IPv4 Routing With OSPF

Keith Bogart

CCIE #4923

✉ kbogart@ine.com
🐦 @keithbogart1
in linkedin.com/in/keith-bogart-2a75042

CCIE Routing & Switching

Course Objectives

+ Help You To Learn The Theory & Operation Of OSPF For IPv4 (OSPFv2)
+ Provide Multi-Vendor Configuration Examples

Course Prerequisites

+ Experience With IPv4 Addressing & Subnetting
+ Knowledge Of IP Routing Basics

An Introduction To OSPF

ine.com

Topic Overview

+ Introduction To OSPF
+ OSPF Metric
+ Introduction To LSAs
+ OSPF Packet Types

Introduction To OSPF

+ Open standard (RFC 2328)
+ Link State routing protocol
+ SPF (shortest path first) algorithm
+ Assigned to IP protocol 89
+ Supports several methods of authentication
    + Clear Text
    + MD5
    + HMAC with SHA (OS-dependent)

Hashed Message Authentication Code with Secure Hash Algorithm

OSPF Metric

+ OSPF Metric = Cost
+ OSPF Cost based on interface bandwidth
    + The greater the BW...the lower the Cost
+ Best path to any given network = path with lowest OSPF cost
+ Can be configured explicitly or affected by changing interface bandwidth.

OSPF LSAs

+ LSA = Link State Advertisement
+ Carried within an OSPF Link State Update Packet
+ Different types carry different data
+ Age out after 1-hour...refreshed every 30-minutes.

OSPF Packet Types

- + Multiple OSPF Packet Types
    - + Hello
    - + Database Descriptor
    - + Link State Request
    - + Link State Update
    - + Link State Acknowledgment
- + Most packets sent as multicasts
    - + 224.0.0.5 (All OSPF Routers)
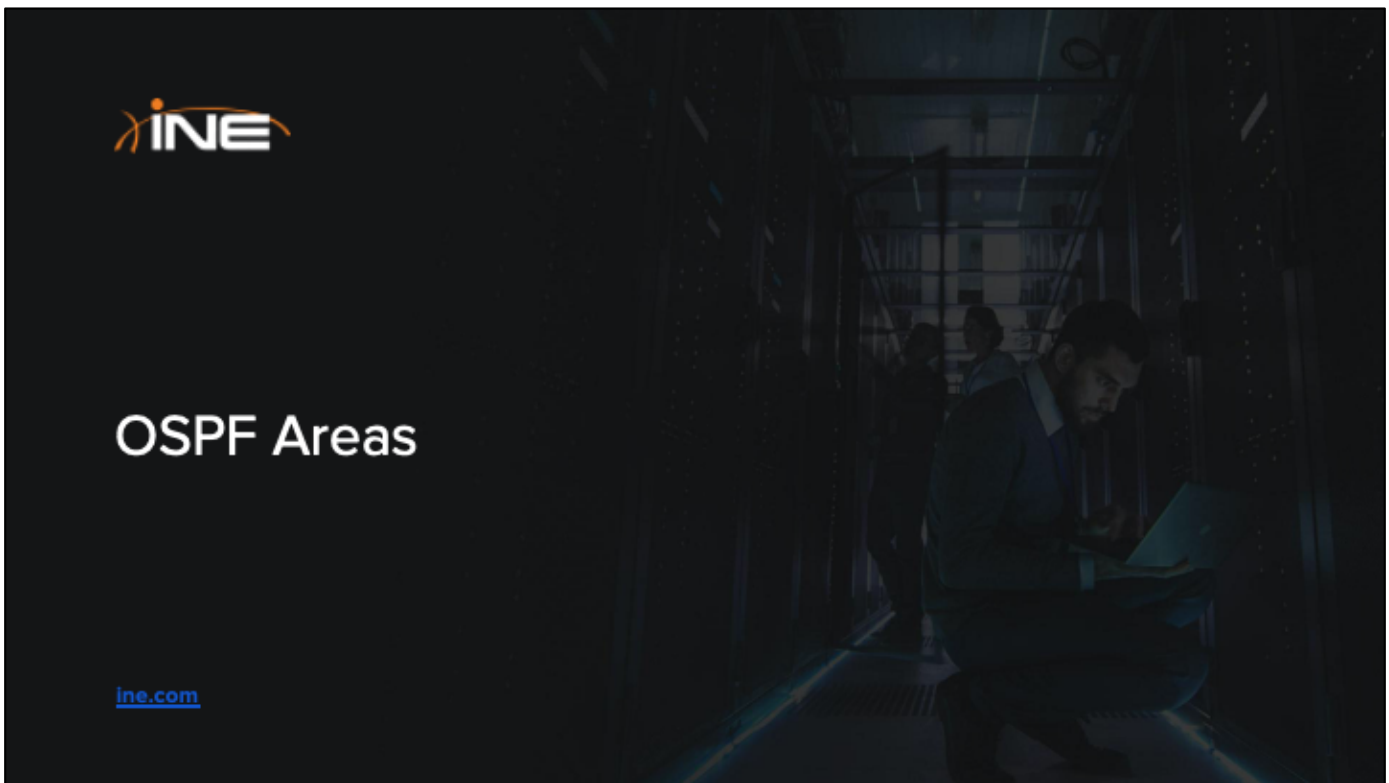   - + 224.0.0.6 (All OSPF Designated Routers)

OSPF Areas

Topic Overview

+ OSPF Area Overview

OSPF Areas

+ OSPF links are defined as belonging to OSPF Areas
+ Areas limit the scope of LSA flooding for most LSA types
+ OSPF domains can be constructed using a single area or multiple areas
+ Filtering and summarization only possible when utilizing multiple areas

OSPF Area Overview

+ In topologies requiring two-or-more areas, one area must be defined as the Backbone Area = Area 0
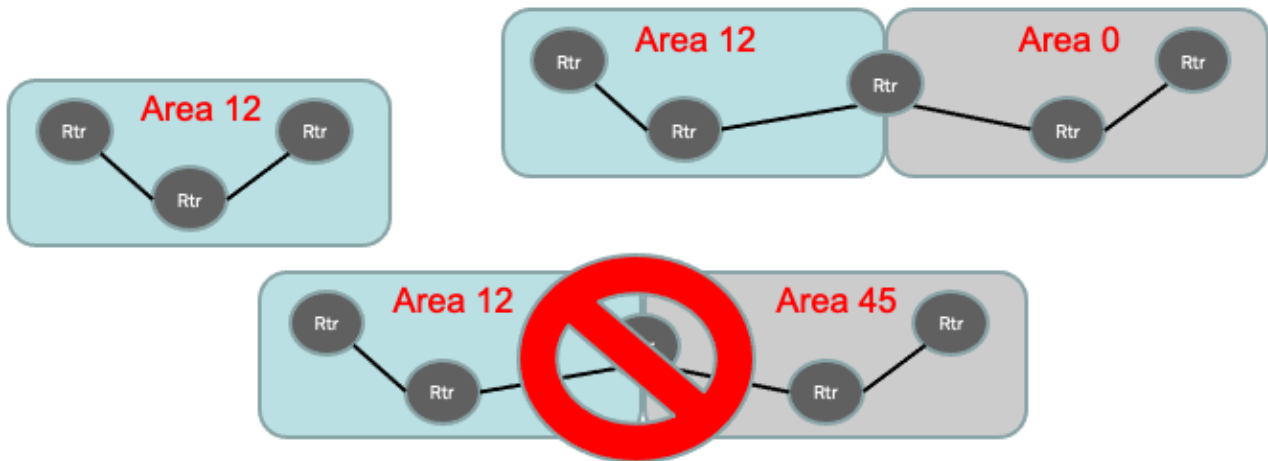
Thanks for Watching!

Router Roles & OSPF Tables

ine.com

Topic Overview

+ Router Roles
+ OSPF Tables

Router Roles

+ OSPF routers perform one-or-more roles
  + OSPF Internal Router

    Area 1 — (Rtr) — Area 1

  + OSPF Backbone Router

    Area 0 — (Rtr) — Area 0

  + ABR

    Area 0 — (Rtr) — Area 1

  + ASBR

    Non-OSPF IGP → (Rtr) — Area X

OSPF Tables

- Neighbor table
  - Contains neighbor information
- Link State Database
  - Contains all LSAs received within an area.
- SPF Tree
- Routing table
  - Contains best route on the basis of link cost

Thanks for Watching!

Topic Overview

+ Hello Packet Details
+ Adjacency Requirements
+ Adjacency Process

Building Adjacencies

+ Adjacency discovery utilizes OSPF "Hello" packets
  + Hello's typically sent to "All OSPF Routers" address of 224.0.0.5
  + May need to manually specify neighbor address depending on network type.
+ Exchange of Database Descriptors
+ Exchange of Link State Requests and Updates
+ Fully Loaded

Adjacency Requirements

+ Several parameters within Hello packets must match between neighbors:
    + Hello & Dead intervals
    + Area ID
    + Subnet mask
    + Stub area flag
    + Authentication
+ Hello interval (10-seconds on most interfaces)
+ Dead interval (4x the Hello interval)

Forming Adjacencies - Overview

Adjacency Process

+ Two types of OSPF relationships
  + 2-Way
  + Full Adjacencies
+ OSPF neighbors go through several steps in order to become fully adjacent:
  + INIT
  + 2-Way
  + EXSTART
  + Exchange
  + Loading
  + Full

Initial OSPF Configuration

ine.com

Topic Overview

- + General Configuration Guidelines
- + Cisco Basic OSPF Configuration
- + Vyos Basic OSPF Configuration
- + Suppressing Hellos
- + Changing Hello Timers
- + Monitoring OSPF Neighbors

General OSPF Configuration Parameters

+ OSPF must know the following things in order to run:
    + Router-Identifier
    + Process-ID (OS-dependent)
    + Which interfaces should run OSPF
    + Area ID for active interfaces
+ Some of the items above can be determined dynamically, others must be specified via configuration
+ OSPF configuration can be applied:
    + Directly upon an interface
    + Within an OSPF process
    + Operating System dependent

Cisco OSPF Initial Configuration (Legacy)

+ OSPF process & networks must be activated

20.20.20.1 /24   190.190.190.1 /29

Gig0/0   Gig1/0

```
router ospf <process-id>
    network 20.20.20.0 0.0.0.255 area <area-id>
    network 190.190.190.1 0.0.0.0 area <area-id>
```

Cisco OSPF Initial Configuration (Current)

+ OSPF process & networks must be activated

20.20.20.1 /24   190.190.190.1 /29
Gig0/0   Gig1/0

interface Gig0/0
    ip ospf <process-id> area <area-id>
interface Gig1/0
    ip ospf <process-id> area <area-id>

Vyos Router Configuration

```
vyos@vyos# set protocols ospf area 0 network  ?
Possible completions:
   <x.x.x.x/x>   OSPF network [REQUIRED]
   124.1.1.0/24
```

+ In this example, the address of 124.1.1.1/24 has already been configured on an interface.

Suppressing Hellos

+ OSPF "network" command (or enabling directly on an interface) causes transmission of Hello packets.
+ Sometimes this is undesirable, so an interface can be made "passive";

    Router(config-router)#passive-interface FastEthernet0/0
    Or...
    Router(config-router)#passive-interface default
    Router(config-router)#no passive-interface FastEthernet1/1

Vyos

```
vyos@vyos# set protocol ospf passive-interface
Possible completions:
   <interface>  Interface to be passive (i.e. suppress routing updates)
   default      Default to suppress routing updates on all interfaces
   eth0
   eth1
   lo
```

ip ospf area 0 Fast1/1

ip ospf area 0 Fast0/0

Hello Message Timers (Cisco)

+ OSPF timers may be changed from default values if desired on most platforms.
+ This is typically done within interface-configuration mode
+ To change hello and dead intervals on an interface, use following commands:
    + **[ip | ipv6] ospf hello-interval** <value>
    + **[ip | ipv6] ospf dead-interval** <value>
+ To verify these intervals:
    + **show [ip | ipv6] ospf interface** <interface>

```
R1#show ip ospf interface Gig0/6 | i Timer
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

Hello Message Timers (Vyos)

```
vyos@vyos# set interfaces ethernet eth0 ip ospf ?
   dead-interval
                Interval after which neighbor is dead
   hello-interval
                Interval between hello packets
```

```
vyos@vyos:~$ sho ip ospf interface eth0
eth0 is up
  ifindex 2, MTU 1500 bytes, BW 10000 Mbit <UP,BROADCAST,RUNNING,MULTICAST>
  Internet Address 124.1.1.1/24, Broadcast 124.1.1.255, Area 0.0.0.0
  MTU mismatch detection: enabled
  Router ID 124.1.1.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  No backup designated router on this network
  Multicast group memberships: OSPFAllRouters OSPFDesignatedRouters
  Timer intervals configured, Hello 18s, Dead 54s, Wait 54s, Retransmit 5
     Hello due in 12.471s
  Neighbor Count is 0, Adjacent neighbor count is 0
```

Monitoring OSPF Neighbors (Cisco)

```
R1#sho ip ospf neighbor

Neighbor ID      Pri    State         Dead Time    Address            Interface
122.122.122.22    1     FULL/BDR      00:00:39     122.122.122.22     GigabitEthernet0/7
112.112.112.2     1     EXSTART/BDR   00:00:39     112.112.112.2      GigabitEthernet0/2
```

```
R1#show ip ospf interface brief
Interface    PID    Area    IP Address/Mask      Cost    State   Nbrs F/C
Gi0/6         1      0      11.11.11.2/24         1      WAIT    0/1
Gi0/7         1      0      122.122.122.2/24      1      DR      1/1
Gi0/2         1      0      112.112.112.1/24      1      DR      1/1
```

Monitoring OSPF Neighbors (Vyos)

```
vyos@vyos:~$ show ip ospf neighbor

Neighbor ID     Pri State          Dead Time Address      Interface            RXmtL RqstL DBsmL
18.1.1.4          1 Full/DR          33.205s 124.1.1.4     eth0:124.1.1.1            0     0     0
```

```
vyos@vyos:~$ sho ip ospf interface
eth0 is up
  ifindex 2, MTU 1500 bytes, BW 10000 Mbit <UP,BROADCAST,RUNNING,MULTICAST>
  Internet Address 124.1.1.1/24, Broadcast 124.1.1.255, Area 0.0.0.0
  MTU mismatch detection: enabled
  Router ID 124.1.1.1, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State Backup, Priority 1
  Backup Designated Router (ID) 124.1.1.1, Interface Address 124.1.1.1
  Multicast group memberships: OSPFAllRouters OSPFDesignatedRouters
  Timer intervals configured, Hello 10s, Dead 40s, Wait 40s, Retransmit 5
    Hello due in 7.472s
  Neighbor Count is 1, Adjacent neighbor count is 1
```

OSPF Router Identifiers
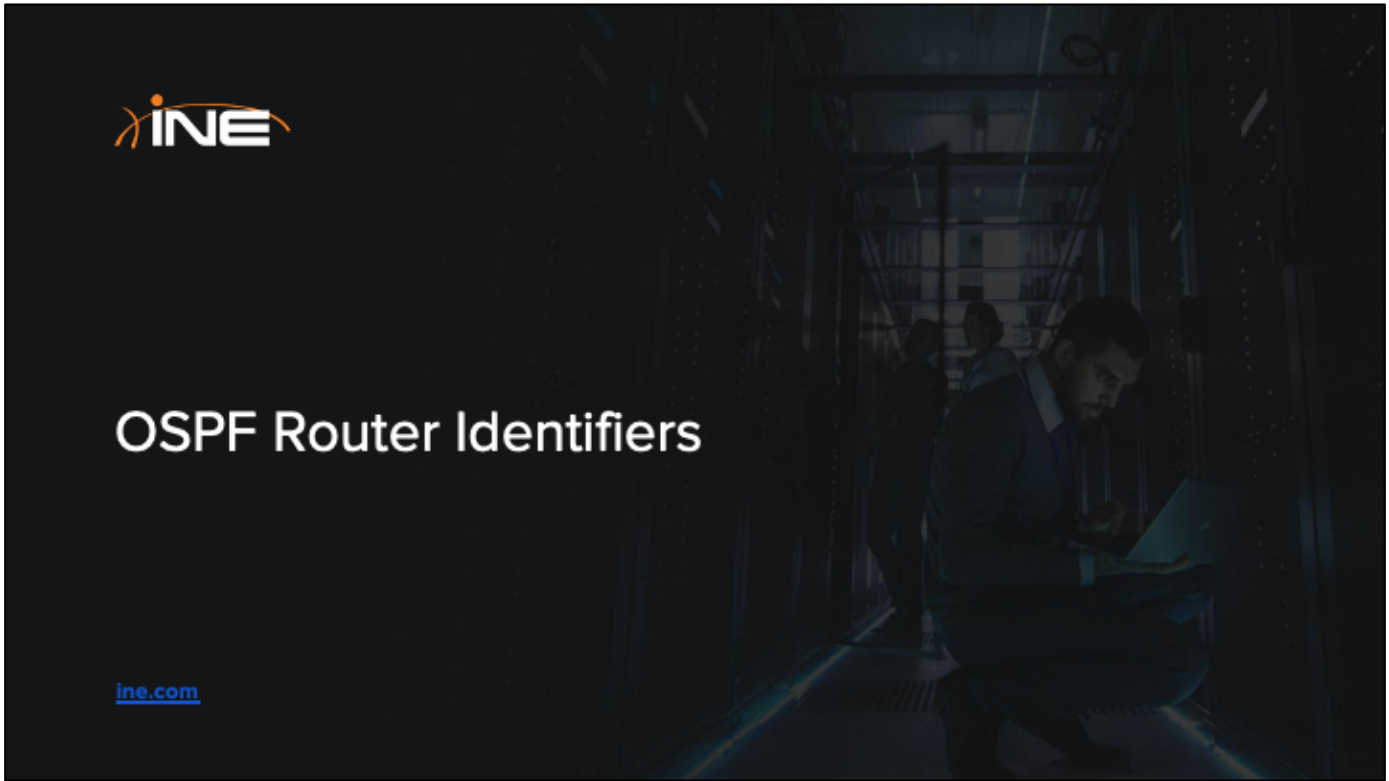
Topic Overview

+ The "What" and "Why" Of OSPF Router-IDs
+ Configuring Router-IDs
+ Router-ID Caveats

Router-ID

+ Characteristics of the OSPF Router-ID:
    + Typically dynamically generated
    + Can be locally configured
    + Should be unique to the OSPF domain
+ Why is it important?
    + Determines the outcome of most OSPF elections by default
    + Identifies the originator of LSAs

OSPF Router-ID

+ OSPF elects a router-id when the process comes up

+ Router-id is elected based on the following hierarchy;

  + Router-id command under the OSPF instance
  + Highest IP address of a loopback interface
  + Highest IP address of a physical interface
  + Can be configured using the `router-id <router-id>` command under the routing instance

Cisco Router-ID Configuration

+ Configuring OSPF router-id
    + Router(config-router)# router-id <router-id>
+ Verifying OSPF router-id
    + Router# show ip ospf
+ Both the router-id and interface priority need OSPF instance reset to take effect after they are changed
    + Router# clear ip ospf process

Vyos Router-ID Configuration

+ ## Configuring the Router-ID

```
vyos@vyos# set protocols ospf parameter router-id
Possible completions:
  <x.x.x.x>    Override the default router identifier
  0.0.0.100
```

```
vyos@vyos# commit
[ protocols ospf ]
For this router-id change to take effect, save config and restart ospfd
```

+ ## Verifying the Router-ID

```
vyos@vyos# show protocols ospf parameters router-id
 router-id 0.0.0.100
[edit]
vyos@vyos#
```

Vyos Router-ID Caveats

+ ## Resetting the Router-ID

    + Vyos does not currently have a method to restart the OSPF process.
    + You must either:
        + Restart the entire router or...
        + Disable and then enable all OSPF-enabled interfaces

```
vyos@vyos# set interface ethernet eth1 disable
[edit]
```

```
vyos@vyos# delete interface ethernet eth1 disable
```

Thanks for Watching!

OSPF DR/BDR Election

Topic Overview

+ Purpose Of Designated Routers
+ DR Election Process
+ OSPF Priority & Configuration
+ Verifying OSPF Designated Routers

Designated Router Election

+ OSPF elects a DR (designated router) and a BDR (backup designated router) in broadcast and non-broadcast multi-access networks
+ DR is responsible for sending an update to the neighbors that are received from other neighbors
+ Special multicast address used for sending routing updates to DR/BDR: 224.0.0.6

OSPF Interface Priority

+ All OSPF-enabled interfaces have an OSPF Priority
+ Priority is the first value inspected to determine DR/BDR election
+ OSPF router with the highest priority becomes the DR
    + Default interface priority = 1
+ OSPF router with lower priority than DR becomes BDR
+ Other OSPF routers are known as DROTHERS

Designated Router Election

+ If OSPF priority has not been configured, highest router-id is referenced
+ OSPF router with highest router-id becomes DR
+ OSPF router with lower router-id than DR's becomes the BDR

Modifying OSPF Interface Priority

Cisco

```
Router(config)#int gig 5
Router(config-if)#ip ospf priority 99
```

Vyos

```
vyos@vyos# set interface ethernet eth1 ip ospf priority
Possible completions:
   <0-255>        Priority (default 1)
   1
```

Cisco OSPF DR/BDR Verification

+ Cisco verification commands

```
Router#show ip ospf interface brief
Interface    PID   Area              IP Address/Mask    Cost   State  Nbrs F/C
Gi5          1     0                 9.9.9.1/24         1      DR     1/1
Gi4          1     23                1.1.1.11/24        1      DR     1/1
```

```
Router#show ip ospf interface gigabit5 | include Designated
  Designated Router (ID) 0.0.0.1, Interface address 9.9.9.1
  Backup Designated router (ID) 0.0.0.211, Interface address 9.9.9.100
    Adjacent with neighbor 0.0.0.211  (Backup Designated Router)
```

Vyos OSPF DR/BDR Verification

+ ## Vyos verification commands

```
vyos@vyos:~$ sho ip ospf neighbor

Neighbor ID     Pri State              Dead Time Address          Interface
0.0.0.1          99 Full/DR             36.396s 9.9.9.1            eth1:9.9.9.100
```

```
vyos@vyos:~$ show ip ospf interface eth1 | grep Designated
  Backup Designated Router (ID) 0.0.0.211, Interface Address 9.9.9.100
  Multicast group memberships: OSPFAllRouters OSPFDesignatedRouters
```

Topic Overview

+ Why OSPF Uses Network Types
+ Broadcast Networks
+ Point-To-Point Networks
+ Non-Broadcast Networks
+ Point-To-Multipoint
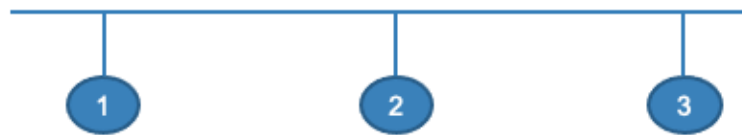+ Influencing OSPF Network Type Determination

Why OSPF Uses Network Types

+ OSPF classifies networks for two reasons:
    + Determines how OSPF can form adjacencies
    + Determines the contents of OSPF LSAs
+ Network type is determined by Layer-2 encapsulation method

Broadcast Networks

+ Broadcast
    + Default setting for Ethernet links
    + Allows for dynamic neighbor discovery
    + Requires election of DR and BDR
    + Timers:
        + Hello interval: 10-seconds
        + Dead interval: 40-seconds

+ Point-to-Point

+ Default setting for HDLC and PPP links
+ Allows for dynamic neighbor discovery
+ Does NOT elect DR or BDR
+ Timers:
    + Hello interval: 10-seconds
    + Dead interval: 40-seconds

Non-Broadcast

+ Non-Broadcast Multi-Access
    + Default setting for Frame-Relay interfaces
    + Requires manual neighbor configuration
    + Requires election DR and BDR
    + Timers:
        + Hello interval: 30-seconds
        + Dead interval: 120-seconds
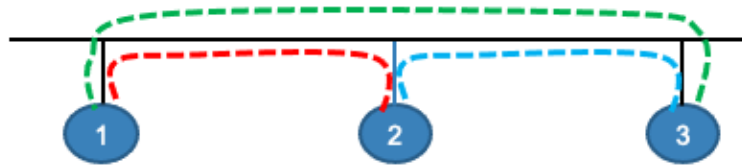
Point To Multipoint

+ Point-To-Multipoint
    + Must be manually configured
    + Allows for dynamic neighbor discovery
    + Typically used in WAN hub-and-spoke environments
    + No requirement for DR or BDR
    + Timers:
        + Hello interval: 30-seconds
        + Dead interval: 120-seconds
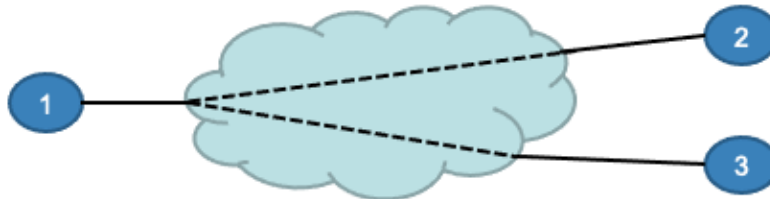
Loopback interfaces are categorized as "Loopback". Virtual-Links also have their own classification.

Network Type Configuration

+ Some Operating Systems allow to you change the default network type on interfaces

+ Network type must match on both sides of the link or OSPF database inconsistencies will occur

```
CSR2(config)#int gig 7
CSR2(config-if)#ip ospf network ?
  broadcast               Specify OSPF broadcast multi-access network
  non-broadcast           Specify OSPF NBMA network
  point-to-multipoint     Specify OSPF point-to-multipoint network
  point-to-point          Specify OSPF point-to-point network
```

```
vyos@vyos# set interface ethernet eth1 ip ospf network
Possible completions:
  broadcast      Broadcast network type
  non-broadcast
                 Non-broadcast network type
  point-to-multipoint
                 Point-to-multipoint network type
  point-to-point
                 Point-to-point network type
```

OSPF Network Type Verification

```
R3#sho ip ospf interface gig0/6
GigabitEthernet0/6 is up, line protocol is up
  Internet Address 23.23.23.2/24, Area 0, Attached via Interface Enable
  Process ID 1, Router ID 31.31.31.2  Network Type BROADCAST  Cost: 1
```

```
vyos@vyos:~$ show ip ospf interface eth1
eth1 is up
  ifindex 3, MTU 1500 bytes, BW 10000 Mbit <UP,BROADCAST,RUNNING,MULTICAST>
  Internet Address 9.9.9.100/24, Broadcast 9.9.9.255, Area 0.0.0.0
  MTU mismatch detection: enabled
  Router ID 0.0.0.211, Network Type POINTOPOINT, Cost: 10
  Transmit Delay is 1 sec, State Point-To-Point, Priority 1
```

Topic Overview

+ OSPF Authentication Overview
+ Authentication Options
+ Authentication Configuration
+ Verifying OSPF Authentication

OSPF Authentication Overview

+ Like other routing protocols, OSPF supports authentication.

+ Configuring OSPF v2 authentication is two-step process:

  + Authentication type and key must be configured

  + Authentication must be activated per-neighbor

```
▶ Frame 3: 114 bytes on wire (912 bits), 114 bytes captured (912 bits) on interface 0
▶ Ethernet II, Src: Vmware_a5:f9:25 (00:50:56:a5:f9:25), Dst: IPv4mcast_05 (01:00:5e:00:00:05)
▶ Internet Protocol Version 4, Src: 9.9.9.1, Dst: 224.0.0.5
▼ Open Shortest Path First
  ▼ OSPF Header
      Version: 2
      Message Type: Hello Packet (1)
      Packet Length: 48
      Source OSPF Router: 0.0.0.1
      Area ID: 0.0.0.0 (Backbone)
      Checksum: 0xc6ed [correct]
      Auth Type: Null (0)
      Auth Data (none): 0000000000000000
  ▶ OSPF Hello Packet
  ▶ OSPF LLS Data Block
```

Authentication Options

+ OSPFv2 supports three types of authentication
  + Plaintext (not recommended)
  + MD5
  + HMAC with SHA (OS-dependent)
+ Plaintext or MD5 can be configured within interface config
+ HMAC with SHA requires configuration of a Key Chain
+ Area-wide authentication also supported, but only with MD5

Plain-Text Authentication

```
▶ Ethernet II, Src: Vmware_a5:f9:25 (00:50:56:a5:f9:25), Dst: IPv4mcast_05 (01:00:5e:00:00:05)
▶ Internet Protocol Version 4, Src: 9.9.9.1, Dst: 224.0.0.5
▼ Open Shortest Path First
  ▼ OSPF Header
        Version: 2
        Message Type: Hello Packet (1)
        Packet Length: 48
        Source OSPF Router: 0.0.0.1
        Area ID: 0.0.0.0 (Backbone)
        Checksum: 0xc6ec [correct]
        Auth Type: Simple password (1)
        Auth Data (Simple): INE
  ▶ OSPF Hello Packet
  ▶ OSPF LLS Data Block
```

```
0000   01 00 5e 00 00 05 00 50   56 a5 f9 25 08 00 45 c0    ··^····P  V··%··E·
0010   00 64 12 95 00 00 01 59   b3 dd 09 09 09 01 e0 00    ·d·····Y  ········
0020   00 05 02 01 00 30 00 00   00 01 00 00 00 00 c6 ec    ·····0··  ········
0030   00 01 49 4e 45 00 00 00   00 00 ff ff ff 00 00 0a    ··INE···  ········
0040   12 63 00 00 00 28 09 09   09 01 09 09 09 64 00 00    ·c···(··  ·····d··
0050   00 d3 7f b4 00 08 00 01   00 04 00 00 00 01 00 12    ········  ········
0060   00 04 00 00 00 0b 80 00   00 08 00 00 00 09 00 00    ········  ··
0070   00 0b                                                ··
```

MD5 Authentication

```
▶ Ethernet II, Src: Vmware_a5:f9:25 (00:50:56:a5:f9:25), Dst: IPv4mcast_05 (01:00:5e:00:00:05)
▶ Internet Protocol Version 4, Src: 9.9.9.1, Dst: 224.0.0.5
▼ Open Shortest Path First
   ▼ OSPF Header
        Version: 2
        Message Type: Hello Packet (1)
        Packet Length: 44
        Source OSPF Router: 0.0.0.1
        Area ID: 0.0.0.0 (Backbone)
        Checksum: 0x0000 (None)
        Auth Type: Cryptographic (2)
        Auth Crypt Key id: 1
        Auth Crypt Data Length: 16
        Auth Crypt Sequence Number: 1601376441
        Auth Crypt Data: 5fb59bc402a4fc6fb2ee0254629e2985
   ▶ OSPF Hello Packet
   ▶ OSPF LLS Data Block
```

Cisco OSPF Authentication Configuration

- **Plain Text Authentication**

  ```
  interface FastEthernet0/0
    ip ospf authentication
    ip ospf authentication-key Cisco
  ```

- **MD5 Authentication (classic)**

  ```
  interface Fastethernet0/0
    ip ospf authentication message-digest
    ip ospf message-digest-key 1 md5 Cisco
  ```

Vyos Authentication Configuration

```
vyos@vyos# set interface ethernet eth1 ip ospf authentication
Possible completions:
 > md5          MD5 parameters
   plaintext-password
                Plain text password
```

Plain-Text Authentication

```
 yos@vyos# set interface ethernet eth1 ip ospf authentication plaintext-password
Possible completions:
   <text>        Plain text password (8 characters or less)
```

MD5 Authentication

```
vyos@vyos# set interface ethernet eth1 ip ospf authentication md5 key-id 1 md5-key INE
```

Cisco HMAC-SHA Authentication

+ HMAC authentication requires configuration of a key chain

```
R2(config)#key chain OSPF-Demo
R2(config-keychain)#key 1
R2(config-keychain-key)#key-string INE
R2(config-keychain-key)#cryptographic-algorithm ?
  hmac-sha-1     HMAC-SHA-1 authentication algorithm
  hmac-sha-256   HMAC-SHA-256 authentication algorithm
  hmac-sha-384   HMAC-SHA-384 authentication algorithm
  hmac-sha-512   HMAC-SHA-512 authentication algorithm
  md5            MD5 authentication algorithm
```

Notice that a key-chain can also be used to activate MD5 authentication.

Applying HMAC Authentication (Cisco)

+ Key-Chain authentication must be activated within the interface

```
R2(config)#interface gigabit0/0
R2(config-if)#ip ospf authentication key-chain ?
  WORD  Key-chain name

R2(config-if)#ip ospf authentication key-chain OSPF-Demo
```

Verifying Authentication (Cisco)

```
Router#sh run interface gigabit5

Current configuration : 232 bytes
!
interface GigabitEthernet5
 ip address 9.9.9.1 255.255.255.0
 ip ospf authentication message-digest
 ip ospf message-digest-key 1 md5 INE
 ip ospf priority 99
 ip ospf 1 area 0
```

```
Router#show ip ospf interface gigabit5
GigabitEthernet5 is up, line protocol is up
  Internet Address 9.9.9.1/24, Interface ID 11, Area 0
  Attached via Interface Enable
  Process ID 1, Router ID 0.0.0.1, Network Type BROADCAST, Cost: 1
```
```
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 0.0.0.211  (Backup Designated Router)
Suppress hello for 0 neighbor(s)
Cryptographic authentication enabled
  Youngest key id is 1
```

Verifying Authentication (Vyos)

```
vyos@vyos:~$ show configuration
interfaces {
    ethernet eth0 {
        hw-id 00:50:56:a5:80:0a
    }
    ethernet eth1 {
        address 9.9.9.100/24
        hw-id 00:50:56:a5:72:b5
        ip {
            ospf {
                authentication {
                    md5 {
                        key-id 1 {
                            md5-key ****************
                        }
```

Thanks for Watching!

Topic Overview

+ Overview Of OSPF LSAs
+ OSPF Link State Update Format
+ Useful Commands For Viewing LSAs
+ Router LSA Details
+ Network LSA Details

OSPFv2 LSAs

+ LSA = Link State Advertisement
+ Carried within an OSPF Link State Update Packet
+ Different types carry different data
+ Age out after 1-hour...refreshed every 30-minutes.

OSPF LSU Packet Format

Version	Type	Packet Length
OSPF Router ID		
OSPF Area ID		
Checksum		Authentication Type
Authentication Data		
Authentication Data		
LS age	Options	LS type
Link State ID		
Advertising Router		
LS sequence number		
LS checksum		Length
Link State Data		

Useful Cisco IOS & Vyos Commands

+ Show ip ospf database
    + Overview of the entire IPv4 OSPF database
    + Displays LSA headers for each connected area
    + Additional command keywords allow one to view details of LSAs and filter for desired output
        + Show ip ospf database <router | network | summary | asbr-summary | external | nssa-external >
        + Show ip ospf database <adv-router | self-originate>

When using Cisco IOS, one can issue the command, "show ip ospf database adv-router <router-id>" to view all OSPF LSAs in the database that were originated by a particular router.

With Vyos, the "adv-router" keyword is not available until AFTER one has selected an LSA type (such as "router" or "summary").

Router LSA

- + LSA Type-1 (Router LSA)
- + Describes the state of connected links
- + Bits to indicate special capabilities of router.
  - + ABR
  - + ASBR
- + Confined to local area only

Network LSA

- Type-2 (Network LSA)
- Only created by Designated Routers
- Describes:
  - All adjacent neighbors of DR
  - Subnet mask of link
- Confined to local area only

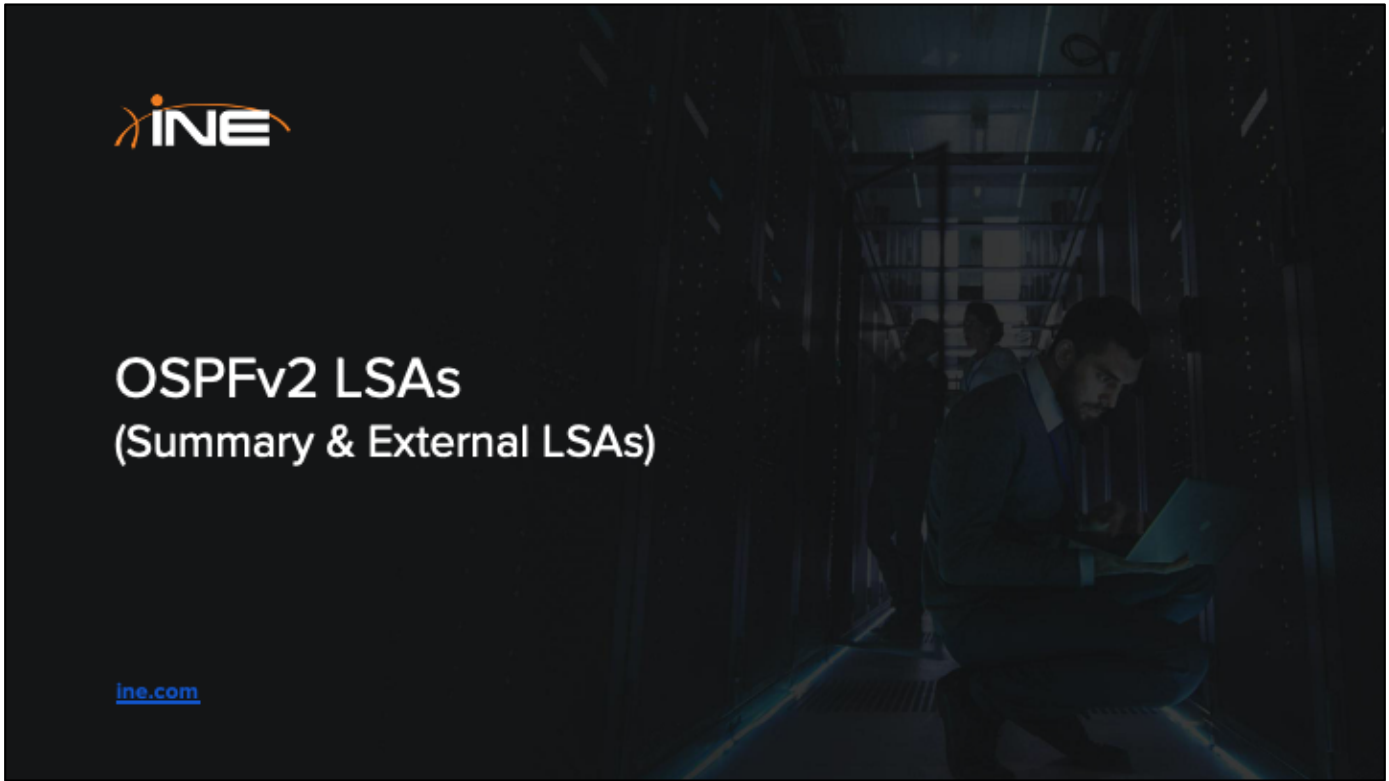**Thanks for Watching!**

OSPFv2 LSAs
(Summary & External LSAs)

Topic Overview

+ OSPF Summary LSAs
+ OSPF External LSAs
+ OSPF ASBR-Summary LSAs

Summary LSA

+ Type-3 (Summary LSA)
+ Describes summarized info of links from one area into an adjacent area
+ Created by ABRs
+ Confined to local area only, but other ABRs may modify and continue to forward.

External LSA

+ Type-5 (External LSA)
+ Advertises non-OSPF routes into OSPF
+ Created by ASBRs
+ Propagated throughout entire OSPF domain.

ASBR Summary LSA

+ Type-4 (ASBR-Summary LSA)
+ Advertises the ASBR into remote areas
+ Created by ABRs
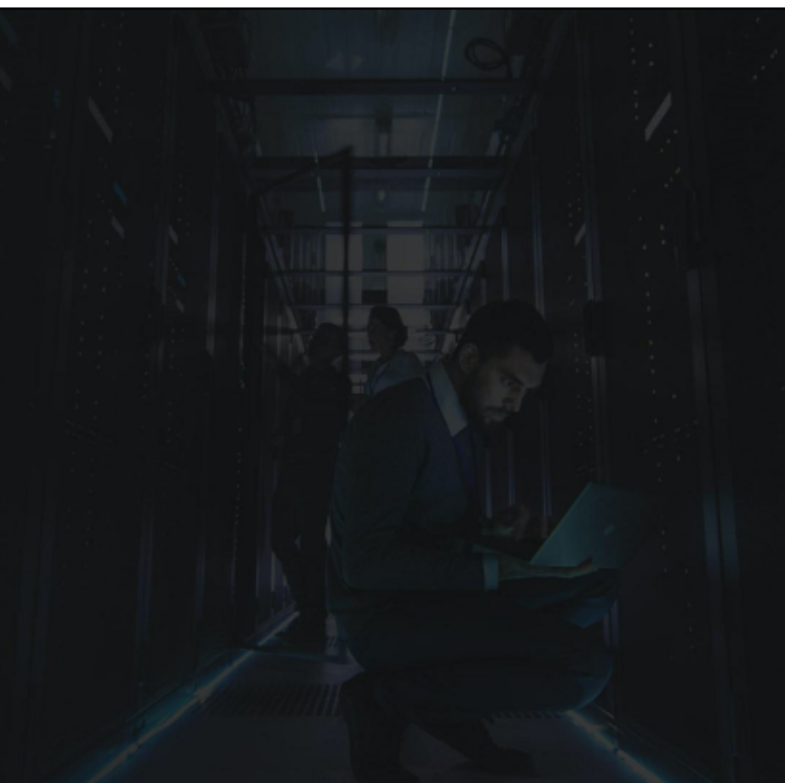+ Allows routers that are not in same area as ASBR to forward traffic to it.

Thanks for Watching!

OSPF Virtual-Links

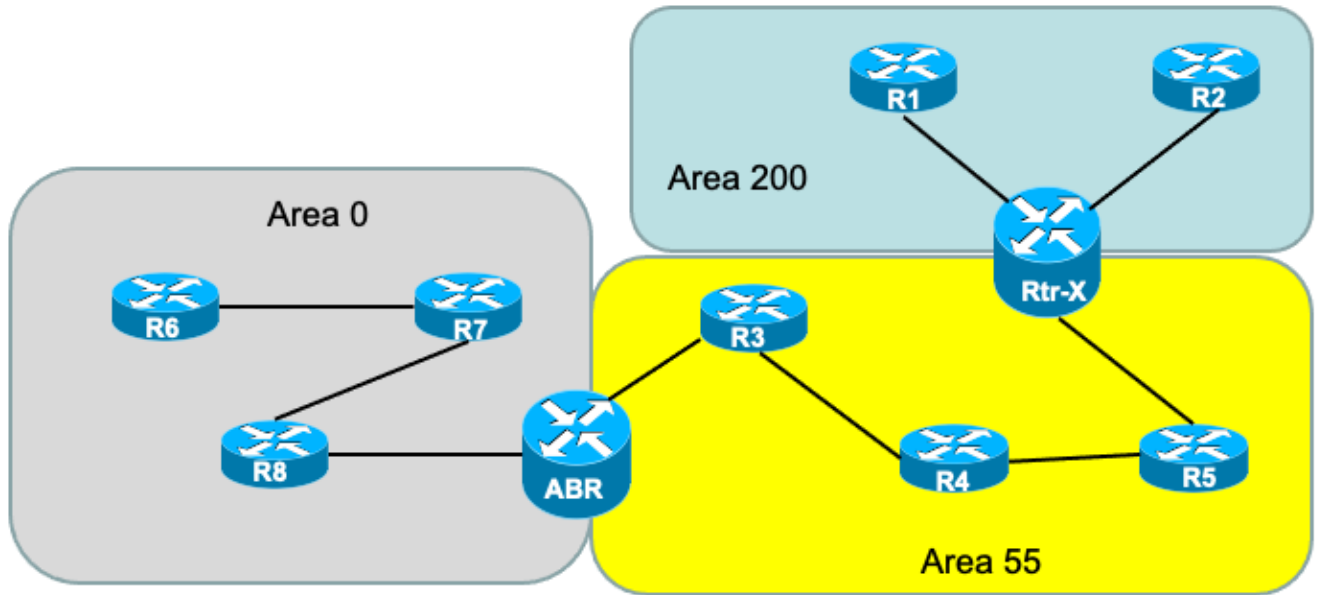ine.com

Topic Overview

- + Purpose Of Virtual Links
- + Virtual Link Configuration
- + Adding Authentication To Virtual Links
- + Monitoring Virtual Links

The Purpose Of Virtual Links

Why Virtual Links?

- + A router may only exchange routes between areas if it serves as an ABR.
- + Virtual Links allow us to convert a router into an ABR that does not physically touch Area-0.
- + A Virtual Link can be thought of as a tunnel, extending from an actual ABR…through a "transit area" and terminating on another router that sits between two non-backbone areas.

Cisco IOS Virtual Link Configuration

+ To configure a virtual link between ABRs, use the following command:
  + **area** area-number **virtual-link** remote-RID (OSPF subcommand)
+ The *area-number* in the above command refers to the transit area number.
+ The remote-RID can be verified from the local router using the command:
  + **show ip ospf border-routers**

Vyos Virtual Link Configuration

+ To configure a virtual link between ABRs, use the following command:
    + **Set protocols ospf area <x> virtual-link <router-id>**
+ The *area-number* in the above command refers to the transit area number.
+ The remote-RID can be verified from the local router using the command:
    + **show ip ospf border-routers**

Virtual Links With Authentication (Cisco)

+ To configure authentication on virtual links, use the following commands:

  + **area** area-num **virtual-link** remote-RID **authentication authentication-key** key-value

    Plaintext Authentication

  + **area** area-num **virtual-link** remote-RID **authentication message-digest message-digest-key** key-number **md5** key-value

    MD5 Authentication

Virtual Links With Authentication (Vyos)

+ To configure authentication on virtual links, use the following commands:

  + Set protocols ospf area <x> virtual-link <router-id> authentication plaintext-password <*password*>

    Plaintext Authentication

  + Set protocols ospf area <x> virtual-link <router-id> authentication md5 key-id <1-255> md5-key <password>

    MD5 Authentication

Monitoring Virtual Links (Cisco)

```
R6#sho ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address      Interface
0.0.0.22           0   FULL/  -            -        2.2.4.22     OSPF_VL9
```

```
R6#show ip ospf virtual-link
Virtual Link OSPF VL9 to router 0.0.0.22 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 2, via interface GigabitEthernet0/1
 Topology-MTID    Cost    Disabled    Shutdown      Topology Name
      0             2         no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
    Adjacency State FULL (Hello suppressed)
```

Monitoring Virtual Links (Vyos)

```
vyos@vyos:~$ sho ip ospf neighbor

Neighbor ID     Pri State            Dead Time Address        Interface
2.2.2.2           1 Full/Backup       31.326s 10.77.77.22      eth2:10.77.77.2
0.0.0.1           1 Full/Backup       30.573s 10.44.44.1       eth1:10.44.44.2
0.0.0.1           1 Full/DROther      34.473s 10.44.44.1       VLINK0
```

```
vyos@vyos:~$ show ip ospf interface
VLINK0 is up
  ifIndex 0, MTU 1500 bytes, BW 0 Mbit <UP>
  Internet Address 10.44.44.2/24, Peer 10.44.44.1, Area 0.0.0.0
  MTU mismatch detection: enabled
  Router ID 0.0.0.2, Network Type VIRTUALLINK, Cost: 10
  Transmit Delay is 1 sec, State Point-To-Point, Priority 1
  No backup designated router on this network
  Multicast group memberships: <None>
  Timer intervals configured, Hello 10s, Dead 40s, Wait 40s, Retransmit 5
    Hello due in 6.305s
  Neighbor Count is 1, Adjacent neighbor count is 1
```

Influencing OSPF Path Selection

ine.com

Topic Overview

+ OSPF Cost
+ OSPF Route Selection
+ Controlling ECMP
+ Influencing The OSPF Metric Algorithm
+ Influencing OSPF Cost
+ Originating Default Routes

OSPF Cost

+ OSPF metric = Cost

+ Cost inversely proportional to interface bandwidth
    + The greater the bandwidth the lower the cost
    + Path to remote destination with lowest aggregate cost is considered best

+ Default OSPF cost formula:
    + 100,000,000/Bw (bps)

OSPF Route Selection

+ **Important Note**:
    + When multiple routes to the same prefix are learned by OSPF, OSPF always considers the route-type prior to metric comparison
+ Intra-Area > Inter-Area > External
+ External route selection
    + O E1's/O N1s are preferred over O E2's/ O N2's
    + If both are 1's (or both are 2's) then lowest cost wins
        + O N1 (cost 35) > O E1 (cost 55)
    + If cost is the same, O N's > O E's
        + O N1 (cost 35) > O E1 (cost 35)
        + O N2 (cost 35) > O E2 (cost 35)

Controlling ECMP

+ ECMP = Equal Cost Multi-Path (routing)
+ OSPF (like most other Routing Protocols) can potentially learn multiple, equal-cost paths to a prefix
+ Default is up to four (4) allowed into IP Routing Table (most platforms)
    + Router(config-router)#**maximum-paths 6**

Influencing The OSPF Metric Algorithm

+ Default OSPF metric calculation does not differentiate between links above FastEthernet
+ For networks with FastEthernet (and above) this can lead to inferior path selection

```
Router(config-router)#auto-cost reference-bandwidth ?
  <1-4294967>   The reference bandwidth in terms of Mbits per second
```

```
vyos@Vyos-1# set protocol ospf auto-cost reference-bandwidth
Possible completions:
   <1-4294967>   Reference bandwidth cost in Mbits/sec (default 100)
```

+ Command should be implemented consistently across all OSPF routers.

Influencing OSPF Cost Values

+ OSPF path selection can be artificially influenced by modifying cost values

+ Two ways to accomplish this:

  + Change bandwidth value on an interface

```
Router(config)#int fast 0/0
Router(config-if)#bandwidth ?
  <1-10000000>   Bandwidth in kilobits
```

```
vyos@Vyos-1# set interface ethernet eth1 ip ospf bandwidth
Possible completions:
  <1-100000>   Bandwidth in megabits/sec (for calculating OSPF cost)
```

  + Set OSPF cost value directly on an interface

```
Router(config-if)#ip ospf cost ?
  <1-65535>   Cost
```

```
vyos@Vyos-1# set interface ethernet eth1 ip ospf cost
Possible completions:
  <1-65535>    OSPF interface cost
```

Originating Default Routes

+ OSPF does not create a default route the same way as other routing protocols.
+ A special OSPF command is used for this purpose:

```
Router(config)#router ospf 1
Router(config-router)#default-information originate ?
  always        Always advertise default route
```

```
vyos@Vyos-1# set protocol ospf default-information originate
Possible completions:
  always        Always advertise default route
  metric        OSPF default metric
  metric-type   OSPF metric type for default routes
  route-map     Route map reference
```

ip route 0.0.0.0 0.0.0.0 1.1.1.1

Internet

1.1.1.1
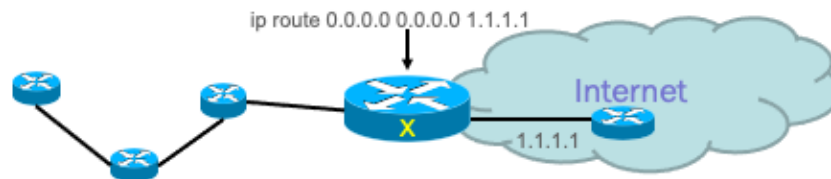
Inter-Area OSPF Route Summarization

Topic Overview

+ The Benefits Of OSPF Summarization
+ Inter-Area Summarization Configuration

Summarization & Areas

+ ## Why segment your OSPF domain into multiple areas?
    + Answer-1: To minimize SPF runs due to churning of Type-1 and Type-2 LSAs
    + Answer-2: To provide the ability to selectively filter and/or summarize routes between different regions of the topology

+ ## Based on the topology shown here do all OSPF routers really NEED to know about all prefixes?
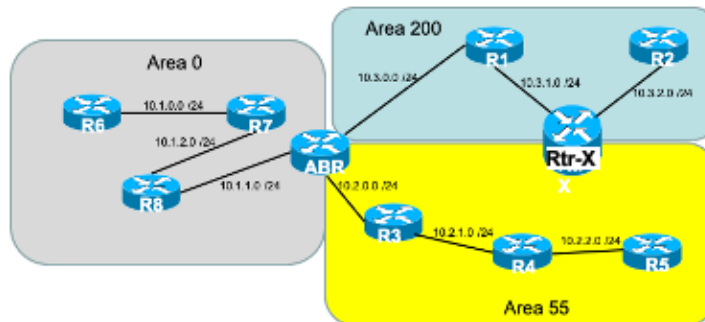
OSPF Route Summarization Overview

+ OSPF allows summarization at ABRs and ASBRs only.
+ The reason is same as with route filtering: All routers in an area must have the same LSDB.
+ Summarization is done for LSAs, not routes.
+ Summarization can be done to reduce the size of the LSDB, save memory, and save CPU resources.
+ Summarization can also be used for path manipulation.

OSPF Interarea Summarization Configuration

+ Area referenced after "range" must be the area where the specific prefixes reside

    + Learned by ABR via Type-1 or Type-2 LSAs

Cisco

```
CSR1(config-router)#area 1 range 10.1.0.0 ?
  A.B.C.D  IP mask for address

CSR1(config-router)#area 1 range 10.1.0.0 255.255.252.0
```

Vyos

```
vyos@Vyos-1# set protocol ospf area 1 range
Possible completions:
 > <x.x.x.x/x>  Summarize routes matching prefix (border routers only)
 >

[edit]
vyos@Vyos-1# set protocol ospf area 1 range 10.2.0.0/22
```

OSPF Area Range Command

+ The ABR assigns a metric for the summary route's Type 3 LSA, by default, to match the best metric among all subordinate subnets.
+ The `area range` command can also explicitly set the cost of the summary.
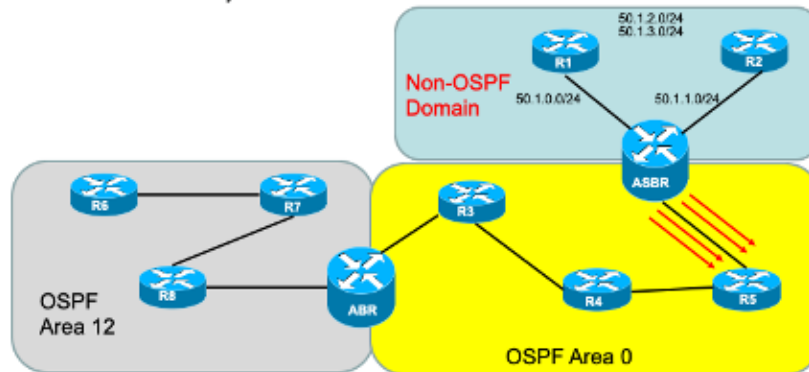+ The summary route created is a Type 3 route.

Topic Overview

+ Summarizing OSPF External Routes
+ Configuration

Summarizing External Routes

+ OSPF can convert non-OSPF routes into OSPF using redistribution on ASBRs

+ Each external route is advertised within a Type-5 or Type-7 External LSA

+ In some cases, summarization can be useful on ASBRs

OSPF External Route Summarization

+ To configure summarization at ASBR, use the following command:
    + **summary-address** prefix mask (OSPF subcommand)
+ Redistribution must be configured first
+ With this, a Type 5 LSA (or Type-7) is created for the summary route.
+ Other vendors (Vyos, Juniper, etc) do not support summarization of external routes.

OSPF Inter-Area LSA Suppression

Topic Overview
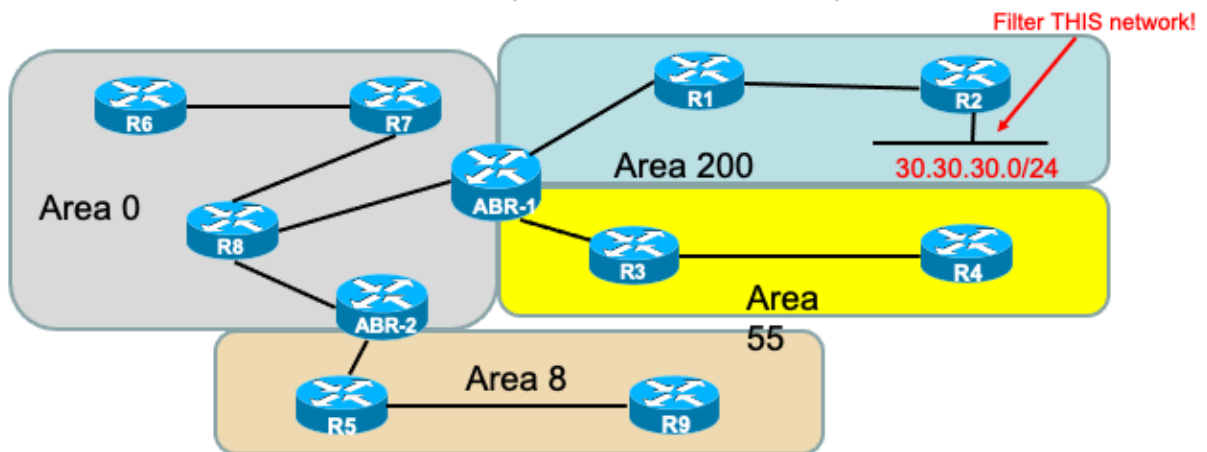
+ Filtering/Suppressing Inter-Area OSPF LSAs
+ Rules & Guidelines
+ Selective Directional Summarization

LSA Suppression Rules

+ The routers inside an area should have the same LSAs in their LSDB.
+ If all routers inside an area do not have the same LSDB, SPF logic will fail.
+ OSPF can prevent the origination of certain types of LSAs at area boundaries
+ Type-3 LSAs are suppressed at ABRs; Type-5 LSAs are suppressed at ASBRs.

Inter-Area Filtering Choices

+ Two methods exist to filter/suppress OSPF inter-area routes.
+ The method selected depends on what you need:

Filter THIS network!

Area 0

Area 200

R6 — R7

R1 — R2

30.30.30.0/24

ABR-1

R8

R3 — R4

Area 55

ABR-2

Area 8

R5 — R9

Filtering On Originating ABR

```
CSR1(config-router)#area 200 range 30.30.30.0 255.255.255.0 not-advertise
vyos@Vyos-1# set protocol ospf area 200 range 30.30.30.0/24 not-advertise
```

Area 0

Area 200

30.30.30.0/24

Area 55

Area 8
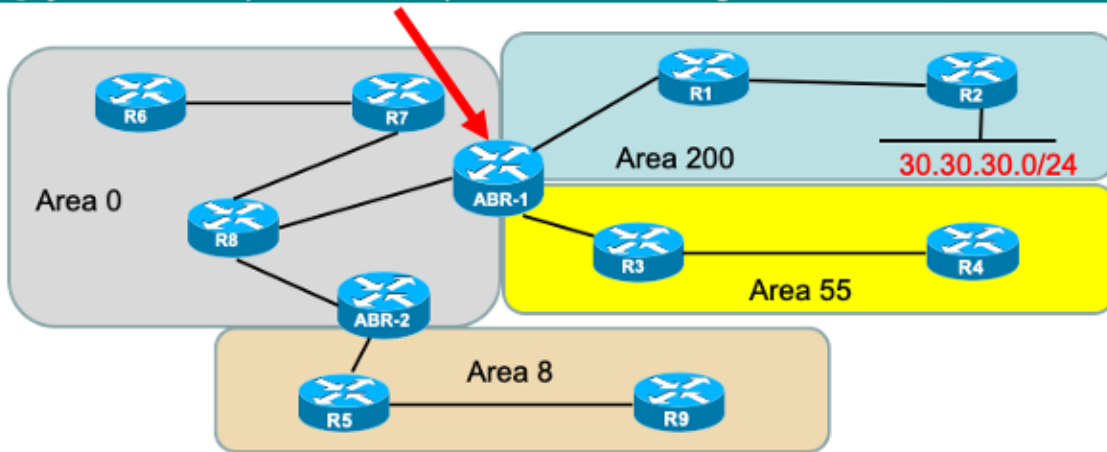
Selective Directional Summary Filtering

+ With Cisco routers, another command exists that allows selective, directional filtering of Summary LSAs on ABRs
+ Can be implemented on any ABR
+ To do this, one must first know how to construct a Prefix-List:
    + **ip prefix-list \<name> \<permit | deny> [prefix/mask]**
    + **area \<number> filter-list prefix \<name> in | out**

Selective Transmission Of Filtered Summary LSAs

```
CSR1(config)#ip prefix-list No-30 deny 30.30.30.0/24
CSR1(config)#ip prefix-list No-30 permit 0.0.0.0/0 le 32
CSR1(config)#router ospf 1
CSR1(config-router)#area 55 filter-list prefix No-30 in
```

R6

R7

Area 0

R8

ABR-1

Area 200

R1

R2

30.30.30.0/24

R3

R4

Area 55

ABR-2

Area 8

R5

R9
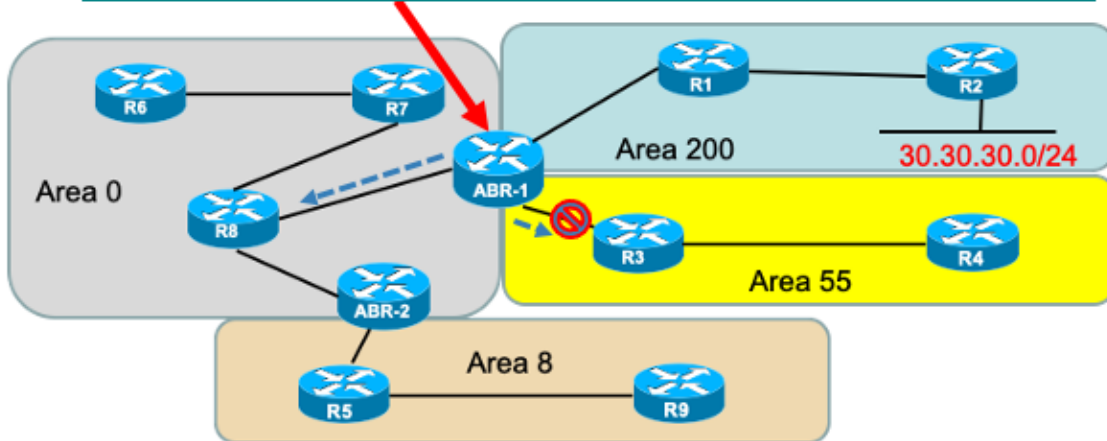
Thanks for Watching!

OSPF Selective External Redistribution

ine.com

Topic Overview

+ Cisco OSPF External Route Filtering
+ Using Summarization For Filtering
+ Vyos External Route Filtering

Cisco OSPF External Route Filtering

+ Type-5 and Type-7 LSAs are filtered *prior to origination* on ASBRs.
+ Type-5 or Type-7 LSAs cannot be filtered after the ASBR has created-and-flooded them.
+ To configure Type-5 or Type-7 filtering on ASBRs, we use the "distribute-list out" command:
    + **distribute-list** <[ACL#] | Prefix [name] | route-map [name]> **out** <protocol><ASN | process-id>

Using Cisco Summarization For Filtering

+ Translation of Type-7 to Type-5 LSAs can be controlled on the NSSA ABR

+ In the example below, ABR is receiving Type-7 LSAs for 1.x.x.x and 2.x.x.x prefixes and we wish to only translate those for 1.x.x.x

(On ABR)
Router ospf 1
  network x.x.x.x area 0
  network y.y.y.y area 3
  area 3 nssa
  summary-address 2.0.0.0 255.0.0.0 not-advertise

Area 0

Type-5: 1.0.0.0/8

ABR

Type-7: 1.0.0.0/8   ASBR

Area 3
NSSA

Type-7: 2.0.0.0/8

Vyos External Route Filtering

+ Vyos filters external routes that are matched against "deny" statements in Access-Lists

```
set policy access-list 1 rule 10 action 'deny'
set policy access-list 1 rule 10 source inverse-mask '0.0.0.255'
set policy access-list 1 rule 10 source network '10.55.55.0'
set policy access-list 1 rule 20 action 'permit'
set policy access-list 1 rule 20 source any
set protocols ospf access-list 1 export 'rip'
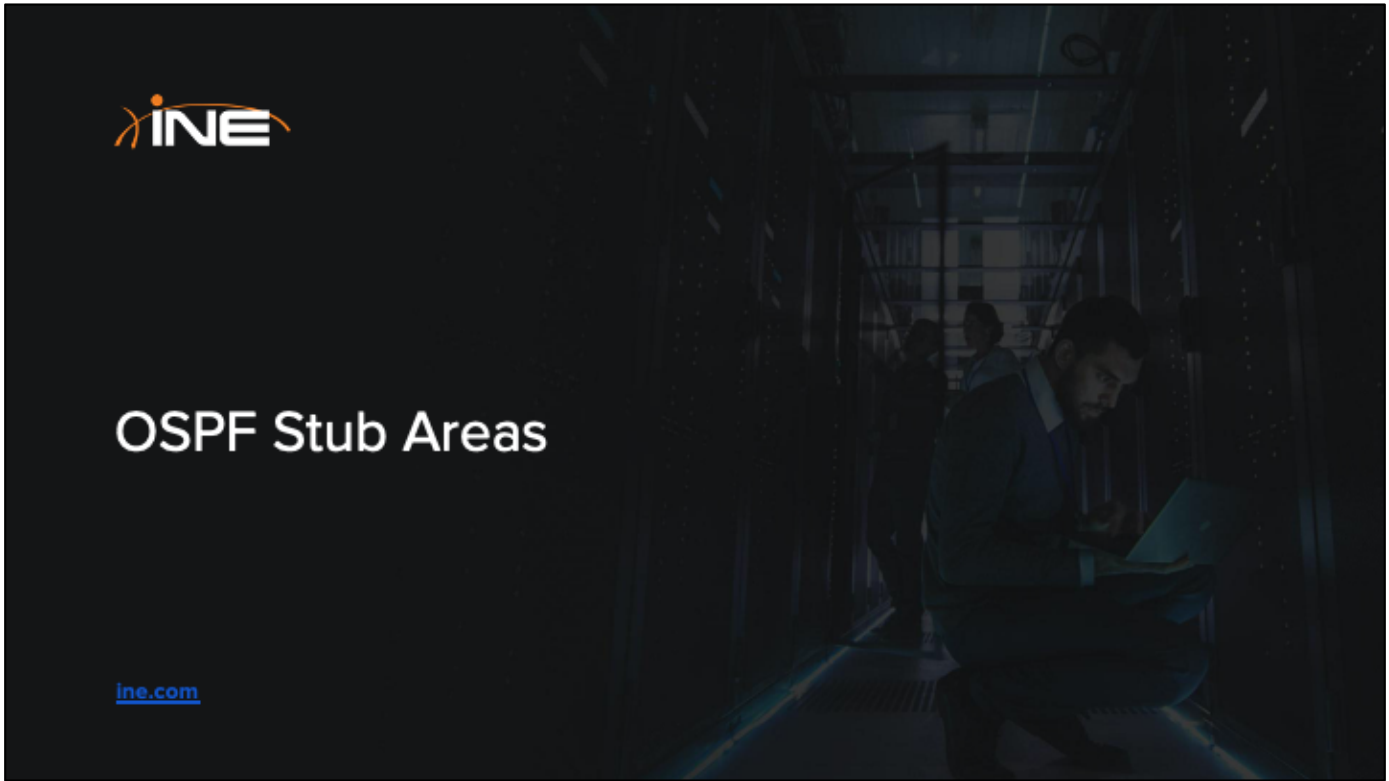```

Non-OSPF protocol that you previously redistributed.

**Thanks for Watching!**

# OSPF Stub Areas

ine.com

## Topic Overview

+ The Problem With Normal OSPF Areas
+ Introduction To Stub Areas
+ Totally Stubby Areas
+ Stub & Totally Stubby Configuration
+ Not-So-Stubby-Areas
+ Totally NSSA

# The Problem With Normal Areas

+ By default, OSPF areas allow all LSAs to be flooded within the area (Type-1 through Type-5)
+ In some cases this can lead to unnecessary bloat of the Link State Database and IP Routing Table
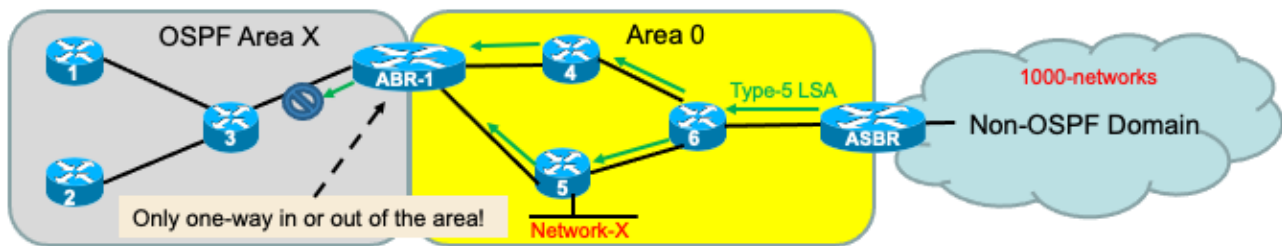    + Does OSPF Area-X really need to know about ALL external routes?
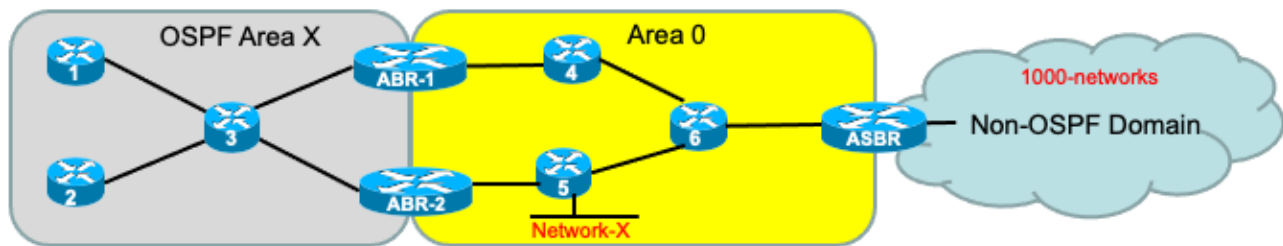
# Introduction To Stub Areas

+ OSPF stub areas prevent ABRs from forwarding External LSAs (Type-5) into an area
+ ABRs instead provide a default route for external reachability
+ OSPF adjacencies across a stub area require the setting of the "stub" flag in OSPF Hellos
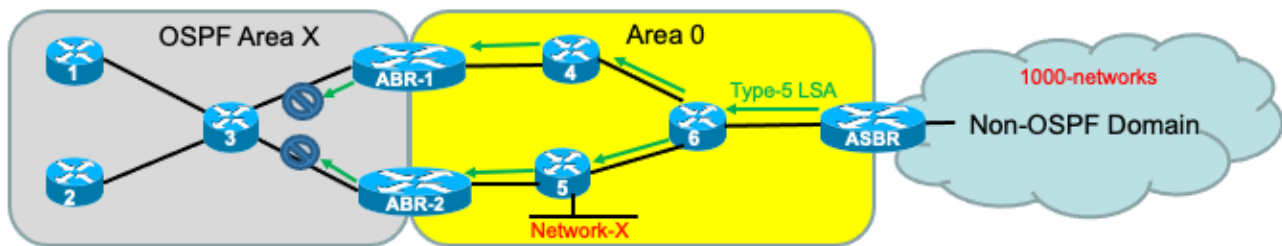
# Totally Stubby Areas

+ Some OSPF areas can further restrict importing Type-3 Summary LSAs
+ This is called a Totally Stubby Area
    + ABR will still generate a default route
    + Still requires OSPF adjacencies to set the "stub" flag in OSPF Hellos
    + ABR requires additional keyword to suppress Summary LSAs



OSPF Area X — Area 0 — Type-5 LSA — 1000-networks — Non-OSPF Domain — ABR-1 — ASBR — Only one-way in or out of the area! — Network-X

# Stub Area Configuration

```
CSR1(config)#router ospf 1
CSR1(config-router)#area 1 stub
```

```
CSR1#sho ip ospf
 Routing Process "ospf 1" with ID 0.0.0.11
<Output omitted for brevity>
    Area 1
        Number of interfaces in this area is 1
        It is a stub area
```

Cisco configuration & verification

```
vyos@Vyos-1# set protocol ospf area 1 area-type stub
```

```
vyos@Vyos-1:~$ show ip ospf
 OSPF Routing Process, Router ID: 0.0.0.1
```

```
Area ID: 0.0.0.1 (Stub)
   Shortcutting mode: Default, S-bit consensus: no
   Number of interfaces in this area: Total: 1, Active: 1
```

Vyos configuration & verification

Totally Stubby Area Configuration

```
CSR1(config)#router ospf 1
CSR1(config-router)#area 1 stub no-summary
```

```
CSR1#sho ip ospf
 Routing Process "ospf 1" with ID 0.0.0.11
```
<Output omitted for brevity>
```
 Area 1
     Number of interfaces in this area is 1
     It is a stub area, no summary LSA in this area
     Generates stub default route with cost 1
```

Cisco configuration & verification

```
vyos@Vyos-1# set protocol ospf area 1 area-type stub no-summary
```

```
vyos@Vyos-1:~$ show ip ospf
 OSPF Routing Process, Router ID: 0.0.0.1
```
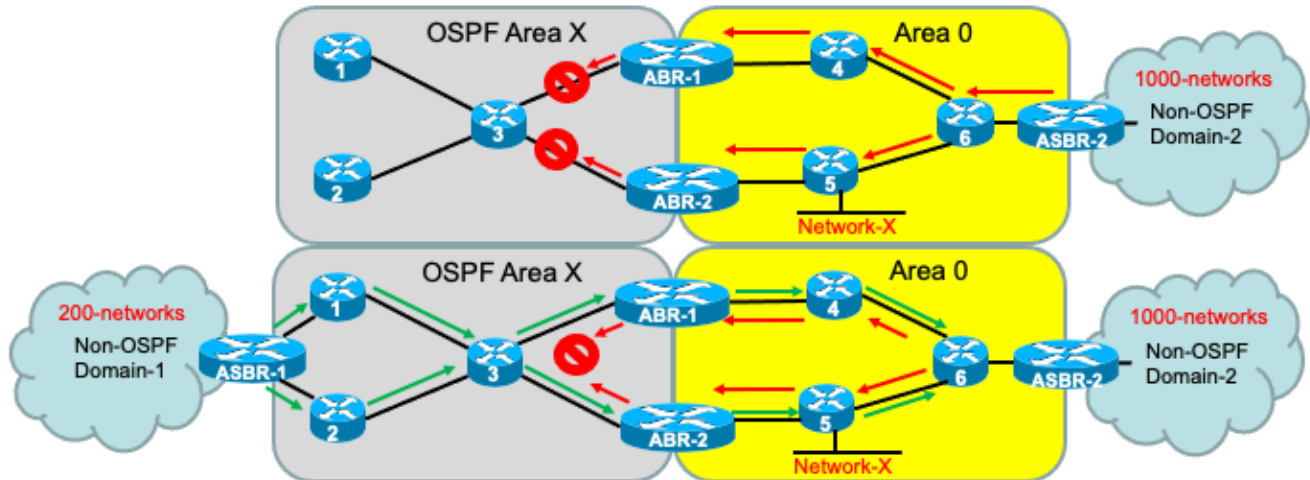<Output omitted for brevity>
```
Area ID: 0.0.0.1 (Stub, no summary)
   Shortcutting mode: Default, S-bit consensus: no
   Number of interfaces in this area: Total: 1, Active: 1
```

Vyos configuration & verification

Not So Stubby Areas

+ In some cases, an OSPF area requires a local ASBR but you still want to prevent external LSAs from remote ASBRs
+ This type of area must be defined as NSSA

NSSA Area Caveats

+ Local ASBRs in NSSA generate Type-7 NSSA-External LSAs
  + ABRs translate these into Type-5 LSAs for flooding into adjacent, normal OSPF areas
  + If multiple ABRs, only highest Router-ID elected to translate Type-7's into Type-5s
+ Unlike Stub and Totally Stubby areas, sometimes ABRs connected to NSSAs do NOT generate a default-route.
  + Cisco ABRs do not generate default-route, but Vyos routers DO
  + Default-route can be generated with additional commands on Cisco ABRs

NSSA Configuration

Cisco configuration

```
CSR1(config)#router ospf 1
CSR1(config-router)#area 1 nssa
```

```
CSR1(config-router)#area 1 nssa default-information-originate
```
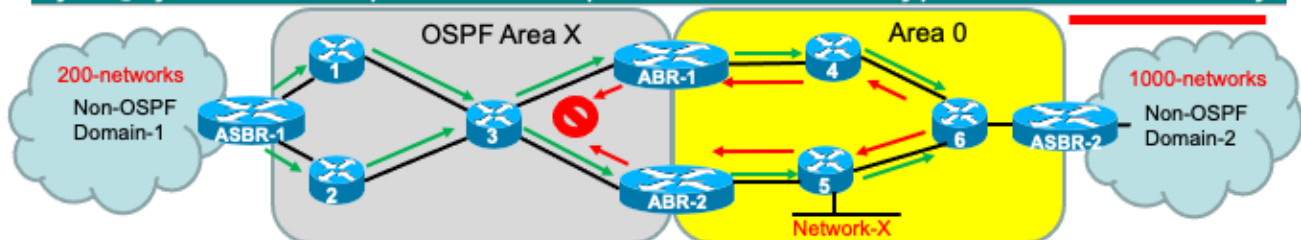
Vyos configuration

```
vyos@Vyos-1# set protocol ospf area 1 area-type nssa
```

Totally NSSA

+ Not So Stubby Areas can also further restrict importation of Type-3 Summary LSAs

+ This becomes a "Totally NSSA".

  + ABRs generate default-route (by default)

+ Requires additional keywords on local ABRs

```
CSR1(config)#router ospf 1
CSR1(config-router)#area 1 nssa no-summary
```

```
vyos@Vyos-1# set protocol ospf area 1 area-type nssa no-summary
```


OSPF Area X

200-networks
Non-OSPF
Domain-1

ASBR-1

1
3
2

ABR-1
ABR-2

Area 0

4
6
5

Network-X

1000-networks
Non-OSPF
Domain-2

ASBR-2

Thanks for Watching!

Changing OSPF Administrative Distance

ine.com

Topic Overview

- Overview Of Administrative Distance
- Default Values Of Admin Distance
- Viewing Admin Distance
- Modifying OSPF Administrative Distance

What Is Administrative Distance

+ Administrative Distance (AD)
  + Defines trustworthiness of a routing protocol
  + 8-bit numbering system
  + Ranges from 0 through 255

Default Administrative Distance Values

Protocols	AD Value
Connected	0
Static	1
EIGRP ← Cisco Only →	90 (EIGRP External = 170)
OSPF	110
IS-IS	115
RIP	120
iBGP/eBGP	200/20
Unreachable	255

Viewing Administrative Distance

+ ## Commands exist to allow you to view the AD values of different protocols

    + Show [ip|ipv6] protocols (Cisco)
    + Show [ip|ipv6] route (Cisco/Vyos)

```
CSR1#sho ip protocols | i Protocol|Distance
Routing Protocol is "application"
    Gateway          Distance        Last Update
  Distance: (default is 4)
Routing Protocol is "ospf 1"
    Gateway          Distance        Last Update
  Distance: (default is 110)
Routing Protocol is "bgp 134"
    Gateway          Distance        Last Update
  Distance: external 20 internal 200 local 200
Routing Protocol is "eigrp 100"
  EIGRP-IPv4 Protocol for AS(100)
     Distance: internal 90 external 170
    Gateway          Distance        Last Update
  Distance: internal 90 external 170
```

Modifying Administrative Distance

+ There are times you may need to modify the default AD value within a routing protocol
    + Prevent routing loops
    + Change the preference of routing protocols
+ Different methods, and restrictions, exist within different protocols
+ Remember that you are only changing a locally-significant value

Changing OSPF Administrative Distance (Cisco)

+ Modifying the AD of selective OSPF internal, inter-area and/or external routes:

```
R2(config)#router ospf 1
R2(config-router)#distance 180 192.168.1.253 0.0.0.0 1
```

+ Modifying the AD of all OSPF internal, inter-area and/or external routes:

```
R2(config-router)#distance ospf external 100 inter-area 120 intra-area 140
```

The first command will work against any kind of OSPF route regardless of its type, as long as it is matched by the ACL.
-
Both commands can exist simultaneously and, if so, the command that applies AD against selective OSPF routes (matched by an ACL) will take precedence.

Changing OSPF Administrative Distance (Vyos)

+ Modifying the AD of all OSPF routes:

```
set protocols ospf distance global '80'
```

+ Modifying the AD of certain classes of OSPF routes:

```
vyos@Vyos-1# set protocol ospf distance ospf
Possible completions:
   external     Distance for external routes
   inter-area   Distance for inter-area routes
   intra-area   Distance for intra-area routes
```

```
set protocols ospf distance ospf inter-area '60'
```

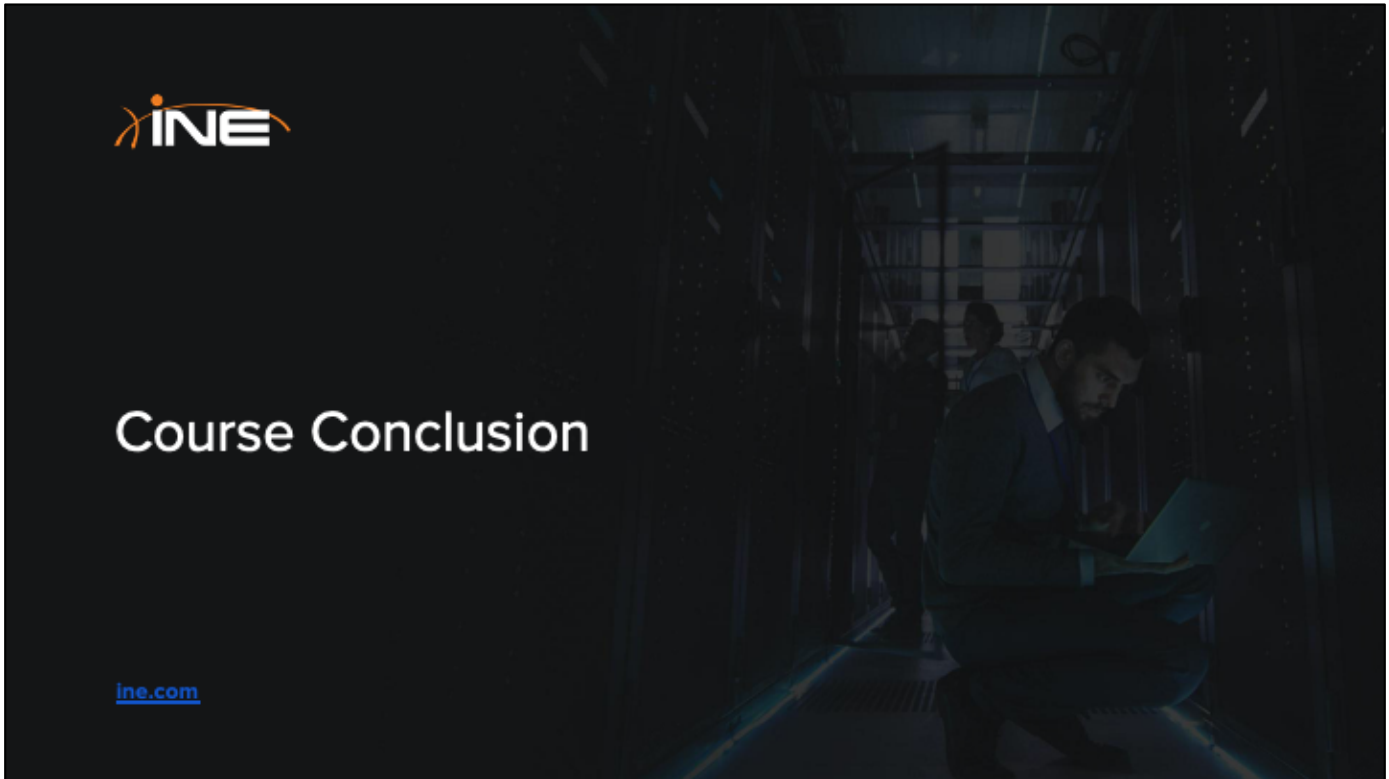Vyos does not allow you to change the AD value of selective routes matched via an Acces-List or Prefix-List.

Thanks for Watching!

Course Conclusion

ine.com

Thanks for Watching!