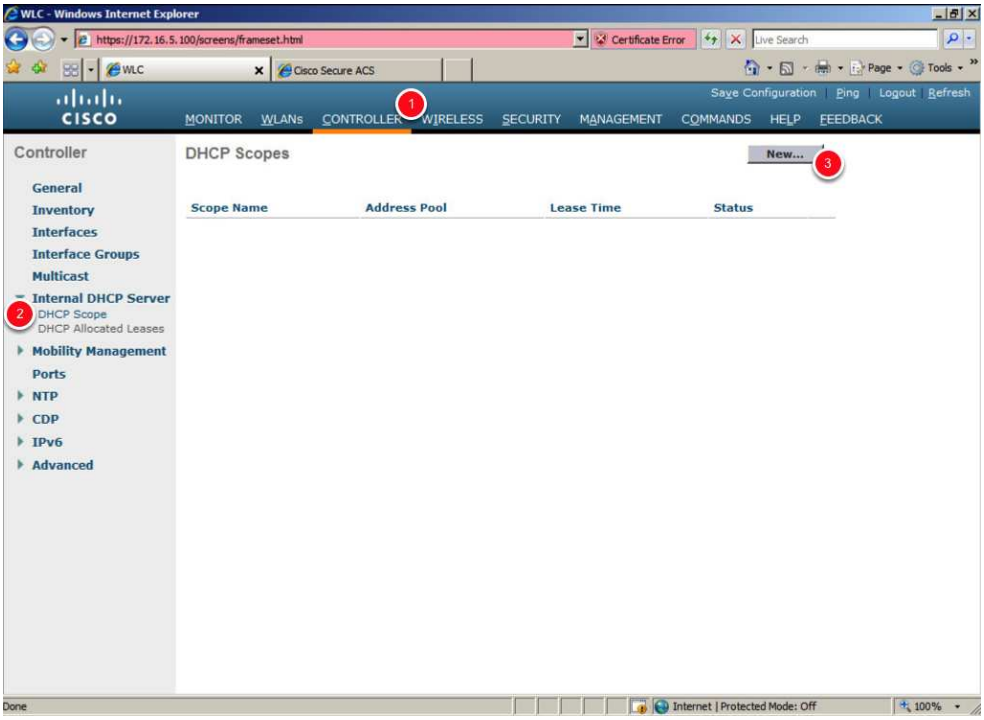
- All permit statements from the ACL define traffic that is not redirected (on wired, it was traffic that was redirected).
- All deny statements from the ACL define traffic that is redirected (on wired, it was traffic that was not redirected).

For CWA to work, make sure to satisfy these requirements:

- The client obtains an IP address (statically or manually).
- DNS and traffic toward ISE portal will not be redirected.
- The client can resolve through DNS the FQDN of the ISE server.

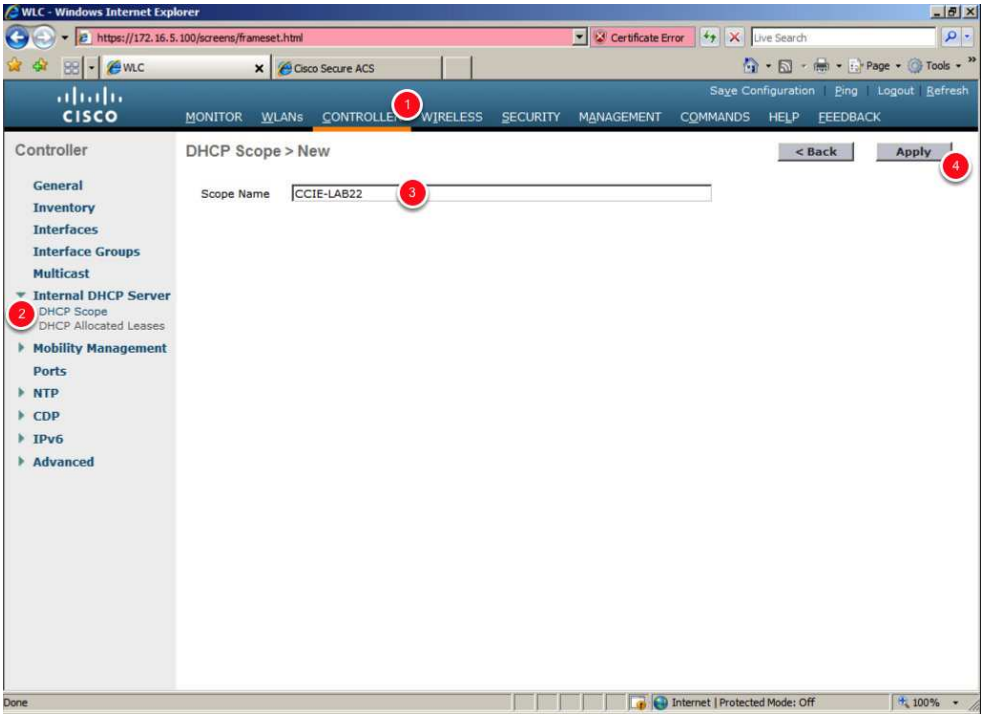
Configuration

Navigate through the WLC GUI to create a new DHCP scope for the WLC internal server:

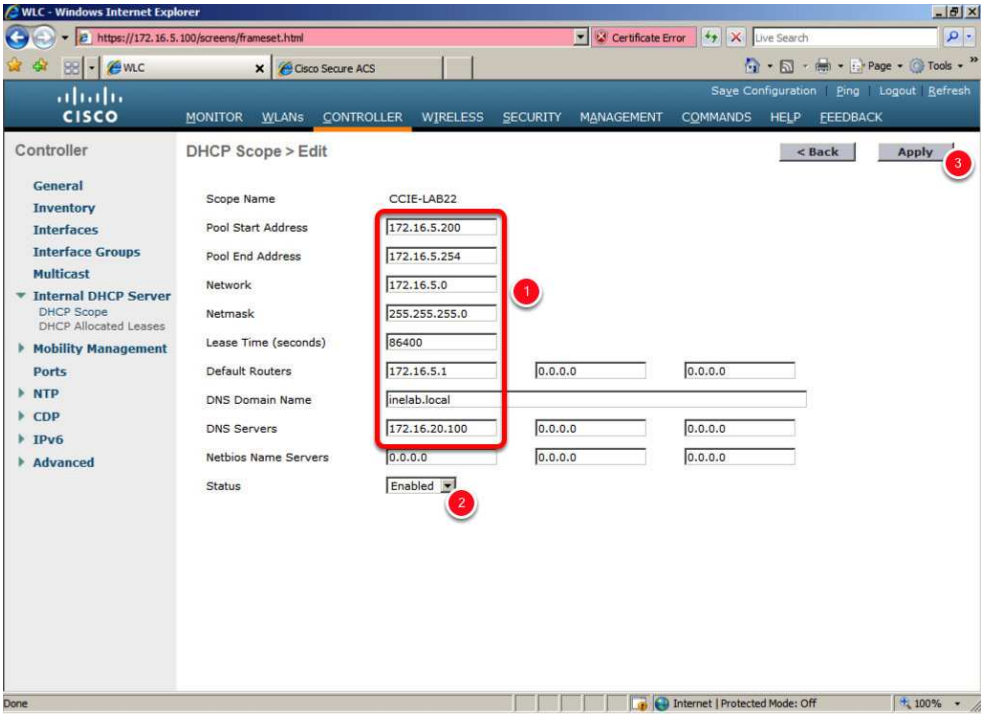


Give the DHCP scope a name:

https://t.me/learningnets

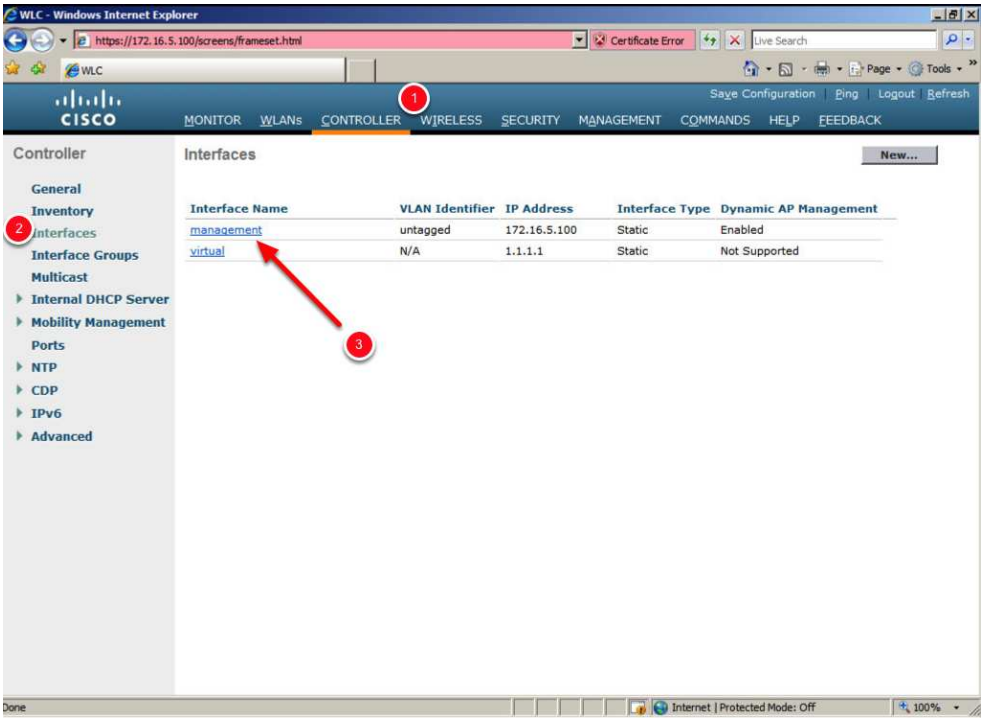


Configure the DHCP settings and enable it; by default it is disabled:

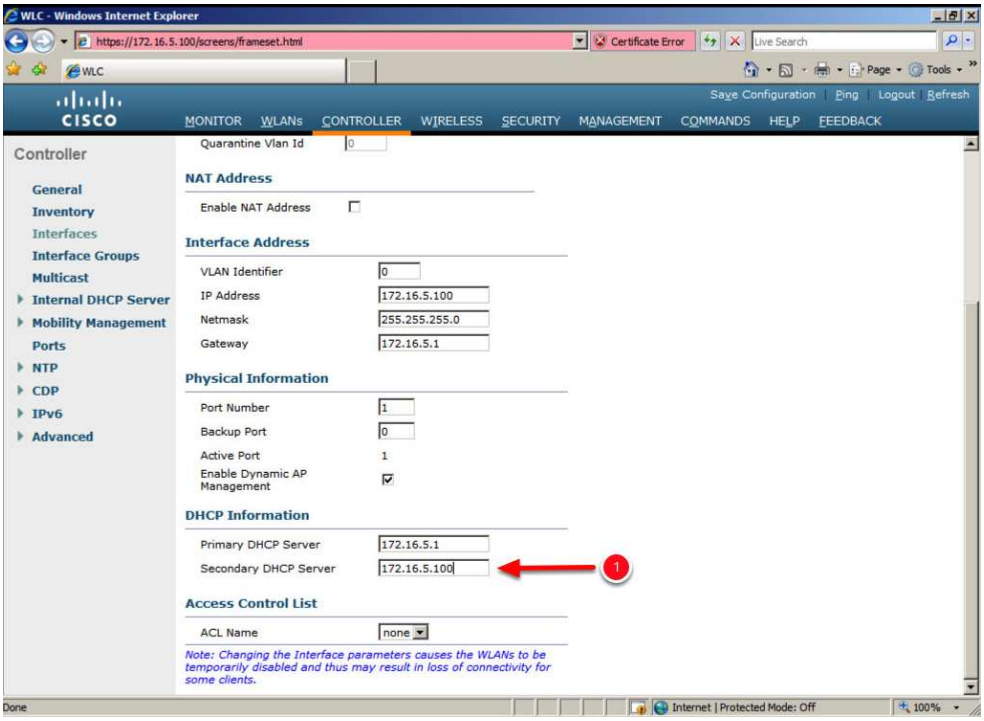


Navigate through the GUI WLC to configure WLC as the DHCP server, at the interface level:

https://t.me/learningnets

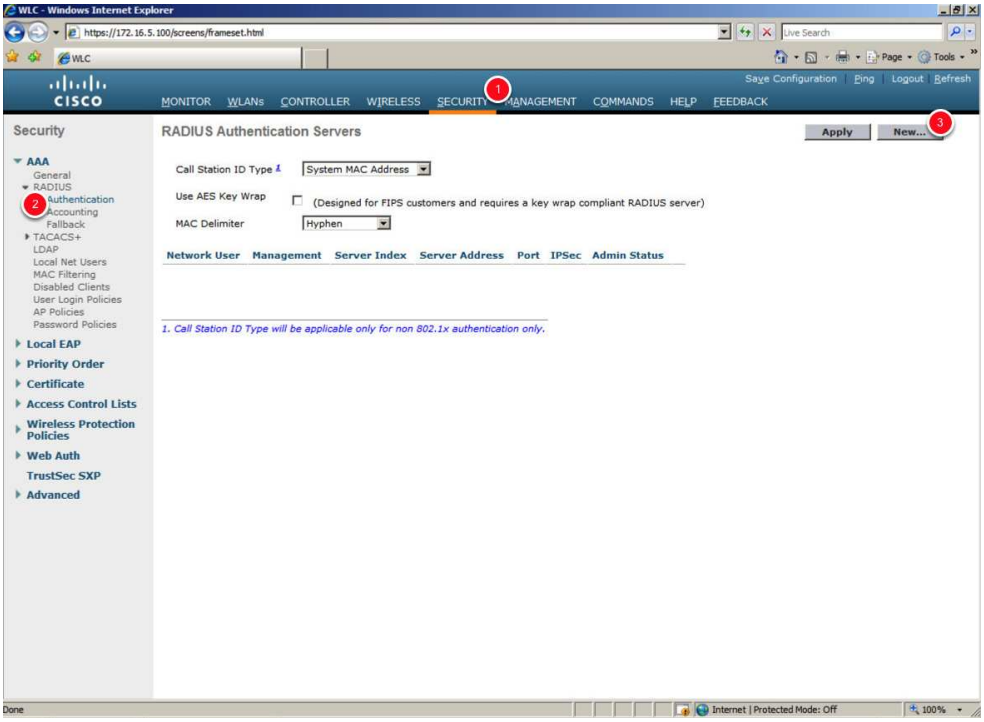


Configure the internal WLC DHCP server as the secondary DHCP server for WLANs using the management interface:

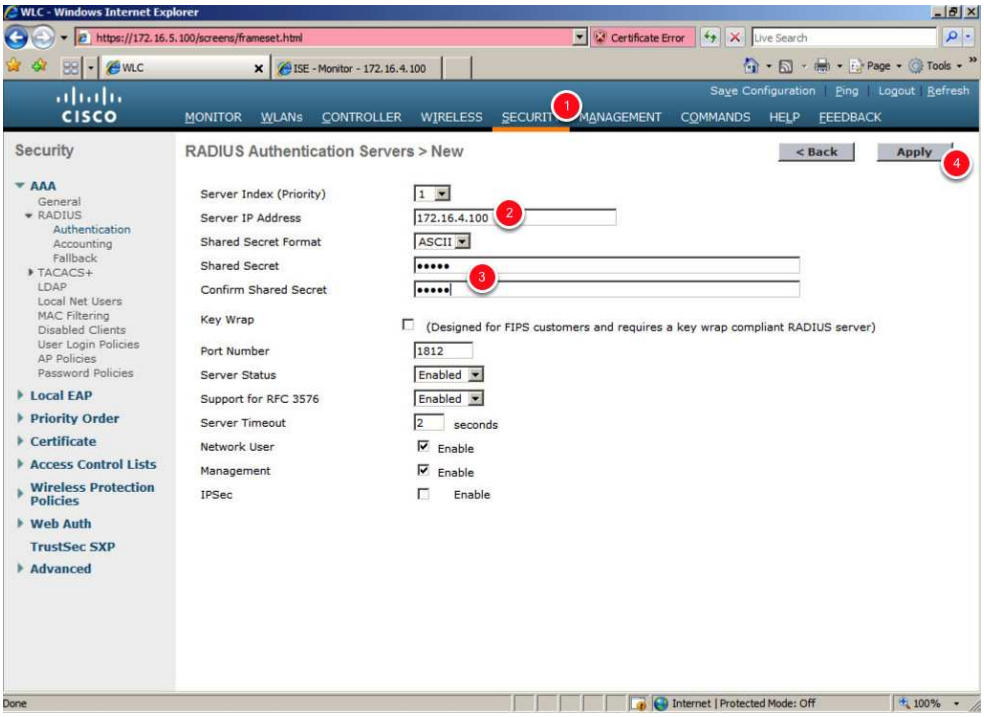


Navigate through the WLC GUI to add the ISE RADIUS server:

https://t.me/learningnets

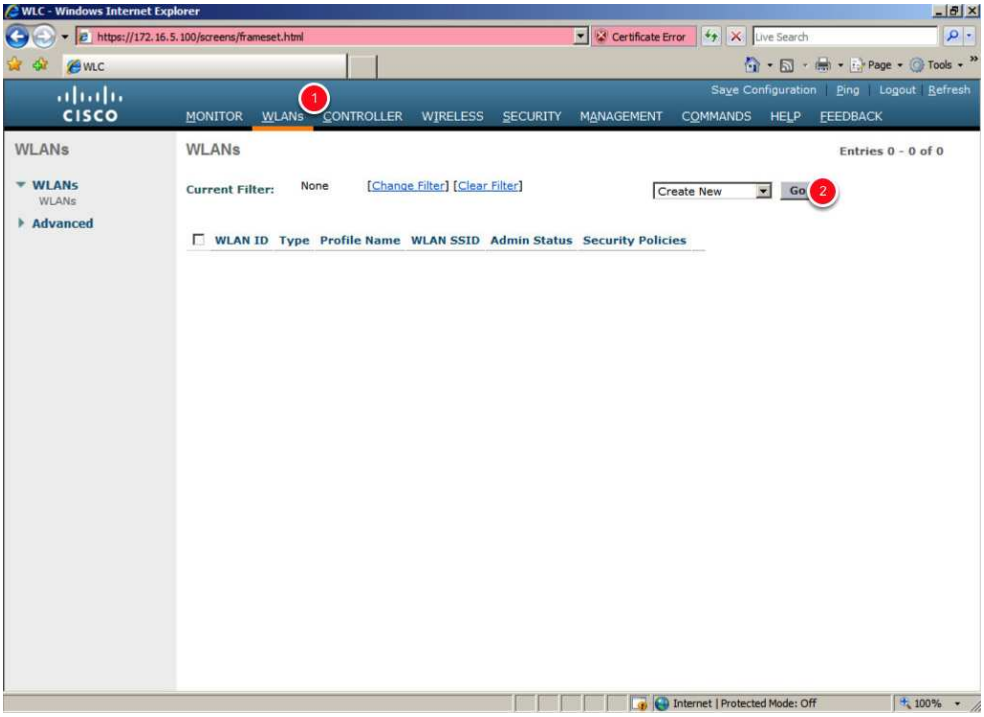


Add the RADIUS server, and ensure that **Support for RFC3576**, which is RADIUS CoA, is enabled:

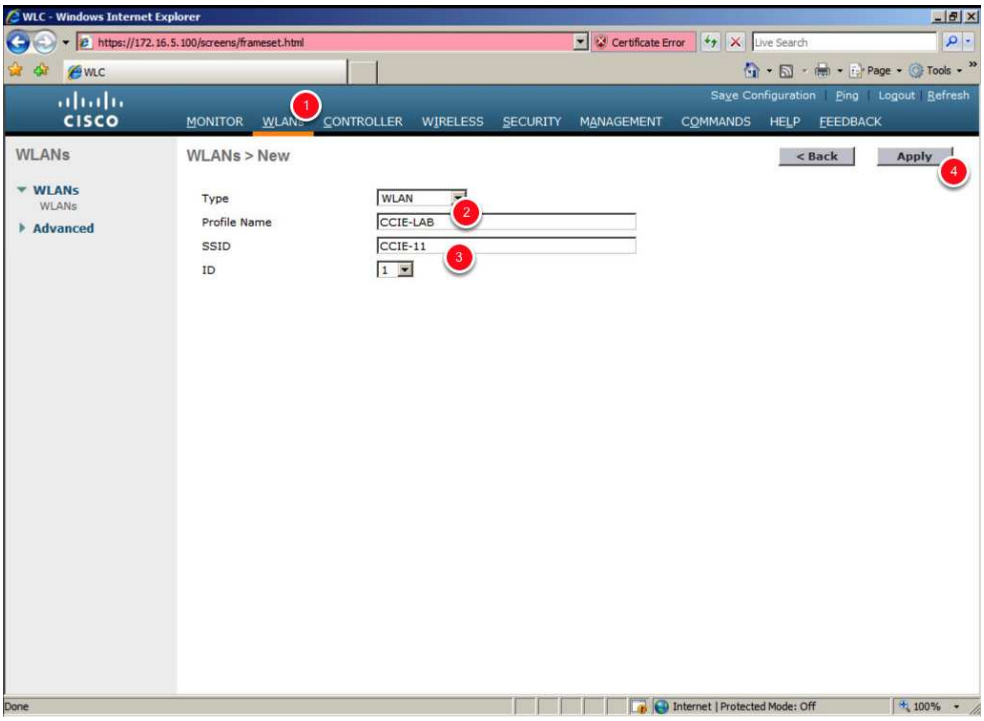


Navigate through WLC GUI to create a new WLAN:

https://t.me/learningnets

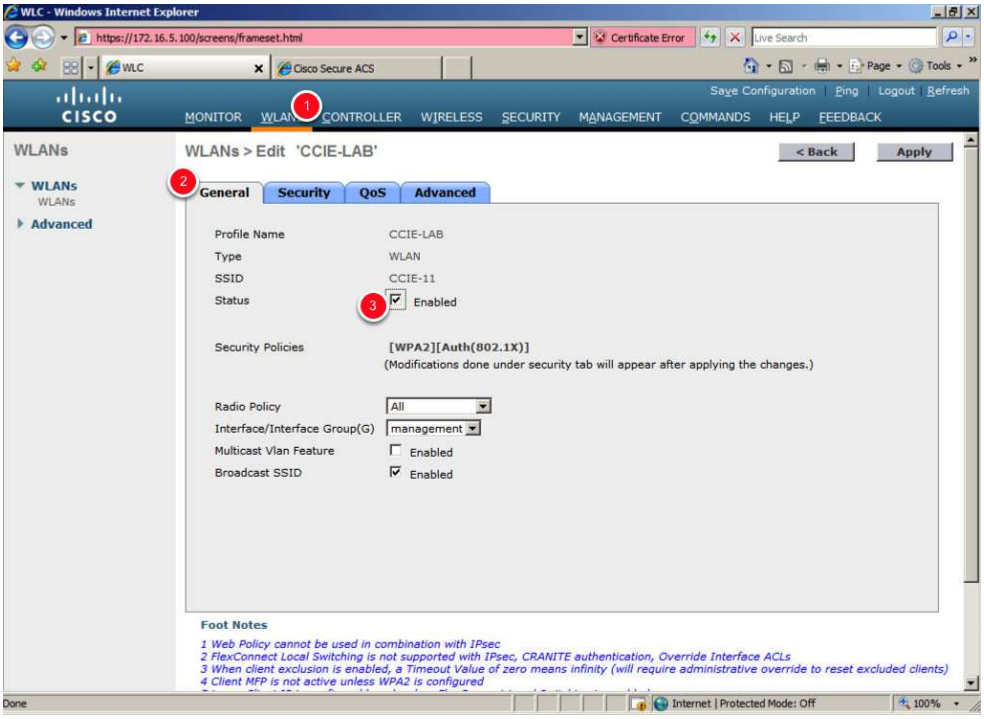


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

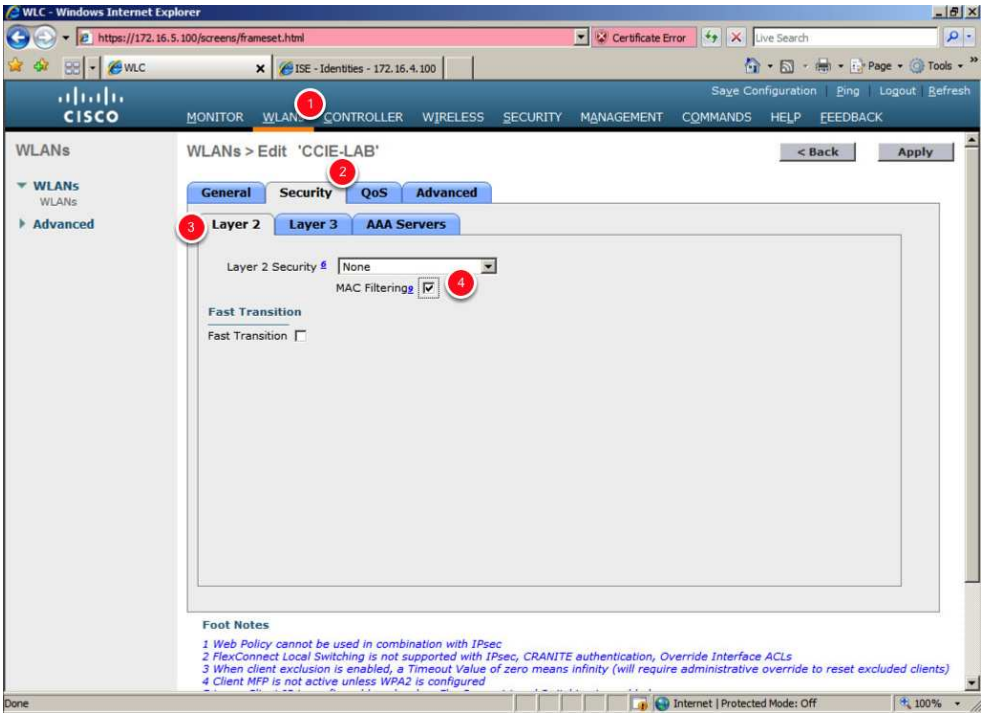


Enable the WLAN and leave it at default to be broadcasted:

https://t.me/learningnets

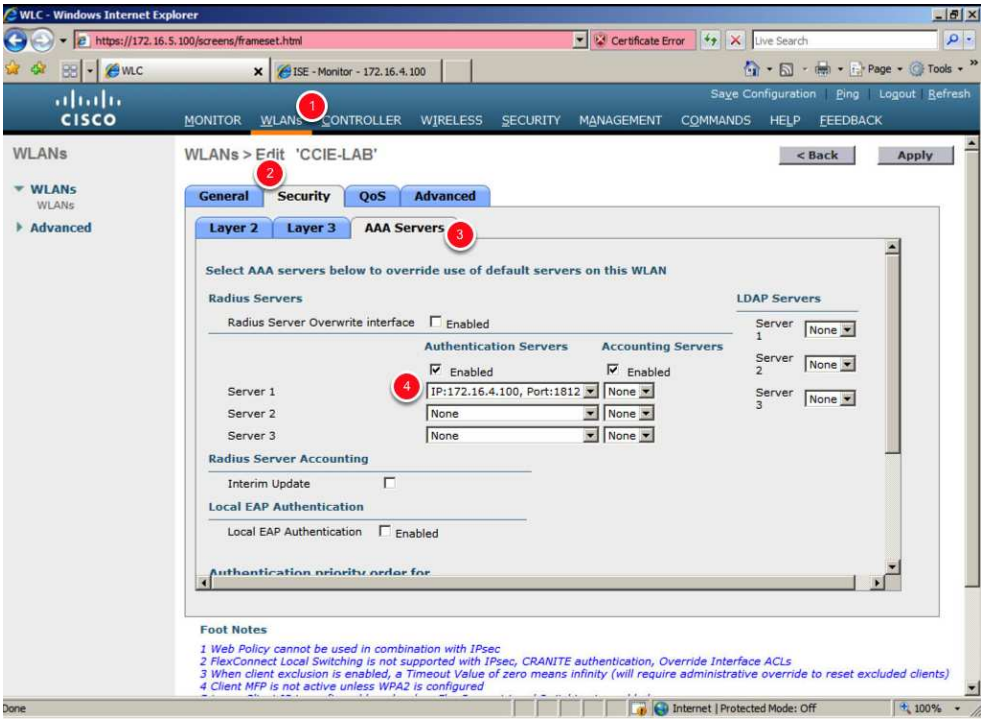


Configure the WLAN for layer2 security with just MAC filtering enabled:

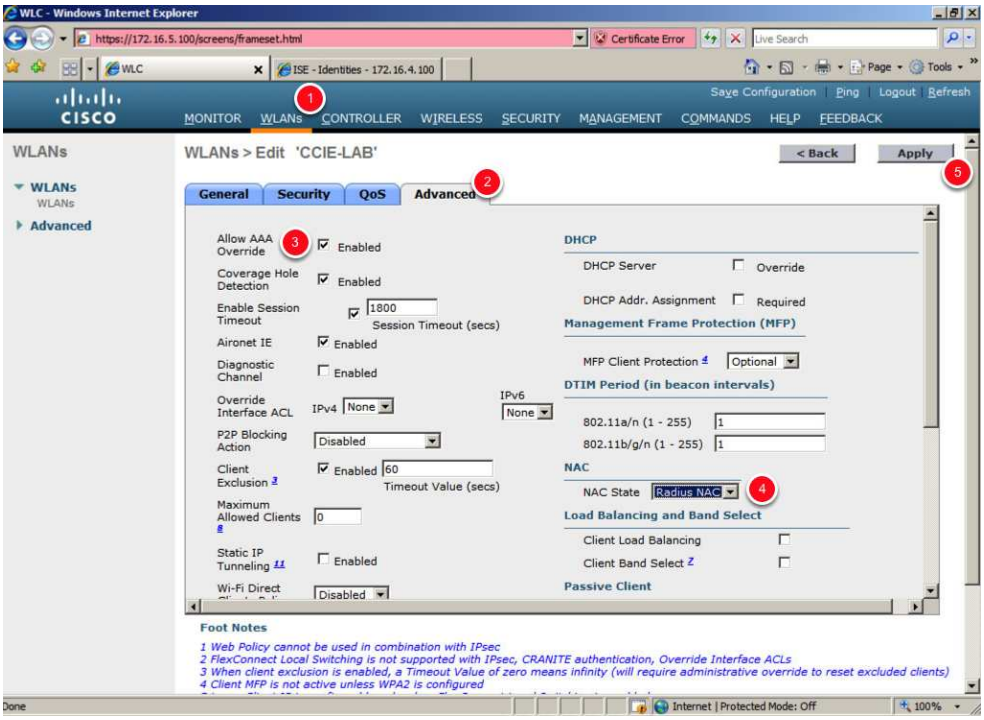


Specify the RADIUS server to be used for authentication:

https://t.me/learningnets

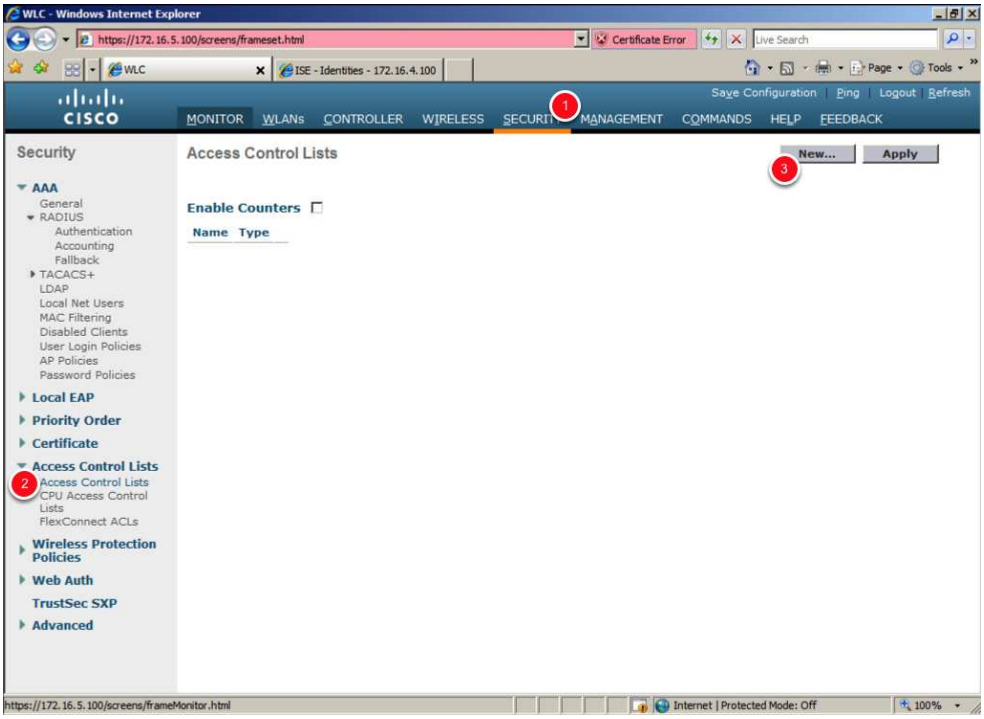


Enable **AAA Override** to apply authorization attributes received from ISE and enable NAC state (for CoA to be activated):

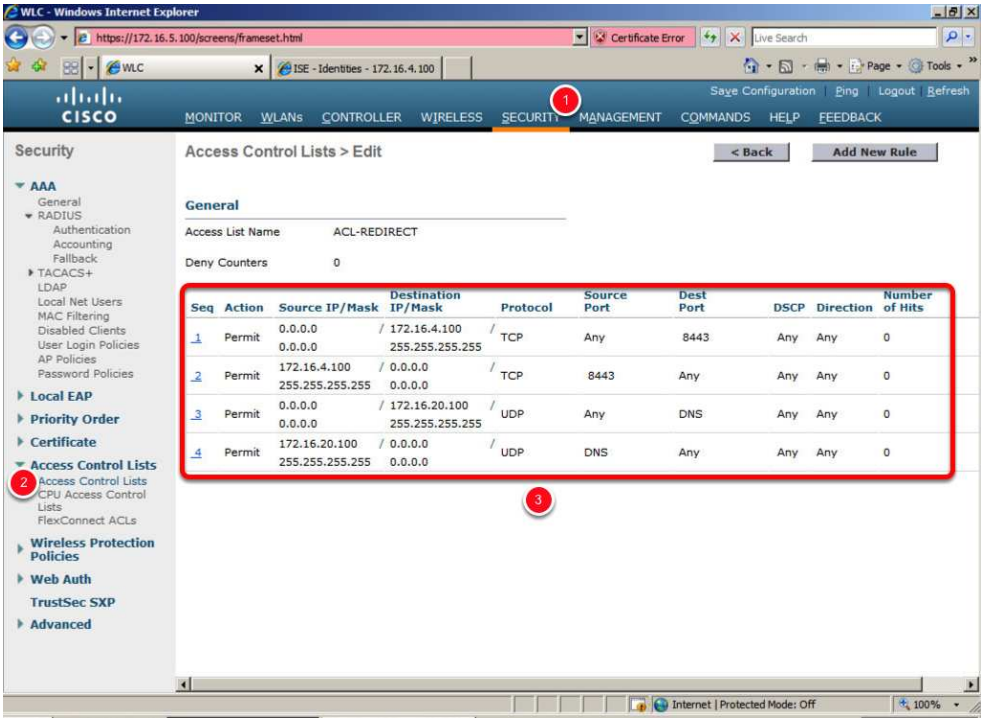


Navigate through the WLC GUI to configure the redirect ACL, which will be used to redirect user traffic toward the ISE captive portal:

https://t.me/learningnets

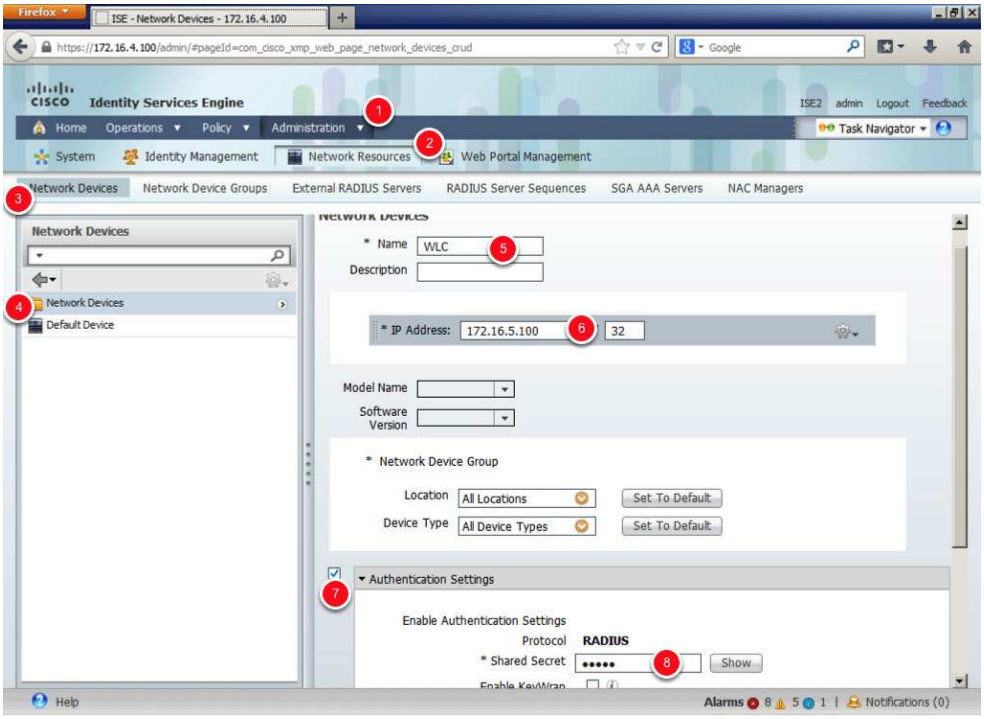


Configure this ACL with permit entries (traffic that is not redirected) for ISE portal and DNS server; all other traffic will match on the implicit deny (which will be redirected):

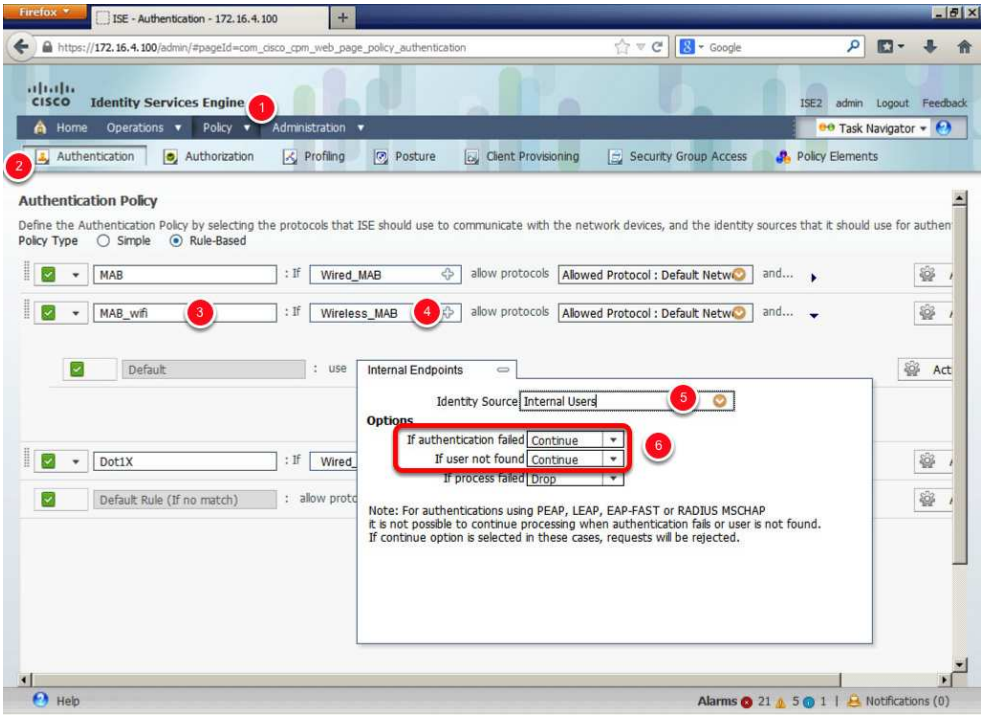


Add WLC as a NAD device into ISE:

https://t.me/learningnets

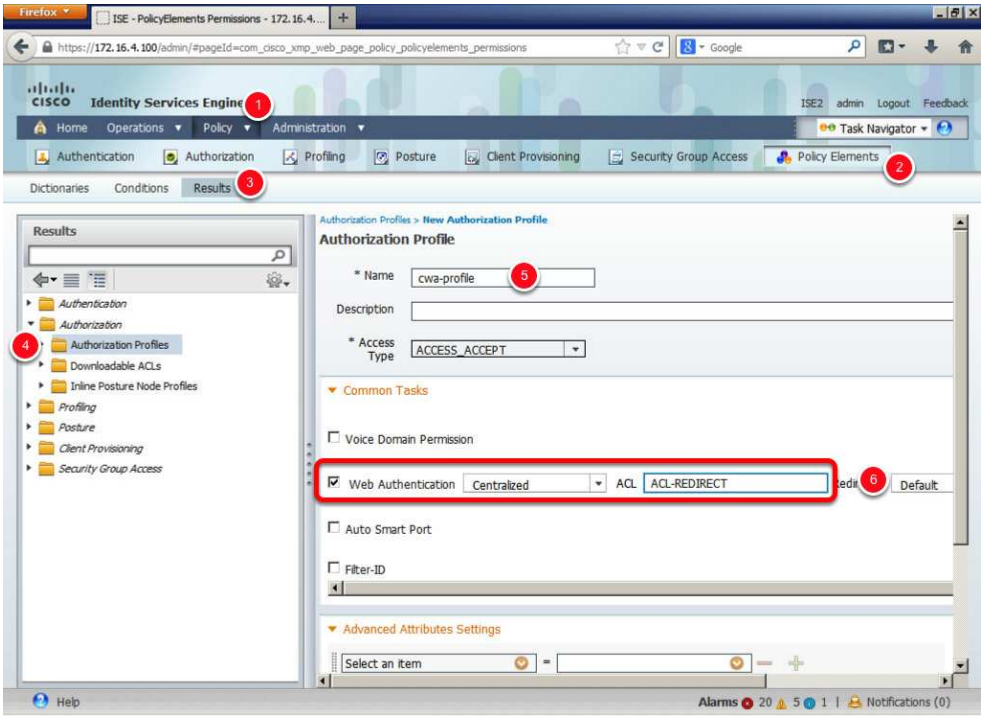


Configure an authentication policy on ISE for wireless MAB, use the Internal User store in order to allow username/password authentication in the portal; configure ISE to assume session is authenticated in case the MAC address is unknown, so that the CWA profile can be applied:

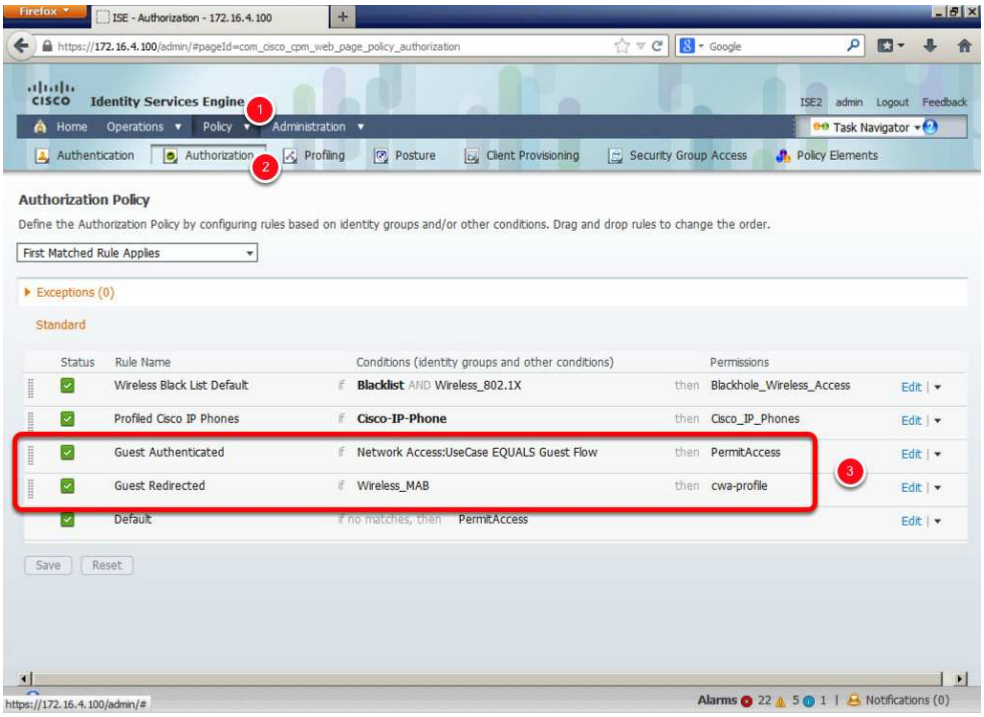


Configure an authorization profile on ISE to trigger the CWA process through the redirect-url and redirect-acl; the name of the ACL should match the name configured on WLC:

https://t.me/learningnets

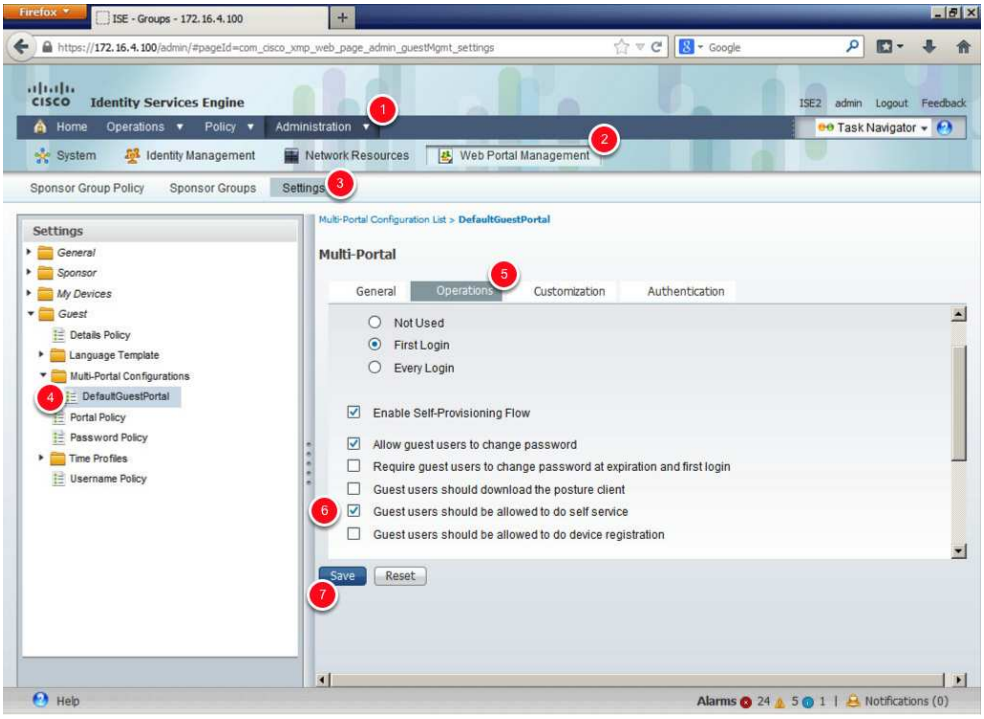


Configure two authorization policies, one for redirecting the user in the captive portal (CWA) and one for allowing user access if it has been authenticated in the portal. Because policies are processed top-down, ensure that the CWA policy is second in order; otherwise, you will find yourself in a loop where you will always be redirected to the captive portal:

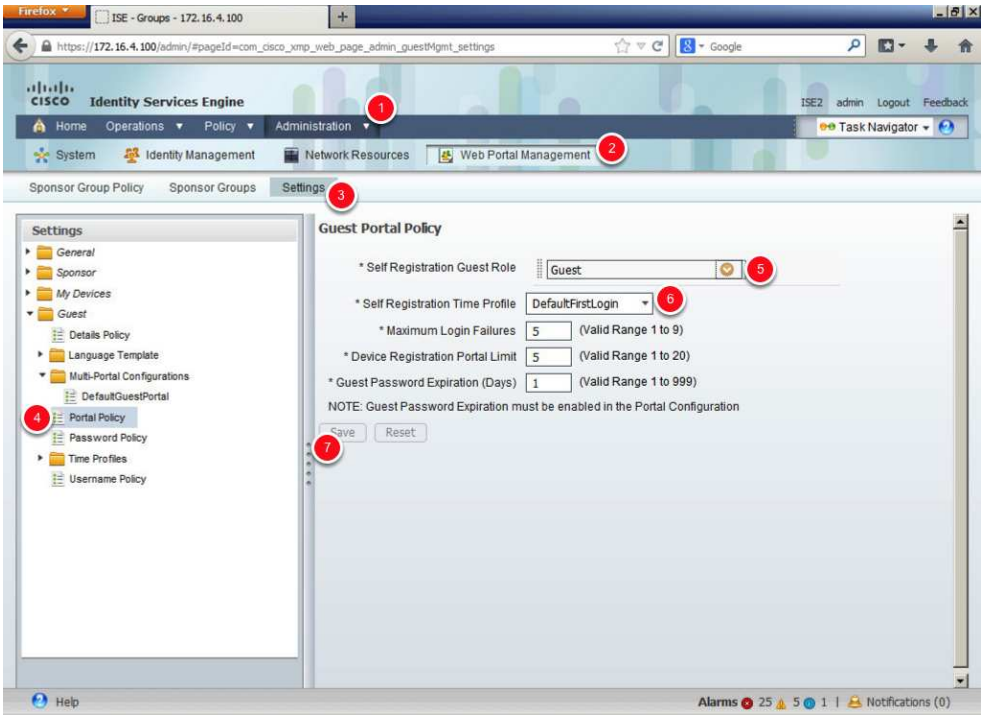


Enable self-service for guest portal, so that users are allowed to create their login credentials:

https://t.me/learningnets



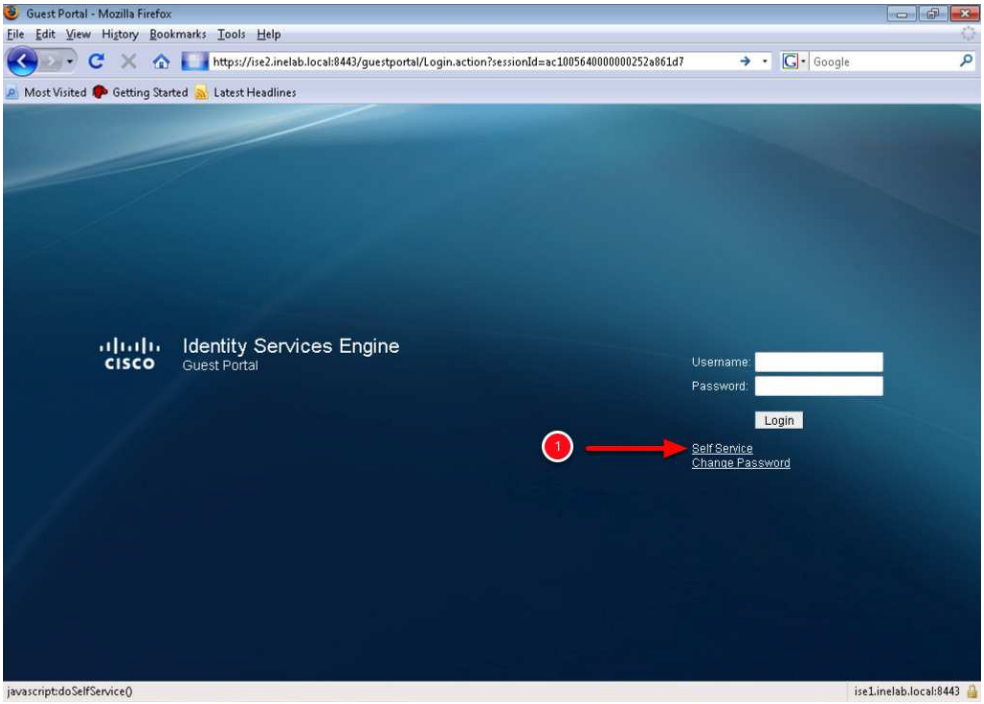
Configure the guest portal with a guest role (in which group will the user be included) and a time profile (defining how long the credentials will be valid):



Verification

Configure WiFi TEST-PC to associate to the configured SSID. Try browsing to a URL, such as example to ACS2; you will be redirected in the captive portal. Use the self-service functionality:

https://t.me/learningnets



Verify the client state before authenticating in the web portal, and also note the redirect-ACL and redirect-URL being applied:

```
(Cisco Controller) >show client summary

Number of Clients................................ 1

MAC Address       AP Name           Status        WLAN/RLAN      Auth Protocol         Port Wired
----------------- ----------------- ------------- -------------- ---- ---------------- ---- -----
ac:f1:df:13:bd:8f AP0006.f63b.9ada  Associated    1              No   802.11n(2.4 GHz) 1    N/A
!
!(Cisco Controller) >show client detail ac:f1:df:13:bd:8f
Client MAC Address.............................. ac:f1:df:13:bd:8f
Client Username ................................ N/A
AP MAC Address.................................. f0:29:29:04:21:f0
AP Name......................................... AP0006.f63b.9ada
Client State.................................... Associated
Client NAC OOB State............................ Access
Wireless LAN Id................................. 1
802.11u......................................... Not Supported
BSSID........................................... f0:29:29:04:21:f0
Connected For .................................. 196 secs
Channel......................................... 1
IP Address...................................... 172.16.5.200
IPv6 Address.................................... fe80::3c16:4073:f3a4:c679
Association Id.................................. 1
Authentication Algorithm........................ Open System
Reason Code..................................... 1
Status Code..................................... 0
```

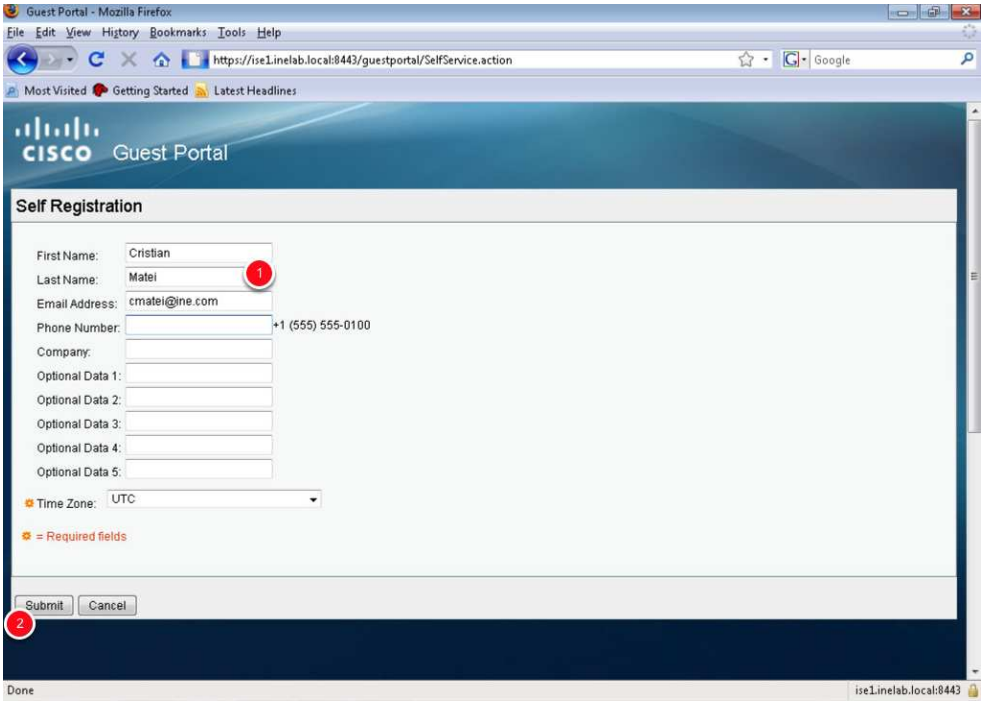
https://t.me/learningnets

```
Session Timeout.............................. 1800
Client CCX version........................... 5
Client E2E version........................... No E2E support
Diagnostics Capability....................... Supported
S69 Capability............................... Not Supported
QoS Level.................................... Silver
802.1P Priority Tag.......................... disabled
CTS Security Group Tag....................... Not Applicable
KTS CAC Capability........................... No
WMM Support.................................. Enabled
  APSD ACs.....................................  BK  BE  VI  VO
Power Save................................... OFF
Current Rate................................. 6.0
Supported Rates.............................. 1.0,2.0,5.5,11.0,6.0,9.0,
    ..........................................  12.0,18.0,24.0,36.0,48.0,
    ..........................................  54.0
Mobility State............................... Local
Mobility Move Count.......................... 0
Security Policy Completed.................... No
Policy Manager State......................... POSTURE_REQD
Policy Manager Rule Created.................. Yes
AAA Override ACL Name........................ ACL-REDIRECT
AAA Override ACL Applied Status.............. Yes
AAA Override Flex ACL Name................... none
AAA Override Flex ACL Applied Status......... Unavailable
AAA URL redirect............................. https://ISE2.inelab.local:8443/guestportal/gateway?sessionId=ac100

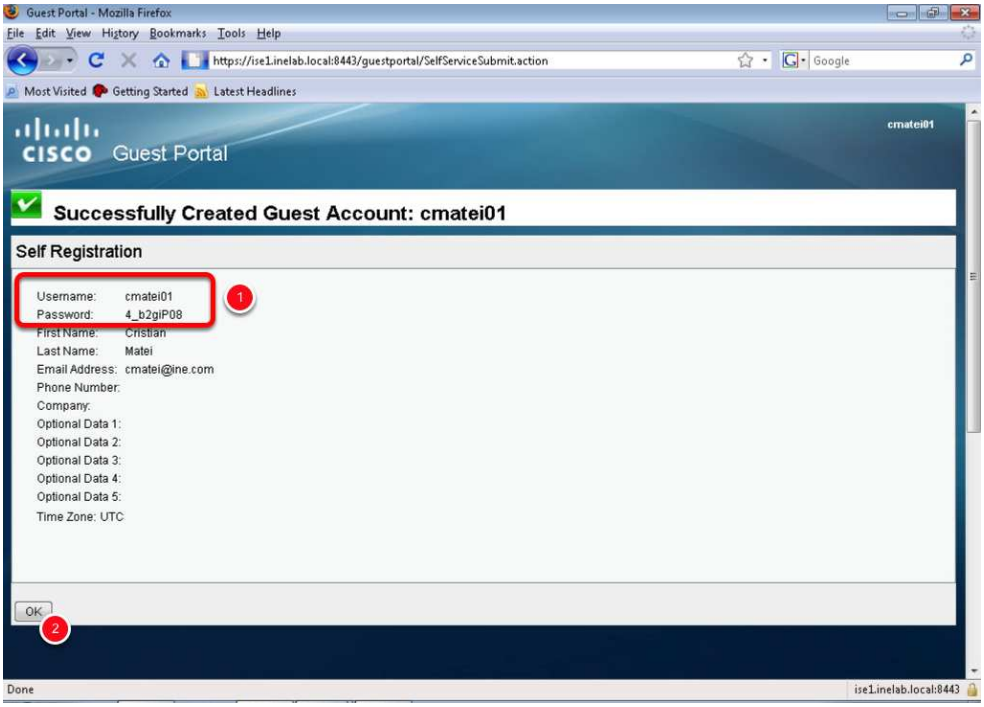
Audit Session ID............................. ac1005640000000252a861d7
IPv4 ACL Name................................ none
```

Configure the minimum required information for ISE to generate a username (and password):

https://t.me/learningnets

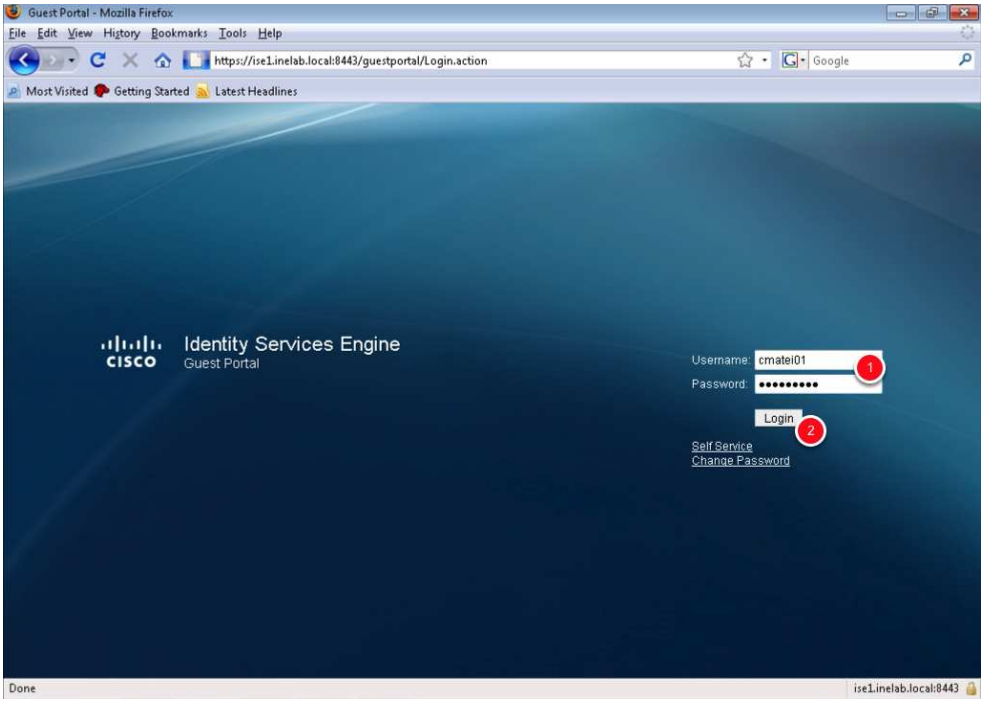


Memorize or save the generated username/password; you will need it to authenticate in the portal:

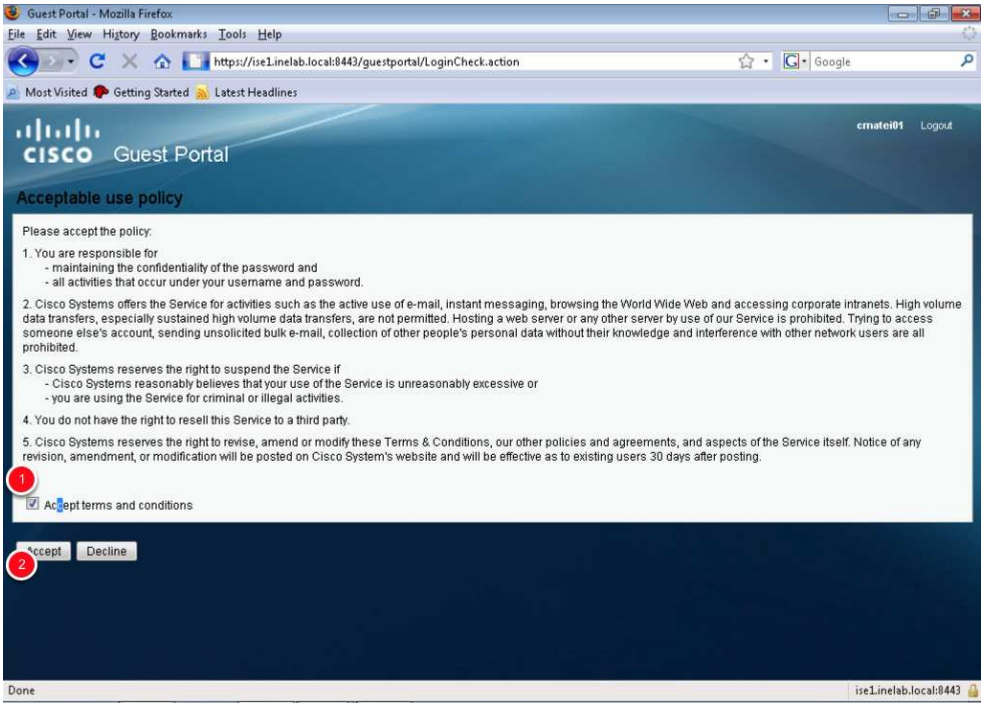


Authenticate using the generated username/password:

https://t.me/learningnets

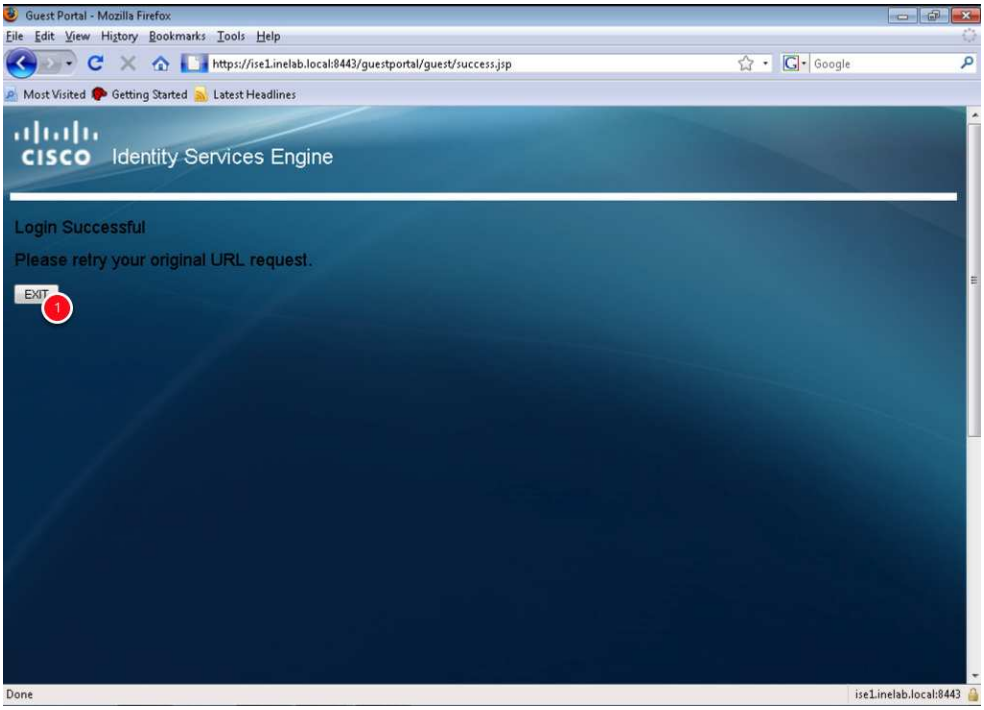


Accept the AUP:



At this moment, login should be successful (access to ACS should work now):

https://t.me/learningnets



After the user has been successfully authenticated, verify again the client state, redirect-URL is no longer applied; even though redirect-ACL is still applied, it has no effect without the redirect-URL:

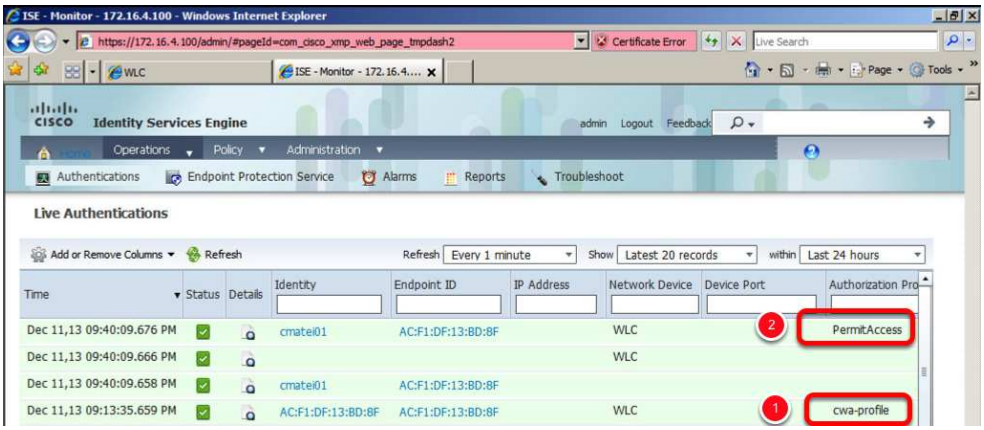
```
(Cisco Controller) >show client detail ac:f1:df:13:bd:8f
Client MAC Address.............................. ac:f1:df:13:bd:8f
Client Username ................................ N/A
AP MAC Address.................................. f0:29:29:04:21:f0
AP Name......................................... AP0006.f63b.9ada
Client State.................................... Associated
Client NAC OOB State............................ Access
Wireless LAN Id................................. 1
802.11u......................................... Not Supported
BSSID........................................... f0:29:29:04:21:f0
Connected For .................................. 114 secs
Channel......................................... 1
IP Address...................................... 172.16.5.200
IPv6 Address.................................... fe80::3c16:4073:f3a4:c679
Association Id.................................. 1
Authentication Algorithm........................ Open System
Reason Code..................................... 1
Status Code..................................... 0
Session Timeout................................. 3600
Client CCX version.............................. 5
Client E2E version.............................. No E2E support
Diagnostics Capability.......................... Supported
S69 Capability.................................. Not Supported
```

https://t.me/learningnets

```
--More-- or (q)uit
QoS Level..................................... Silver
802.1P Priority Tag........................... disabled
CTS Security Group Tag........................ Not Applicable
KTS CAC Capability............................ No
WMM Support................................... Enabled
  APSD ACs.....................................  BK  BE  VI  VO
Power Save.................................... OFF
Current Rate.................................. m0
Supported Rates............................... 1.0,2.0,5.5,11.0,6.0,9.0,
    ..........................................  12.0,18.0,24.0,36.0,48.0,
    ..........................................  54.0
Mobility State................................ Local
Mobility Move Count........................... 0
Security Policy Completed..................... Yes
Policy Manager State.......................... RUN
Policy Manager Rule Created................... Yes
AAA Override ACL Name......................... ACL-REDIRECT
AAA Override ACL Applied Status............... Yes
AAA Override Flex ACL Name.................... none
AAA Override Flex ACL Applied Status.......... Unavailable
AAA URL redirect.............................. none

Audit Session ID.............................. ac1005640000000252a861d7
IPv4 ACL Name................................. none
```

Verify the RADIUS authentication logs from ISE:



Verify the DHCP server configuration and leases:

```
(Cisco Controller) >show dhcp detailed CCIE-LAB22
Scope: CCIE-LAB22

Enabled....................................... Yes
Lease Time.................................... 86400 (1 day )
```

https://t.me/learningnets

```
Pool Start..................................... 172.16.5.200
Pool End....................................... 172.16.5.254
Network........................................ 172.16.5.0
Netmask........................................ 255.255.255.0
Default Routers................................ 172.16.5.1  0.0.0.0  0.0.0.0
DNS Domain..................................... inelab.local
DNS............................................ 172.16.20.100  0.0.0.0  0.0.0.0
Netbios Name Servers........................... 0.0.0.0  0.0.0.0  0.0.0.0
!
!(Cisco Controller) >show dhcp leases

       MAC                IP          Lease Time Remaining
ac:f1:df:13:bd:8f    172.16.5.200     23 hours 53 minutes 59 seconds
```

https://t.me/learningnets

CCIE Security v4 Technologies Workbook - Section 8: Wireless Security

WLC Wireless IPS

Task

- Configure WLC as follows:
  - activate a open authentication SSID named **CCIE-XX**, where **X** equals the rack number you've been assigned.
  - enable the internal DHCP server, assign IP addresses from range 172.16.5.200-172.16.5.254, and assign domain inelab.local and DNS server of 172.16.20.100.
  - don't fire any alerts for EAPOL flooding attacks.
- Configure IPS as follows:
  - use an IP address of 136.1.59.100 and SPAN VLAN 45 traffic to its Gi0/1 interface.
  - activate shunning for wireless clients.

> **A Note on Task Initial Configuration Files**: For this task, you must load the initial configuration files for the section, which can be found in the **Section 8 Introduction** or by clicking the **Resources** button.
>
> If have completed previous tasks in this section, the simplest way to revert back to initial configurations is by using the command `configure replace nvram:startup-config` instead of reloading the devices. This command, however, is not supported by the WLC or LAP; for LAP you can issue the command `clear lwapp private-config` and reload, while for WLC you can either not save the configuration and reload or load the initial configuration named **wlc-default.cfg** from the TFTP server 172.16.5.1 (from **Commands**-->**Download File** menu).

https://t.me/learningnets

Overview

WLC has a built-in set of 17 signatures (for both management and data-plane), which by default are enabled; if these are matched, an event will be fired and this action is non-configurable. These signatures are strictly related to basic radio attacks, in order to offer greater IPS support, WLC integrates with Cisco IPS dedicated solutions. How this works is that WLC will regularly (at the configured interval) read the IPS event store for blocked/shunned hosts and add these to the client exclusion list if:

- are active wireless clients.
- the respective WLAN has client exclusion enabled.

Note that clients will be excluded for the period configured at the IPS level, regardless of the interval specified at the WLC level; the WLC interval has precedence only for clients excluded by the WLC itself, not signaled from the IPS.

Configuration

SW3:

```
monitor session 1 source vlan 45 rx
monitor session 1 destination interface Gi1/0/9
!
interface GigabitEthernet1/0/9
 no shutdown
```

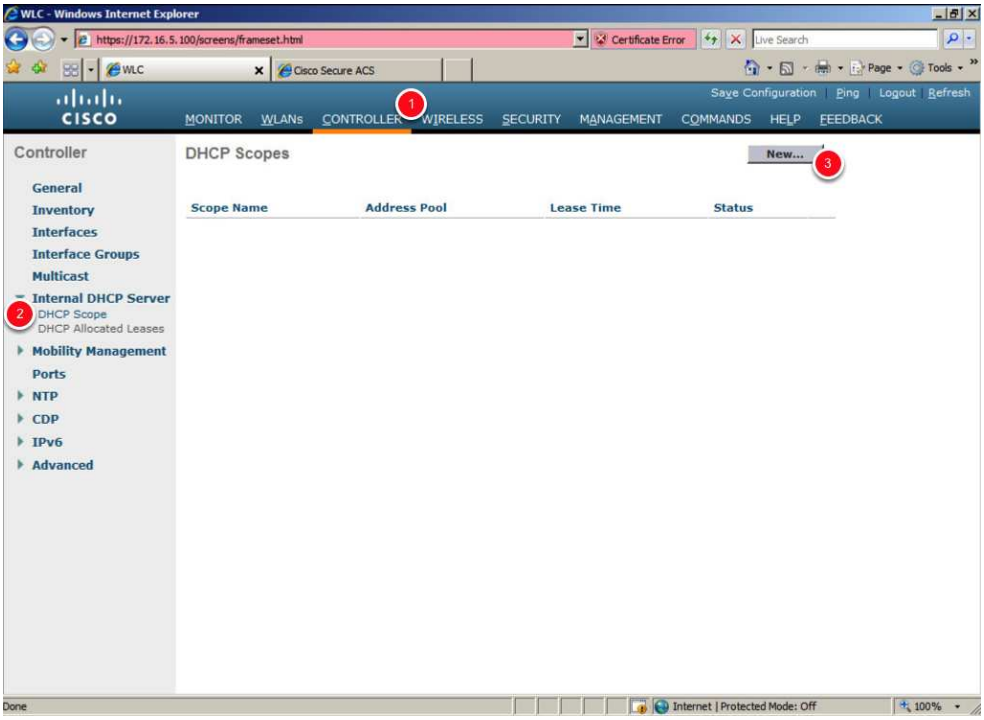
IPS:

```
sensor(config)# service host
sensor(config-hos)# network-settings
sensor(config-hos-net)# host-ip 136.1.59.100/24,136.1.59.13
sensor(config-hos-net)# access-list 172.16.5.100/32
sensor(config-hos-net)# exit
sensor(config-hos)# exit
Apply Changes?[yes]:
!
sensor(config)# service interface
sensor(config-int)# physical-interfaces gigabitEthernet0/1
sensor(config-int-phy)# admin-state enabled
sensor(config-int-phy)# exit
sensor(config-int)# exit
Apply Changes?[yes]:
!
sensor(config)# service network-access
```

https://t.me/learningnets

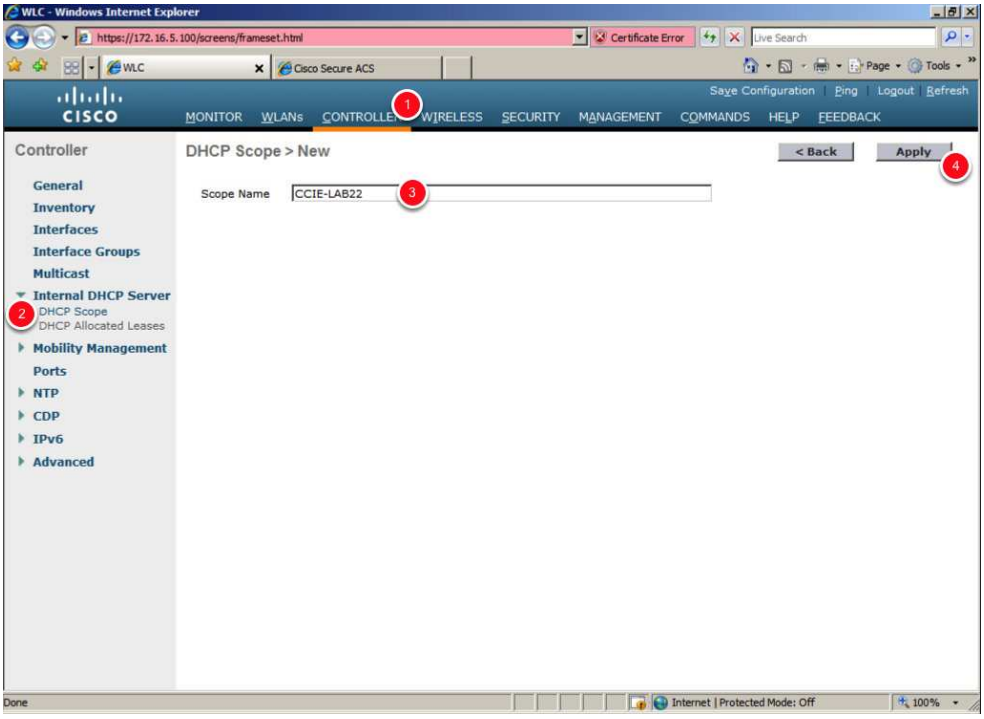
```
sensor(config-net)# general
sensor(config-net-gen)# block-enable true
sensor(config-net-gen)# log-all-block-events-and-errors true
sensor(config-net-gen)# exit
sensor(config-net)# exit
Apply Changes?[yes]:
!
sensor(config)# username wireless password qazwsx1234 privilege viewer
```

Navigate through the WLC GUI to create a new DHCP scope for the WLC internal server:

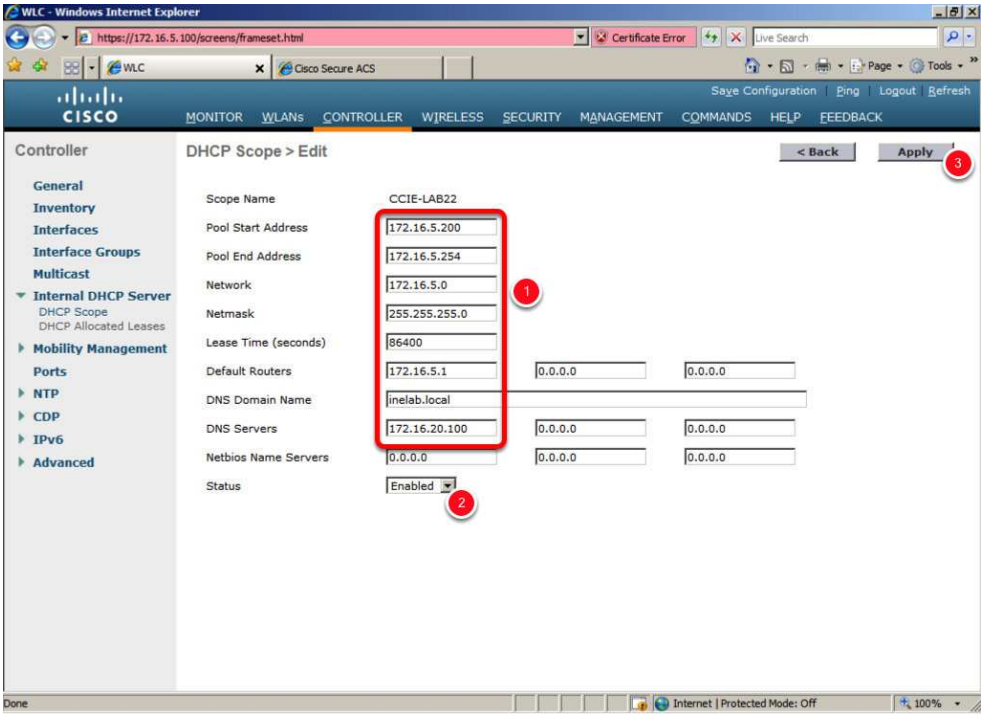


Give the DHCP scope a name:

https://t.me/learningnets

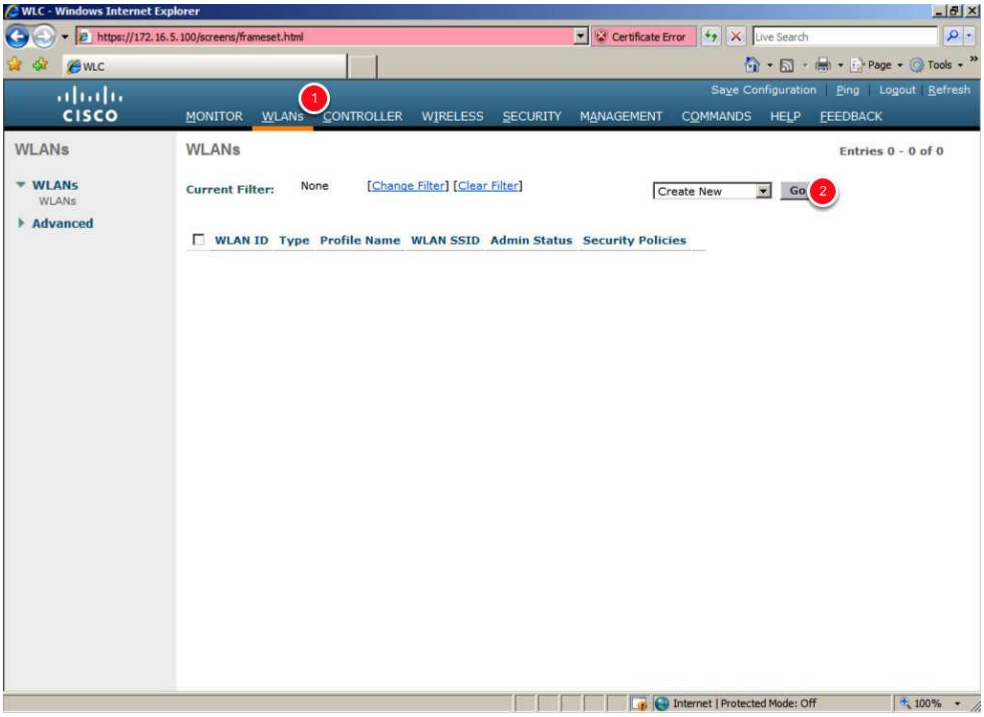


Configure the DHCP settings and enable it, by default it is disabled:

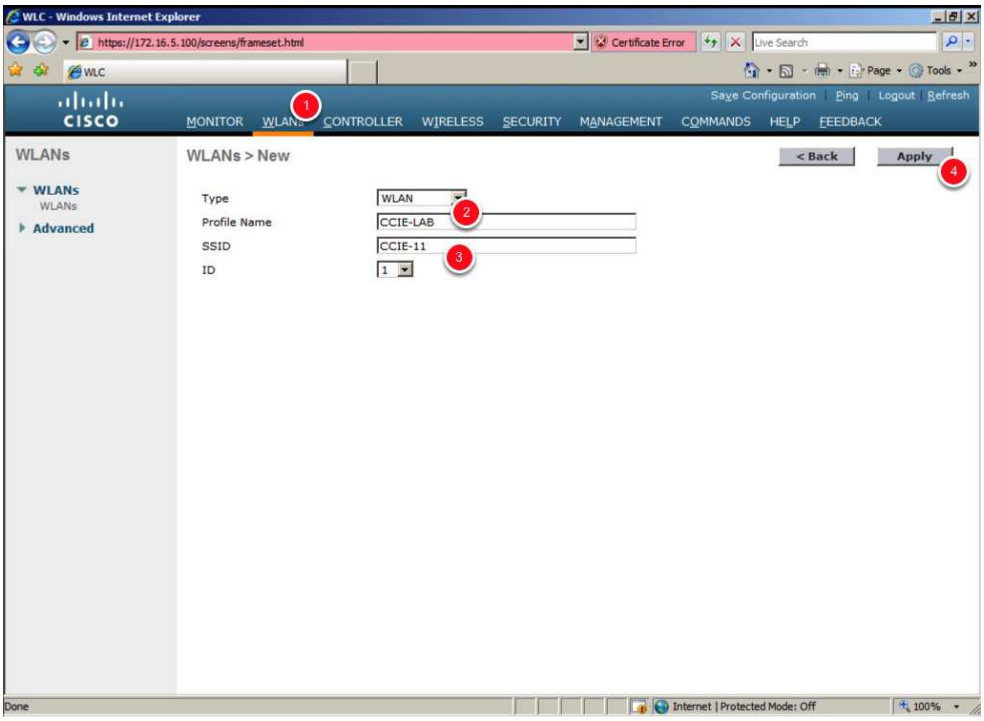


Navigate through the WLC GUI to create a new WLAN:

https://t.me/learningnets

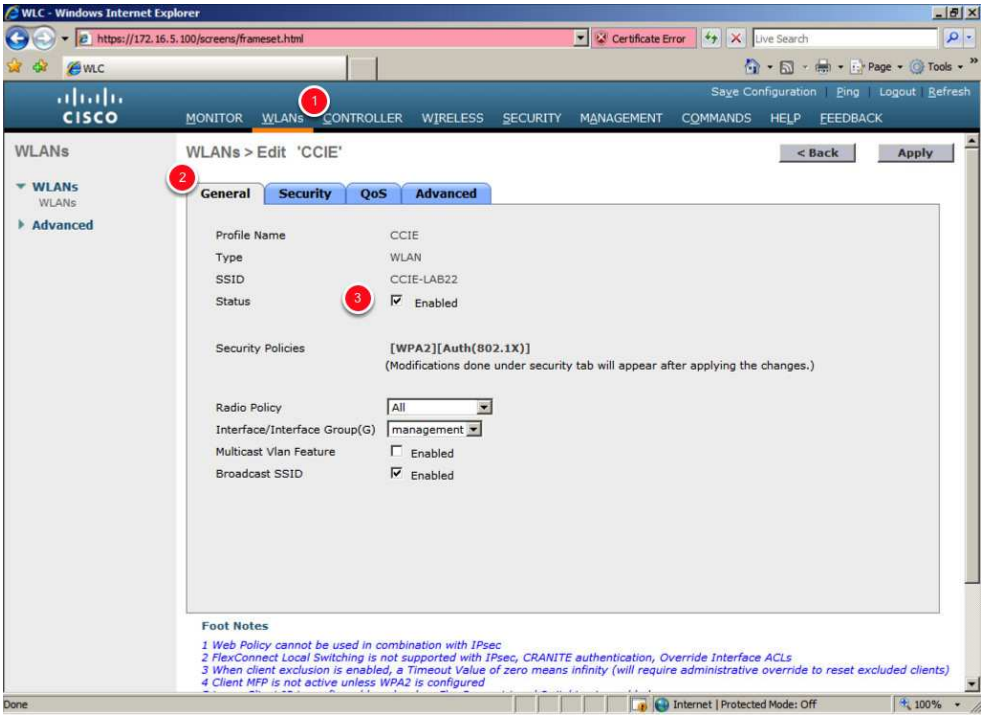


Create the WLAN, and make sure to correctly type the SSID, which is case-sensitive:

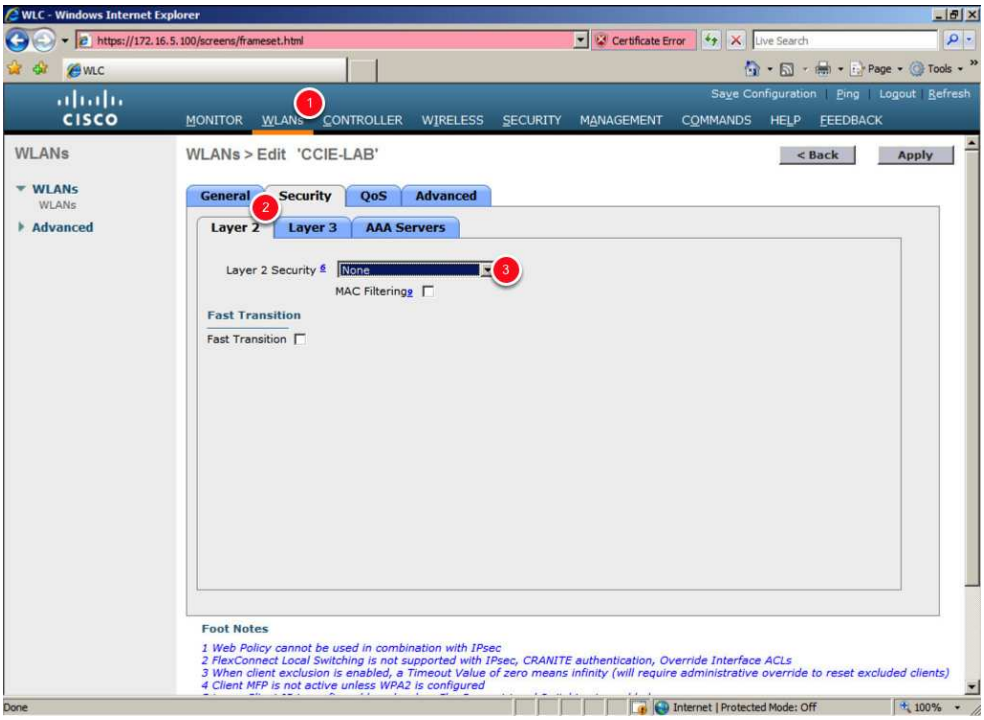


Enable the WLAN and leave WLC to broadcast the SSID:

https://t.me/learningnets

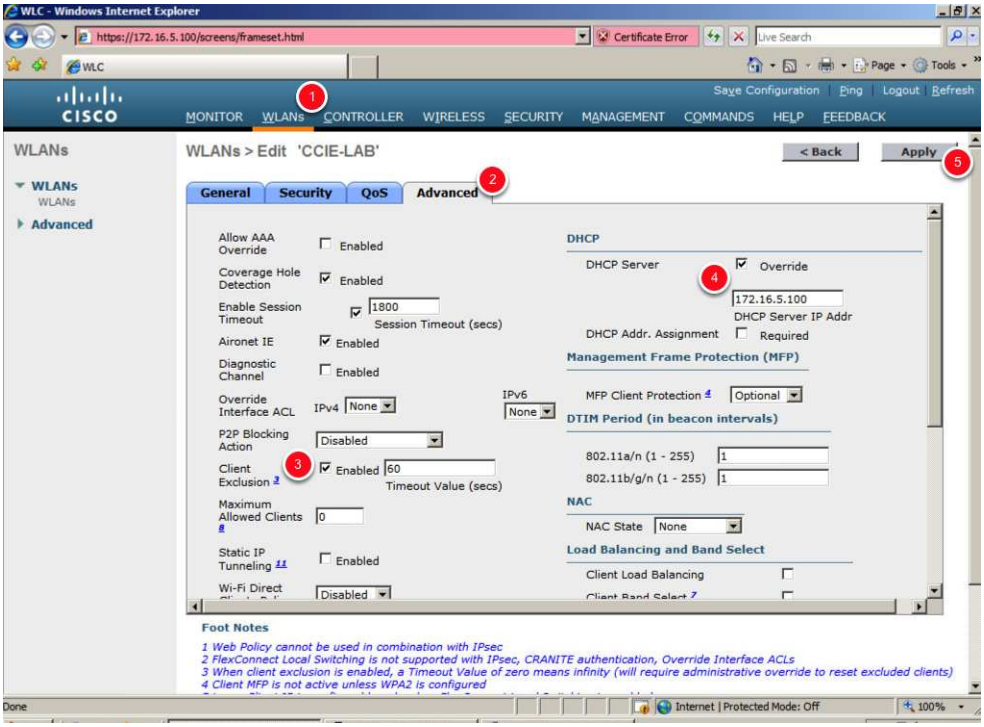


Configure the WLAN so that no authentication is required for joining:



Ensure client exclusion is enabled for the WLAN; specify the internal DHCP server in order to assign IP address to clients of the WLAN:

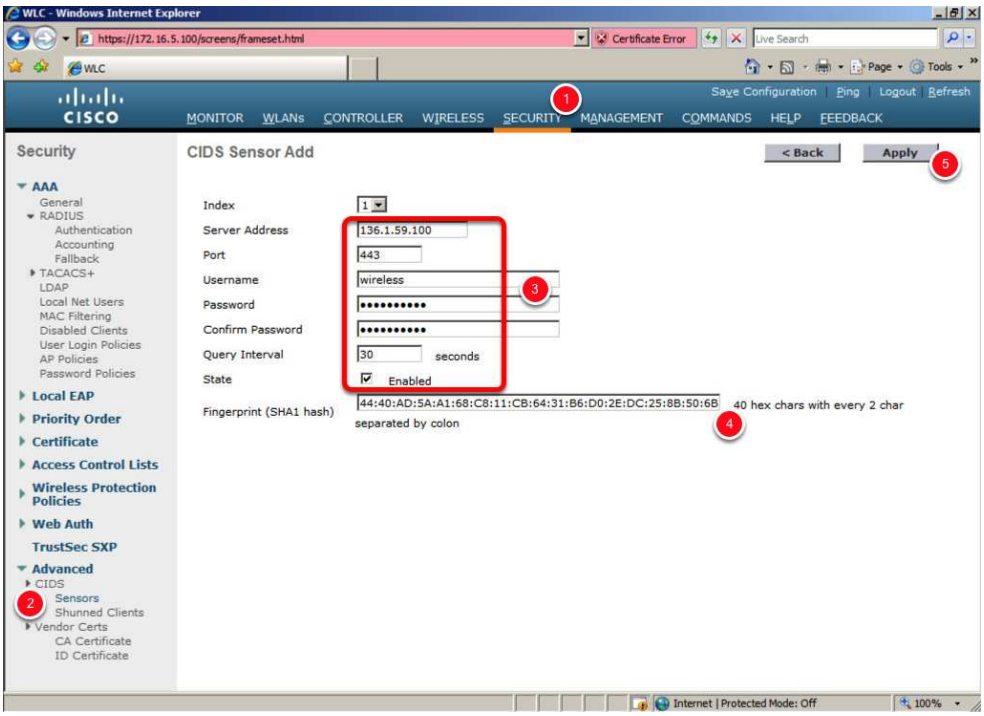
https://t.me/learningnets



Identify the SHA1 hash of the IPS certificate:

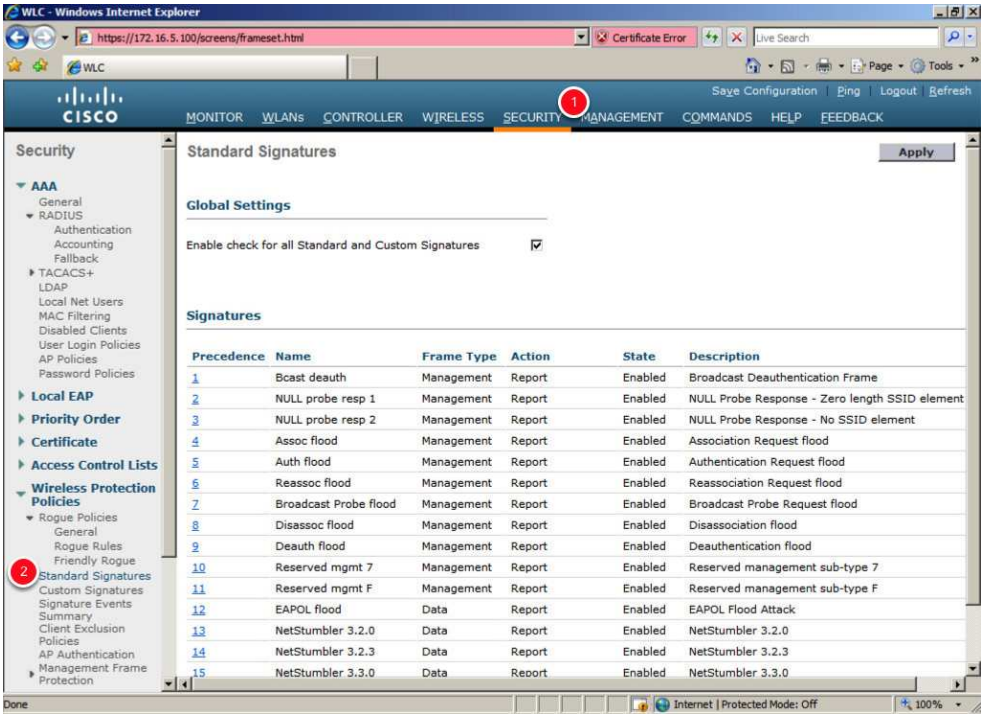
```
sensor# show tls fingerprint
MD5: 9A:D6:2D:17:B1:88:0B:2A:B6:9A:0E:39:FE:EF:86:94
SHA1: 44:40:AD:5A:A1:68:C8:11:CB:64:31:B6:D0:2E:DC:25:8B:50:6B:40
```

Configure the WLC integration with IPS:

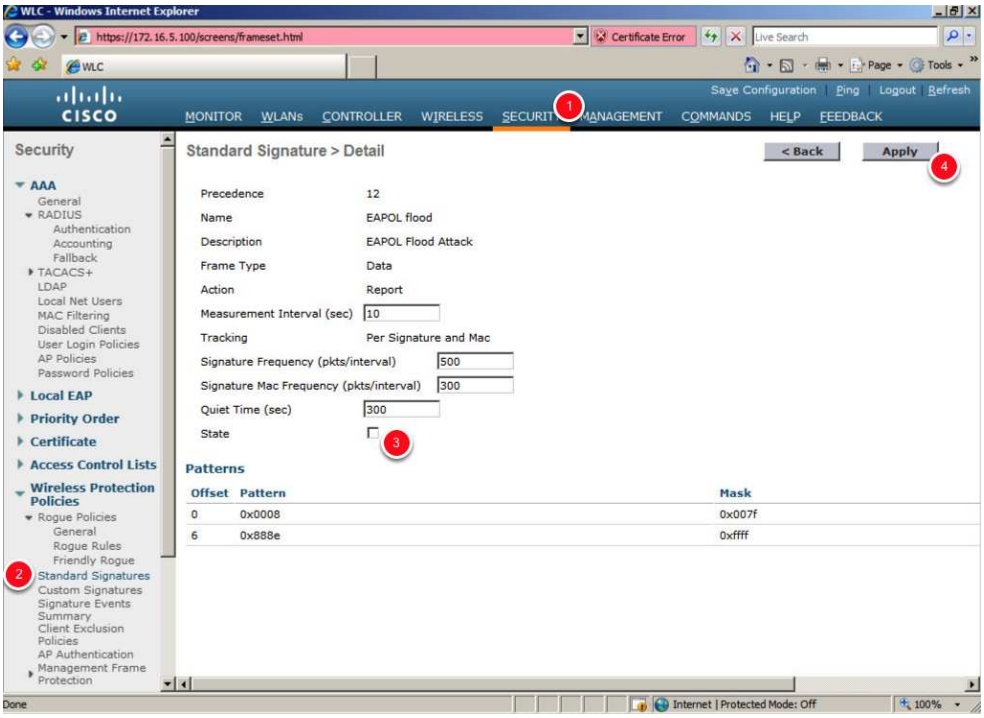


https://t.me/learningnets

Navigate through the GUI WLC in order to identify that by default all standard signatures are enabled:



Disable the EAPOL flood attack signature:



Verification

Verify that WLC is successfuly connected to IPS:

https://t.me/learningnets

```
(Cisco Controller) >show wps cids-sensor summary


Configured IDS Sensors

Index  Server Address    Port    State      Intvl   Last Query
-----  ---------------   ------  ---------  ------  --------------
1      136.1.59.100      443     Enabled    30      Success
```

Verify that EAPOL flood attack signature is disabled:

```
(Cisco Controller) >show wps signature detail 12

Signature-ID..................................... 12
Precedence....................................... 12
Signature Name................................... EAPOL flood
Type............................................. standard
FrameType........................................ data
State............................................ disabled

Action........................................... report
Tracking......................................... per Signature and Mac
Signature Frequency.............................. 500 pkts/interval
Signature Mac Frequency.......................... 300 pkts/interval
Interval......................................... 10 sec
Quiet Time....................................... 300 sec
Description...................................... EAPOL Flood Attack
Patterns:
                0(Header):0x0008:0x007f
                6(Body):0x888e:0xffff
```

In order to test the client-exclusion functionality, enable a signature on IPS and set its action to request-block-host and alert:

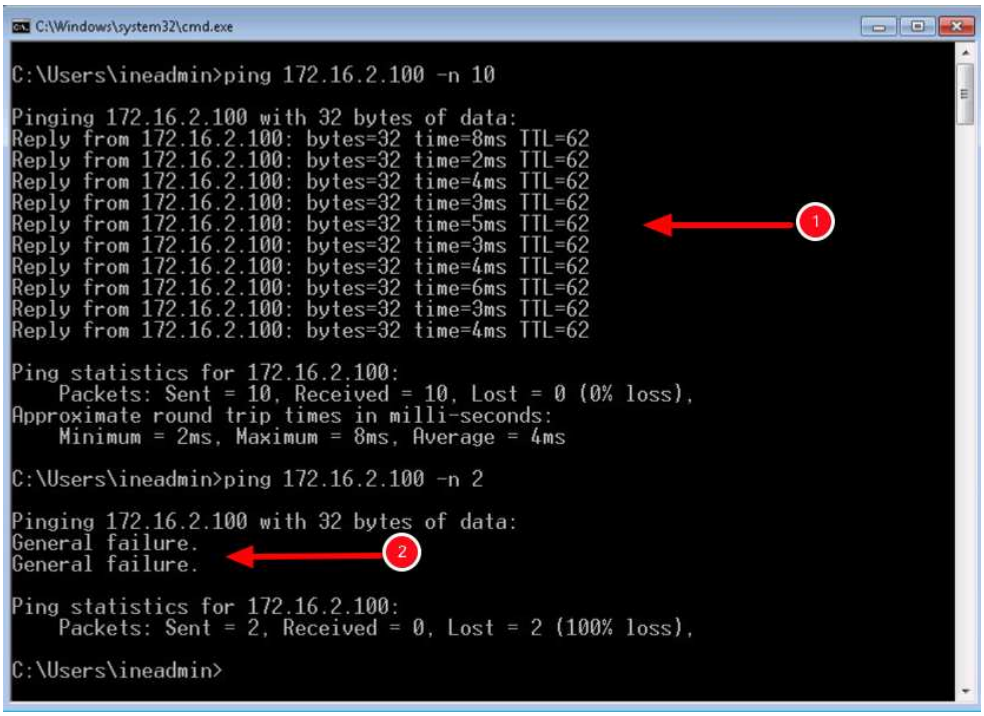
IPS:

```
sensor(config)# service signature-definition sig0
sensor(config-sig)# signatures 2004 0
sensor(config-sig-sig)# status
sensor(config-sig-sig-sta)# enabled true
sensor(config-sig-sig-sta)# exit
sensor(config-sig-sig)# engine atomic-ip
sensor(config-sig-sig-ato)# event-action request-block-host|produce-alert
sensor(config-sig-sig-ato)# exit
```

https://t.me/learningnets

```
sensor(config-sig-sig)# exit
sensor(config-sig)# exit
Apply Changes?[yes]:
```

Initiate ICMP Echo traffic from WiFi TEST-PC towards ACS2 in order to trigger the signature; first packets are successful, until client is excluded:



Verify that client has been blocked for the period specified on the IPS (default of 30 minutes), regardless of the configuration on the WLC side:

```
(Cisco Controller) >show wps shun-list

Total number of Entries ..... 1

IP Address        Last MAC address     Expire     Sensor IP/Index
---------------   ------------------   ---------  ----------------
172.16.5.200      ac:f1:df:13:bd:8f    30         136.1.59.100/1
```

https://t.me/learningnets