360RSW05-WKB-TA03

Cisco 360 CCIE R&S Exercise Workbook Introduction

The Cisco 360 CCIE® R&S Exercise Workbook contains 20 challenging scenarios at the Cisco CCIE level that can be used for rigorous self-paced practice.

Each lab provides an extensive answer key, Mentor Guide support, and verification tables and is designed to maximize learning by providing practical experience. Also, self-paced learning resources such as the Cisco 360 CCIE R&S Reference Library and Cisco 360 CCIE R&S lessons supplement the Exercise Workbook scenarios.

https://t.me/learningnets

Cisco 360 CCIE R&S Exercise Workbook Lab 3 Troubleshooting Section Answer Key

COPYRIGHT 2013, CISCO SYSTEMS, INC. ALL RIGHTS RESERVED. ALL CONTENT AND MATERIALS, INCLUDING WITHOUT LIMITATION, RECORDINGS, COURSE MATERIALS, HANDOUTS AND PRESENTATIONS AVAILABLE ON THIS PAGE, ARE PROTECTED BY COPYRIGHT LAWS. THESE MATERIALS ARE LICENSED EXCLUSIVELY TO REGISTERED STUDENTS FOR THEIR INDIVIDUAL PARTICIPATION IN THE SUBJECT COURSE. DOWNLOADING THESE MATERIALS SIGNIFIES YOUR AGREEMENT TO THE FOLLOWING: (1) YOU ARE PERMITTED TO PRINT THESE MATERIALS ONLY ONCE, AND OTHERWISE MAY NOT REPRODUCE THESE MATERIALS IN ANY FORM, OR BY ANY MEANS, WITHOUT PRIOR WRITTEN PERMISSION FROM CISCO; AND (2) YOU ARE NOT PERMITTED TO SAVE ON ANY SYSTEM, MODIFY, DISTRIBUTE, REBROADCAST, PUBLISH, TRANSMIT, SHARE OR CREATE DERIVATIVE WORKS ANY OF THESE MATERIALS. IF YOU ARE NOT A REGISTERED STUDENT THAT HAS ACCEPTED THESE AND OTHER TERMS OUTLINED IN THE STUDENT AGREEMENT OR OTHERWISE AUTHORIZED BY CISCO, YOU ARE NOT AUTHORIZED TO ACCESS THESE MATERIALS.

https://t.me/learningnets

Table of Contents

https://t.me/learningnets

Answer Key Structure

Section One

The answer key PDF document is downloadable from the web portal.

Section Two

To obtain a comprehensive view of the configuration for a specific section, access the Mentor Guide engine in the web portal.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Exercise Workbook Lab 3 Troubleshooting Section Answer Key

Note	Regardless of any configuration you perform in this lab, it is very important that you conform to the general guidelines that are provided in the “Restrictions and Goals” section. If you do not conform to the guidelines, you could have a significant deduction of points in your final score.

Grading and Duration

- Troubleshooting lab duration: 2 hours
- Troubleshooting lab maximum score: 24 points

Note	You can assess your progress on the self-paced labs in this workbook by adding up the points that are assigned to sections and tasks. Consider taking the full Assessment Labs to assess your readiness level.

Difficulty Level

- Difficulty: Intermediate

Restrictions and Goals

Note	Read this section carefully.

- To receive any credit for a subsection, you must fully complete the subsection as per the requirements. You will *not* receive partial credit for partially completed subsections.
- IPv4 subnets displayed in the scenario diagram belong to network 135.15.0.0/16. *Points will be deducted from multiple sections for failing to assign correct IPv4 addresses.*
- Do not use any static routes.
- Advertise loopback interfaces with their original masks.
- Network 0.0.0.0/0 should not appear in any routing table (**show ip route**).
- Do not use the **ip default-gateway** or **ip default-network** commands.
- Do not introduce any new IP addresses.
- All IP addresses involved in this scenario must be reachable, unless explicitly specified otherwise.

https://t.me/learningnets

- Unless explicitly specified otherwise, addresses and networks that are advertised in the Border Gateway Protocol (BGP) section need to be reachable by all BGP routers but do not have to be reachable by routers that use only interior gateway protocol (IGP).
- Use conventional routing algorithms only, unless specified otherwise.
- Do not create new interfaces to fulfill IGP requirements, and do not summarize unless explicitly asked to do so.
- Do not modify the hostname, console, or vty configuration unless you are specifically asked to do so.
- Do not modify the initial interface or IP address numbering.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Explanation of Each of the Restrictions and Goals

IPv4 subnets displayed in the scenario "IPv4 IGP" diagram belong to network 135.15.0.0/16.

All IP addresses in this exam belong to the 135.15.0.0/16 address space, with the exception of prefixes that are explicitly specified as being part of a different IP space.

Do not use any static routes.

Static routes can be used to solve a range of reachability problems. However, you cannot use them in this lab. You must rely on skillful configuration of all your unicast routing protocols.

Advertise loopback interfaces with their original masks.

The original mask is the mask configured on the loopback interface. Open Shortest Path First (OSPF) by default treats loopback interfaces as host routes and advertises them as /32 prefixes. The requirement to advertise loopback interfaces with their original masks precludes using the default OSPF network type for the loopback interfaces.. You need to provide a solution such as changing the OSPF network type or summarizations. Remember that this rule applies to both the IPv4 and IPv6 networks.

Network 0.0.0.0/0 should not appear in any routing table (show ip route**).**

A 0.0.0.0/0 entry can be used to solve a range of reachability problems. In particular, a 0.0.0.0/0 entry can be used to set up the gateway of last resort. In this exercise, you cannot use any 0.0.0.0/0 entries. An alternative to using the 0.0.0.0/0 route to solve the reachability problem is route summarization.

Do not use the ip default-gateway **or** ip default-network **commands.**

These commands can also be used to solve reachability issues by setting a gateway of last resort. They generate a 0.0.0.0/0 route in the Routing Information Protocol (RIP) environment. You cannot use them in this scenario.

All IP addresses involved in this scenario must be reachable.

This goal is a key goal to observe. It requires that all of your IGPs and routing policy tasks must be configured properly. The key elements of your routing policy include route redistribution and the controlling of routing updates using the **distribute-lists**, **route-maps**, and **distance** commands. A key point to remember about this lab is that the term "redistribution" is not explicitly used. However, you must perform redistribution to ensure that all IP addresses are reachable without the use of static routes or 0.0.0.0/0 routes.

https://t.me/learningnets

Addresses and networks that are advertised in the BGP section need to be reachable by all BGP routers but do not have to be reachable by IGP-only routers.

This statement relaxes the requirement that all IP addresses must be reachable. The BGP prefixes need only be reachable among the routers specified in the BGP section. They can be used in other unicast tables. However, BGP routers need to have the prefixes in the routing tables as well as be able to forward traffic to the addresses known via BGP.

Use conventional routing algorithms.

This restriction prevents you from solving any problems by configuring policy routing. At the heart of this restriction is the interpretation of "conventional routing algorithms." Although this phrase can be interpreted in several different ways, the following interpretation is applied in this workbook:

> Conventional routing algorithms are routing algorithms that apply destination-based prefix lookups in a routing table. Conventional routing algorithms do not use any type of information other than the destination address to make a packet-forwarding decision.

Because of this restrictive interpretation, no form of policy routing can be applied. Whenever you see this restriction, you will need to use dynamic routing protocols to fulfill all packet-forwarding requirements.

.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

1. Switched Network Troubleshooting Section

The "VLAN Propagation" diagram is reproduced here for easy reference.

VLAN Propagation

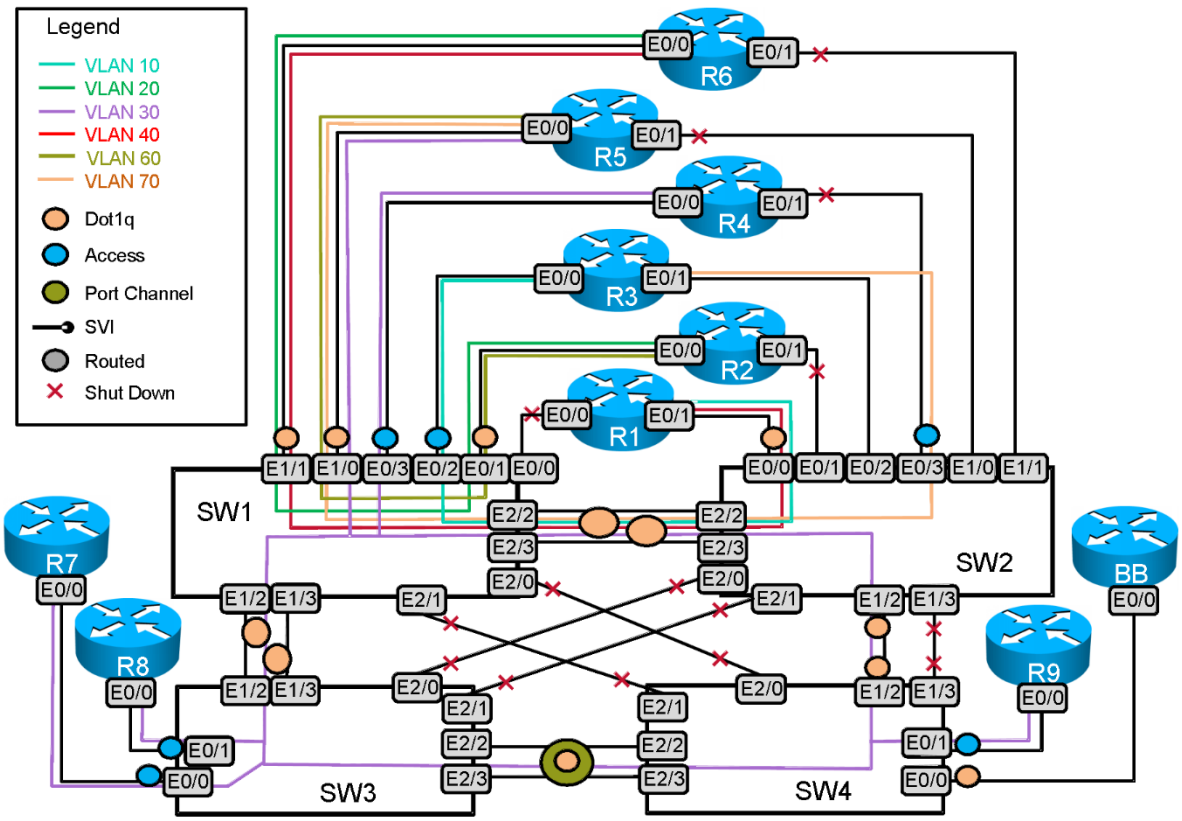

1.1. Symptom: Interface Port-Channel 1 on SW3 and SW4 is down.

Analysis and testing:

```
SW3#show ip int brief | inc Po
Port-channel1          unassigned      YES unset  down                  down
SW3#

SW3#show etherchannel port
                Channel-group listing:
                ----------------------

Group: 1
----------
                Ports in the group:
                -------------------
```


https://t.me/learningnets

```
Port: Et2/2
------------

Port state    = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1           Mode = Passive      Gcchange = -
Port-channel  = null        GC   =   -          Pseudo port-channel = Po1
Port index    = 0           Load = 0x00         Protocol =   LACP

Flags:  S - Device is sending Slow LACPDUs   F - Device is sending fast LACPDUs.
        A - Device is in active mode.        P - Device is in passive mode.

Local information:
                            LACP port     Admin     Oper    Port        Port
Port      Flags   State     Priority      Key       Key     Number      State
Et2/2     SP      indep     32768         0x1       0x1     0x203       0x7C

Age of the port in the current state: 0d:01h:34m:24s

Port: Et2/3
------------

Port state    = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1           Mode = Passive      Gcchange = -
Port-channel  = null        GC   =   -          Pseudo port-channel = Po1
Port index    = 0           Load = 0x00         Protocol =   LACP

Flags:  S - Device is sending Slow LACPDUs   F - Device is sending fast LACPDUs.
        A - Device is in active mode.        P - Device is in passive mode.

Local information:
                            LACP port     Admin     Oper    Port        Port
Port      Flags   State     Priority      Key       Key     Number      State
Et2/3     SP      indep     32768         0x1       0x1     0x204       0x7C

Age of the port in the current state: 0d:01h:34m:24s


SW3#
SW4#
SW4#show etherchannel port
                Channel-group listing:
                ----------------------

Group: 1
----------
                Ports in the group:
                -------------------
Port: Et2/2
------------

Port state    = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1           Mode = Passive      Gcchange = -
Port-channel  = null        GC   =   -          Pseudo port-channel = Po1
Port index    = 0           Load = 0x00         Protocol =   LACP

Flags:  S - Device is sending Slow LACPDUs   F - Device is sending fast LACPDUs.
        A - Device is in active mode.        P - Device is in passive mode.

Local information:
                            LACP port     Admin     Oper    Port        Port
Port      Flags   State     Priority      Key       Key     Number      State
Et2/2     SP      indep     32768         0x1       0x1     0x203       0x7C

Age of the port in the current state: 0d:01h:36m:55s

Port: Et2/3
------------
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
Port state    = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1           Mode = Passive       Gcchange = -
Port-channel  = null        GC   =   -           Pseudo port-channel = Po1
Port index    = 0           Load = 0x00          Protocol =   LACP

Flags:  S - Device is sending Slow LACPDUs   F - Device is sending fast LACPDUs.
        A - Device is in active mode.        P - Device is in passive mode.

Local information:
                            LACP port     Admin     Oper    Port         Port
Port      Flags   State     Priority      Key       Key     Number       State
Et2/3     SP      indep     32768         0x1       0x1     0x204        0x7C

Age of the port in the current state: 0d:01h:36m:56s


SW4#
```

Likely cause: Both sides of EtherChannel are in passive mode.

Remember that the lab requires that only SW3 should actively request negotiation for bundling interfaces Et2/2 and Et2/3.

Let's verify the configuration on SW3:

```
SW3#show run interface Et2/2
Building configuration...

Current configuration : 167 bytes
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
 switchport mode trunk
 duplex auto
 channel-group 1 mode passive
end

SW3#show run interface Et2/3
Building configuration...

Current configuration : 167 bytes
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
 switchport mode trunk
 duplex auto
 channel-group 1 mode passive
end

SW3#
```

Let's verify the configuration of SW4 now:

```
SW4#show run interface Et2/2
Building configuration...

Current configuration : 167 bytes
!
interface Ethernet2/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
```


https://t.me/learningnets

```
 switchport mode trunk
 duplex auto
 channel-group 1 mode passive
end

SW4#show run interface Et2/3
Building configuration...

Current configuration : 167 bytes
!
interface Ethernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 30
 switchport mode trunk
 duplex auto
 channel-group 1 mode passive
end

SW4#
```

Resolution: Configure SW3 in active mode.

If both sides are passive, the EtherChannel port will never come up.

Let's configure SW3 in active mode:

```
SW3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW3(config)#interface range Ethernet2/2 - 3
SW3(config-if-range)#channel-group 1 mode active
SW3(config-if-range)#end
SW3#
```

We immediately see logging showing that the Port-Channel1 comes up:

```
SW3#
*May 24 21:26:22.419: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
SW3#
*May  24  21:26:22.423:  %LINEPROTO-5-UPDOWN:  Line  protocol  on  Interface  Port-
channel1, changed state to up
SW3#
```

Let's verify the port-channel configuration on both switches again:

```
SW3#show etherchannel port-channel
                Channel-group listing:
                ----------------------

Group: 1
----------
                Port-channels in the group:
                ----------------------

Port-channel: Po1    (Primary Aggregator)

------------

Age of the Port-channel   = 0d:01h:45m:20s
Logical slot/port   = 16/1          Number of ports = 2
HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            =   LACP
Port security       = Disabled
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
Fast-switchover     = disabled
Load share deferral = disabled

Ports in the Port-channel:

Index   Load      Port          EC state       No of bits
------+------+------------+------------------+-----------
 0      00          Et2/2             Active   0
 0      00          Et2/3             Active   0

Time since last port bundled:    0d:00h:01m:35s    Et2/2

SW3#
SW4#show etherchannel port-channel
                Channel-group listing:
                ----------------------

Group: 1
----------
                Port-channels in the group:
                ---------------------

Port-channel: Po1    (Primary Aggregator)

------------

Age of the Port-channel   = 0d:01h:45m:56s
Logical slot/port   = 16/1          Number of ports = 2
HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            =   LACP
Port security       = Disabled
Fast-switchover     = disabled
Load share deferral = disabled

Ports in the Port-channel:

Index   Load      Port          EC state       No of bits
------+------+------------+------------------+-----------
 0      00          Et2/2            Passive   0
 0      00          Et2/3            Passive   0

Time since last port bundled:    0d:00h:02m:09s    Et2/3

SW4#
```

The interface Port-Channel 1 is now up:

```
SW3#show ip interface brief  | include Port
Port-channel1          unassigned      YES unset  up                    up
SW3#
```

1.2. Symptom: R9 cannot ping any device in subnet 135.15.20.0/24.

Analysis and testing:

R9 cannot ping R7 and R8.

Here is an example of the pings from R9 to R7 and R8 on the locally connected subnet 135.15.20.0/24:

```
R9#ping 135.15.20.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.20.10, timeout is 2 seconds:
```


https://t.me/learningnets

```
.....
Success rate is 0 percent (0/5)
R9#ping 135.15.20.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.20.20, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R9#
```

Likely cause: VLAN 30 is not fully operational.

Since R7 and R8 are connected to SW3, verify the VLANs that are known by SW3:

```
SW3#show vlan brie

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/2, Et0/3, Et1/0, Et1/1
                                                Et2/0, Et2/1
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
SW3#
```

We can see that SW3 does not have any VLANs that are used in this lab.

Remember that the lab requires that all VLANs be configured on SW4 only.

Therefore SW3 should learn all VLANs from SW4 via VTP.

Let's verify VTP configuration on SW3:

```
SW3#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 7
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 5
VTP Operating Mode              : Server
VTP Domain Name                 : CISCO360
VTP Pruning Mode                : Disabled (Operationally Disabled)
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x01 0xF4 0x10 0x12 0x87 0xEB 0xCB 0xF6
Configuration last modified by 0.0.0.0 at 5-24-13 16:24:25
Local updater ID is 0.0.0.0 (no valid interface found)
VTP version running             : 1
SW3#
```

The VTP domain name is correct.

Let's verify the password:

```
SW3#show vtp password
The VTP password is not configured.
SW3#
```

We can see that the password is missing on SW3.

Resolution: Configure the VTP password CISCO on SW3.

Configure the VTP password on SW3:

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
SW3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW3(config)#vtp password CISCO
Setting device VTP password to CISCO
SW3(config)#end
SW3#
```

Remember that VTP servers send VTP advertisements every 5 minutes. You may want to create a useless VLAN on SW4 and remove it to force SW4 to immediately send a VTP advertisement, so that you will not need to wait for up to 5 minutes.

Let's then check the VTP status on SW3:

```
SW3#show vtp status
VTP Version                     : 3 (capable)
Configuration Revision          : 12
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 11
VTP Operating Mode              : Server
VTP Domain Name                 : CISCO360
VTP Pruning Mode                : Enabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xB0 0x1F 0xBC 0xAA 0x47 0x1F 0xB9 0xAB
Configuration last modified by 0.0.0.0 at 3-3-93 13:19:18
Local updater ID is 135.15.130.1 on interface Lo130 (first layer3 interface found)
```

SW3 shows 11 VLANs now.

Check the VLANs that have been learned by SW3:

```
W3#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Et0/2, Et0/3, Et1/0, Et1/1
                                                Et2/0, Et2/1
10   VLAN0010                         active
20   VLAN0020                         active
30   VLAN0030                         active    Et0/0, Et0/1
40   VLAN0040                         active
60   VLAN0060                         active
70   VLAN0070                         active
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
SW3#
```

R9 can now ping its neighbors R7 and R8 in VLAN 30:

```
R9#ping 135.15.20.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.20.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R9#ping 135.15.20.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.20.20, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
R9#
```


https://t.me/learningnets

1.3. Symptom: Verifying MAC address table aging time.

Analysis and testing:

There is no symptom here that would get our attention except for the description in the troubleshooting ticket.

This lab requires that, for VLAN 30 only, inactive entries in the MAC address table of SW4 should be discarded twice as fast as the default behavior.

By default, the aging time is 300 seconds. You therefore want 150 seconds for VLAN 30 only; VLAN 1 must use the default aging time.

Let's verify the aging time that is configured on SW4.

```
SW4#show mac-address-table aging-time
Vlan    Aging Time
----    ----------
  30     150
   1     150
SW4#
```

Likely cause: The aging time for VLAN 1 is incorrect.

The aging time for VLAN 1 should remain at 300 seconds.

Let's verify the MAC aging time configuration commands on SW4:

```
SW4#show running-config | include aging
mac address-table aging-time 150 vlan 30
mac address-table aging-time 150 vlan 1
SW4#
```

Resolution: Modify the aging time for all VLANs other than VLAN 30.

```
SW4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#mac address-table aging-time 300
SW4(config)#mac address-table aging-time 150 vlan 30
SW4(config)#end
SW4#
```

Let's verify the aging times now:

```
SW4#show mac-address-table aging-time
Vlan    Aging Time
----    ----------
  30     150
   1     300
SW4#
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter into the engine more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

2. IPv4 OSPF Troubleshooting Section

2.1. Symptom: R1 and R4 are not neighbors.

Analysis and testing:

The following output indicates that a neighbor relationship is missing between R1 and R4:

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
135.15.106.1      1   FULL/DR         00:01:53    135.15.16.6     Ethernet0/1.40
135.15.102.1      1   FULL/DR         00:00:33    135.15.12.2     Serial1/0
135.15.103.1      1   FULL/DR         00:00:34    135.15.13.3     Ethernet0/1.10
```

We can see that R1 has built neighbor relationships with R2, R3, and R6, but not with R4, as required by the "Lab IPv4 IGP" diagram.

Likely cause: OSPF is not activated on S1/1.

Let's verify whether OSPF has been activated on S1/1:

```
R1#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et0/1.40     1     0               135.15.16.1/24     100   BDR   1/1
Se1/0        1     12              135.15.12.1/24     647   BDR   1/1
Et0/1.10     1     134             135.15.13.1/24     100   BDR   1/1
R1#
```

Note that the OSPF is not configured on the S1/1 interface on R1.

Verify the configuration of OSPF on R1:

```
R1#show running-config | section ospf
 ip ospf network non-broadcast
 ip ospf network broadcast
 ipv6 ospf 1 area 12
 ip ospf network point-to-multipoint
 ip ospf cost 65535
router ospf 1
 auto-cost reference-bandwidth 1000
 redistribute connected subnets route-map C2O
 network 135.15.12.1 0.0.0.0 area 12
 network 135.15.13.1 0.0.0.0 area 134
 network 135.15.16.1 0.0.0.0 area 0
 network 135.15.41.1 0.0.0.0 area 14
 neighbor 135.15.16.6
ipv6 router ospf 1
R1#
```

Resolution: An incorrect network command exists under router configuration on R1.

Remove the wrong network configuration and configure the network statement for the S1/1 interface:

```
R1(config)#router ospf 1
R1(config-router)#no network 135.15.41.1 0.0.0.0 area 14
R1(config-router)#network 135.15.14.1 0.0.0.0 area 134
```

We can immediately see the neighbor relationship being built between R1 and R4:

https://t.me/learningnets

```
*May 24 22:25:47.600: %OSPF-5-ADJCHG: Process 1, Nbr 135.15.104.1 on Serial1/1 from
LOADING to FULL, Loading Done
```

Let's confirm this with the **show ip ospf neighbor** command:

```
R1#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
135.15.106.1      1   FULL/DR         00:01:49    135.15.16.6     Ethernet0/1.40
135.15.102.1      1   FULL/DR         00:00:33    135.15.12.2     Serial1/0
135.15.104.1      0   FULL/  -        00:01:56    135.15.14.4     Serial1/1
135.15.103.1      1   FULL/DR         00:00:31    135.15.13.3     Ethernet0/1.10
R1#
```

2.2. Symptom: Incorrect OSPF cost on Et0/0.20 of R2.

Analysis and testing:

Remember that the lab requires that OSPF costs should be adapted so that a cost of 1 would be attributed to a 1000-Mb/s link.

First, check the bandwidth on the Ethernet 0/0.20 interface of R2:

```
R2#show interfaces Ethernet 0/0.20
Ethernet0/0.20 is up, line protocol is up
  Hardware is AmdP2, address is aabb.cc00.0200 (bia aabb.cc00.0200)
  Internet address is 135.15.26.2/24
  MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation 802.1Q Virtual LAN, Vlan ID  20.
  ARP type: ARPA, ARP Timeout 04:00:00
  Keepalive set (10 sec)
  Last clearing of "show interface" counters never
R2#
```

Then, verify the OSPF cost on the Et0/0.20 interface on R2:

```
R2#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et0/0.20     1     0               135.15.26.2/24     10    P2MP  1/1
Se1/0        1     12              135.15.12.2/24     64    DR    1/1
R2#
```

A 10-Mb/s interface should therefore get a cost of 100:

Likely cause: The reference bandwidth is incorrect.

We see that Et0/0.20 gets an OSPF cost of 10 instead of 100.

Remember that the **auto-cost reference-bandwidth** command should be configured consistently on all devices.

Resolution: Modify the reference bandwidth.

Without any configuration, the reference bandwidth used by OSPF by default is 100. This means that a bandwidth of 10 Mb/s gets a cost of 10. This lab requires that a cost of 1 be attributed to a bandwidth of 1000 Mb/s.

Verify the OSPF configuration on R2:

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
R2#show running-config | section router ospf
router ospf 1
 redistribute rip subnets route-map R2O
 network 135.15.12.2 0.0.0.0 area 12
 network 135.15.26.2 0.0.0.0 area 0
 distance 121 0.0.0.0 255.255.255.255 RIP
ipv6 router ospf 1
R2#
```

Note that the **auto-cost reference-bandwidth** command is not configured on R2.

Fix the OSPF configuration on R2:

```
R2(config)#router ospf 1
R2(config-router)#auto-cost reference-bandwidth 1000
% OSPF: Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent across all routers.
```

Verify the OSPF interfaces on R2 again:

```
R2#show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et0/0.20     1     0               135.15.26.2/24     100    P2MP  0/0
Se0/0/0.12   1     12              135.15.12.2/24     647   DR    1/1
```

Note that the E0/0.20 interface is assigned with a cost of 100 now.

2.3. Symptom: R2 and R6 are not neighbors.

Analysis and testing:

The following output indicates that a neighbor relationship is missing between R2 and R6:

```
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
135.15.101.1      1   FULL/BDR        00:00:36    135.15.12.1     Serial1/0
R2#
```

And on R6:

```
R6#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
135.15.101.1      1   FULL/DR         00:01:50    135.15.16.1     Ethernet0/0.40
R6#
```

Remember that this lab indicates which subnets should elect a DR and BDR. We can read that subnet 135.15.26.0/24 does not belong to this list.

Therefore, subnet 135.15.26.0/24 on VLAN 20 may be neither an OSPF network type broadcast nor nonbroadcast. It may only be point-to-multipoint or point-to-multipoint nonbroadcast.

Furthermore, the lab indicates which subnets should automatically discover OSPF neighbors.

https://t.me/learningnets

We can read that subnet 135.15.26.0/24 does not belong to this list.

Therefore subnet 135.15.26.0/24 may not be an OSPF network type point-to-multipoint.

It may only be point-to-multipoint nonbroadcast.

Therefore, a **neighbor** command is required, either on R2 or R6.

Likely cause: A** neighbor **command is missing.

Do we have a neighbor statement on R2?

```
R2#show run | section router ospf
router ospf 1
 auto-cost reference-bandwidth 1000
 redistribute rip subnets route-map R2O
 network 135.15.12.2 0.0.0.0 area 12
 network 135.15.26.2 0.0.0.0 area 0
 distance 121 0.0.0.0 255.255.255.255 RIP
ipv6 router ospf 1
R2#
```

There is no OSPF neighbor configuration on R2.

Do we have a neighbor statement on R6?

```
R6#show run | section router ospf
router ospf 1
 auto-cost reference-bandwidth 1000
 redistribute connected subnets route-map C2O
 network 135.15.16.6 0.0.0.0 area 0
 network 135.15.26.6 0.0.0.0 area 0
R6#
```

There is no OSPF neighbor configuration on R6.

Resolution: Configure the OSPF neighbor either on R2 or on R6, or on both.

R2 is configured with the OSPF neighbor statement in this answer key:

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#router ospf 1
R2(config-router)#neighbor 135.15.26.6
R2(config-router)#end
R2#
```

Verify the neighbors on R2:

```
R2#show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
135.15.106.1      0   FULL/  -        00:01:54    135.15.26.6     Ethernet0/0.20
135.15.101.1      1   FULL/BDR        00:00:38    135.15.12.1     Serial1/0
R2#
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter into the engine more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

3. EIGRP Troubleshooting Section

3.1. Symptom: EIGRP logging is on all EIGRP devices every 10 seconds.

Analysis and testing:

We get the following logging messages on R8:

```
R8#
*May  24  23:06:36.719:  %DUAL-5-NBRCHANGE:  EIGRP-IPv4  100:  Neighbor  135.15.20.4
(Ethernet0/0) is down: K-value mismatch
R8#
*May  24  23:06:38.042:  %DUAL-5-NBRCHANGE:  EIGRP-IPv4  100:  Neighbor  135.15.20.10
(Ethernet0/0) is down: K-value mismatch
R8#
*May  24  23:06:39.321:  %DUAL-5-NBRCHANGE:  EIGRP-IPv4  100:  Neighbor  135.15.20.30
(Ethernet0/0) is down: K-value mismatch
R8#
```

Verify the EIGRP neighbors. Here is an example from R8:

```
R8#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
R8#
```

Note that R8 does not show any EIGRP neighbors.

Likely cause: A K-value mismatch exists.

Remember that two EIGRP routers will not become neighbors if they do not have the same K values.

By default, K1 = 1, K2 = 0, K3 = 1, K4 = 0, and K5 = 0. This lab requires that K3 = 3.

Therefore it has to be configured on all EIGRP AS 100 devices.

Resolution: Configure K3 on R8.

You verified all EIGRP AS 100 devices and found that R8 is missing the K-value configuration.

Fix the EIGRP configuration on R8:

```
R8#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R8(config)#router eigrp 100
R8(config-router)# metric weights 0 1 0 3 0 0


R8#show ip protocols | inc weigh
    Metric weight K1=1, K2=0, K3=3, K4=0, K5=0
R8#
```

Verify that R8 shows all neighbors now:

```
R8#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
2   135.15.20.10            Et0/0                    13 00:01:54   18   108  0  8
1   135.15.20.30            Et0/0                    12 00:01:54   11   100  0  9
0   135.15.20.4             Et0/0                    11 00:01:54    7   100  0  11
R8#
```


https://t.me/learningnets

Note that R8 now forms EIGRP AS 100 neighbor relationships with R4, R7, and R9. It does not form the EIGRP AS 100 neighbor relationship with R5.

3.2. Symptom: No EIGRP neighbors on R5.

Analysis and testing:

Verify the EIGRP neighbor relationships on R5:

```
R5#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
R5#
```

Note that R5 does not form any EIGRP neighbor relationships.

Verify the EIGRP configuration on the Et0/0.30 interface on R5.

```
R5#show ip eigrp interfaces detail
EIGRP-IPv4 Interfaces for AS(100)
                              Xmit  Queue     PeerQ             Mean     Pacing  Time
Multicast    Pending
Interface               Peers  Un/Reliable  Un/Reliable  SRTT    Un/Reliable    Flow
Timer   Routes
Et0/0.30                 0        0/0       0/0           0       0/0             0
0
  Hello-interval is 10, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 0/0
  Hello's sent/expedited: 1418/1
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  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 "KEY"
R5#
```

Verify the key chain configuration:

```
R5#show key chain
Key-chain KEY:
    key 1 -- text "CICSO"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
```

Likely cause: The password is misspelled in the key chain.

Note that there is a wrong key, CICSO. The key should be CISCO.

Resolution: Fix the EIGRP password.

```
R5(config)#no key chain KEY
R5(config)#key chain KEY
R5(config-keychain)# key 1
R5(config-keychain-key)#   key-string CISCO
```

Verify the EIGRP neighbor relationships on R5 again:

```
R5#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
3   135.15.20.30            Et0/0.30                 13 00:00:38   11   100  0  12
2   135.15.20.20            Et0/0.30                  6 00:00:38  820  4920  0  10
1   135.15.20.4             Et0/0.30                  7 00:00:38  525  3150  0  16
0   135.15.20.10            Et0/0.30                  6 00:00:38    2   100  0  11
R5#
```

Note that R5 now forms all necessary EIGRP AS 100 neighbor relationships.

3.3. Symptom: No device can reach subnet 135.15.110.0/24.

Analysis and testing:

Here is an example of the connectivity test to 135.15.110.0/24 from R8:

```
R8#ping 135.15.110.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.110.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R8#
```

Note that the ping is not successful.

This lab requires that the subnet 135.15.110.0/24 must be redistributed into EIGRP on R7:

```
R7#show running-config | section router eigrp
router eigrp 100
 metric weights 0 1 0 3 0 0
 default-metric 1 1 1 1 1
 network 135.15.20.0 0.0.0.255
 redistribute connected route-map C2E
 auto-summary
R7#
```

Verify the route-map C2E configuration:

```
R7#show route-map
route-map C2E, permit, sequence 10
  Match clauses:
    interface Ethernet0/0
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
R7#
```

Likely cause: The route map C2E is incorrect.

The route map C2E is configured to redistribute the interface Ethernet0/0. This is useless, since EIGRP has been already activated on the Ethernet0/0 interface with the **network 135.15.20.0 0.0.0.255** command.

This route map C2E should redistribute the Loopback 101 interface.

Resolution: Match the interface Loopback 110 instead of Ethernet0/0 in the route map C2E.

```
R7#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R7(config)#no route-map C2E permit 10
R7(config)#route-map C2E permit 10
R7(config-route-map)# match interface Loopback110
```

Let's verify that this subnet is now reachable:

https://t.me/learningnets

```
R8#ping 135.15.110.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.110.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R8#
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter into the engine more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

4. IPv4 RIP Troubleshooting Section

4.1. Symptom: R2 does not have any RIP routes.

Analysis and testing:

You verified the RIP routing table on R2 and found that there are no RIP prefixes in the routing table:

```
R2#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
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

R2#
```

Remember that this lab requires R2 and R5 to deny all RIP updates that are received directly from each other. In other words, R2 and R5 should only exchange the RIP updates via R3. Does R3 have RIP routes?

Verify the RIP routing table on R3:

```
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
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      135.15.0.0/16 is variably subnetted, 27 subnets, 3 masks
R        135.15.4.0/23 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R        135.15.20.0/24 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R        135.15.25.0/24 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R        135.15.105.0/24 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R        135.15.110.0/24 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R        135.15.120.0/24 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R        135.15.130.0/24 [120/1] via 135.15.35.5, 00:00:05, Ethernet0/1
R3#
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Note that R3 learns RIP prefixes from R5.

Is there a problem with the RIP updates exchanged between R2 and R3?

Likely cause: No RIP updates are exchanged between R2 and R3.

Verify if R2 and R3 exchange RIP updates:

```
R2#show ip protocols | section "rip"
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is  gateway FilterRip
  Sending updates every 30 seconds, next due in 18 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: connected, ospf 1 (internal, external 1 & 2, nssa-external 1 & 2)

  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface              Send  Recv  Triggered RIP  Key-chain
    Ethernet0/0.60         2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
     135.15.0.0
    135.15.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/0.20
    Ethernet0/1
    Ethernet0/2
  Passive Interface(s):
    Ethernet0/3
    Serial1/0
    Serial1/1
    Serial1/2
    Serial1/3
    Loopback102
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 120)
R2#

R3#show ip protocols | section "rip"
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 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: ospf 1 (internal, external 1 & 2, nssa-external 1 & 2)

  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface              Send  Recv  Triggered RIP  Key-chain
    Ethernet0/1            2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    135.15.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
```


https://t.me/learningnets

```
    Serial1/2
    Serial1/3
  Passive Interface(s):
    Loopback103
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    135.15.35.5          120      00:00:04
  Distance: (default is 120)
R3#
```

Note that R2 and R3 do not send any RIP updates because their respective S1/1 interfaces are passive.

Resolution: Fix the passive interface configuration on R2 and R3.

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#router rip
R2(config-router)#no passive-interface Serial1/1
R2(config-router)#end
R2#

R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#router rip
R3(config-router)#no passive-interface Serial1/1
R3(config-router)#end
R3#
```

Verify the RIP routing table on R2:

```
R2#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
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      135.15.0.0/16 is variably subnetted, 27 subnets, 3 masks
R        135.15.4.0/23 [120/2] via 135.15.23.3, 00:00:16, Serial1/1
R        135.15.20.0/24 [120/2] via 135.15.23.3, 00:00:16, Serial1/1
R        135.15.35.0/24 [120/1] via 135.15.23.3, 00:00:16, Serial1/1
R        135.15.105.0/24 [120/2] via 135.15.23.3, 00:00:16, Serial1/1
R        135.15.110.0/24 [120/2] via 135.15.23.3, 00:00:16, Serial1/1
R        135.15.120.0/24 [120/2] via 135.15.23.3, 00:00:16, Serial1/1
R        135.15.130.0/24 [120/2] via 135.15.23.3, 00:00:16, Serial1/1
R2#
```

Note that R2 now receives the RIP routes from R3.

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter into the engine more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

5. IPv4 Redistribution Troubleshooting Section

5.1. Symptom: EIGRP devices cannot ping 135.15.104.1 on R4.

Analysis and testing:

Here is an example of the ping to 135.15.104.1 from R9:

```
R9#ping 135.15.104.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.104.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R9#
```

According to the lab requirements, the Loopback 104 interface on R4 should be redistributed into OSPF on R4.

Verify how R3 learns the Loopback 104 interface from R4:

```
R3#show ip route ospf | include E2
       E1 - OSPF external type 1, E2 - OSPF external type 2
O E2     135.15.101.0/24 [110/101] via 135.15.13.1, 04:11:48, Ethernet0/0
O E2     135.15.104.0/24 [110/104] via 135.15.34.4, 04:13:01, Serial1/0
O E2     135.15.106.0/24 [110/106] via 135.15.13.1, 01:02:01, Ethernet0/0
R3#
```

R3 correctly learns the Loopback 104 network as an OSPF external prefix from R4.

All OSPF devices can reach the 135.15.104.0/24 subnet. Only EIGRP devices R5, R7, R8, and R9 are not able to reach this subnet. The EIGRP routers should learn the 135.15.104.0/24 subnet from R5, which is the redistribution point between EIGRP AS 100 and RIP version 2. Does R5 show the 135.15.104.0/24 subnet in the RIP routing table?

```
R5#show ip route rip | inc 104
R5#
```

Note that R5 does not have the 135.15.104.0/24 subnet in the RIP routing table.

Likely cause: Subnet 135.15.104.0/24 is not redistributed from OSPF into RIP.

Verify the redistribution process from OSPF into RIP on R2:

```
R2#show run | section router rip
router rip
 version 2
 redistribute connected metric 1 route-map C2R
 redistribute ospf 1 metric 1 route-map O2R
 passive-interface default
 no passive-interface Ethernet0/0.60
 no passive-interface Serial1/1
 network 135.15.0.0
 distribute-list gateway FilterRip in
R2#

R2#show route-map O2R
route-map O2R, permit, sequence 10
  Match clauses:
    ip address (access-lists): OSPF
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
```


https://t.me/learningnets

```
R2#

R2#show ip access-lists OSPF
Standard IP access list OSPF
    10 permit 135.15.14.0, wildcard bits 0.0.0.255 (8 matches)
    20 permit 135.15.34.0, wildcard bits 0.0.0.255 (4 matches)
    30 permit 135.15.12.0, wildcard bits 0.0.0.255
    40 permit 135.15.13.0, wildcard bits 0.0.0.255 (4 matches)
    50 permit 135.15.16.0, wildcard bits 0.0.0.255 (6 matches)
    60 permit 135.15.26.0, wildcard bits 0.0.0.255 (5 matches)
    70 permit 135.15.106.0, wildcard bits 0.0.0.255 (4 matches)
    80 permit 135.15.101.0, wildcard bits 0.0.0.255 (4 matches)
    90 permit 135.15.103.0, wildcard bits 0.0.0.255 (3 matches)
R2#
```

Note that the subnet 135.15.104.0/24 is not permitted by the access list OSPF on R2.

R3 is configured with the same access list.

Resolution: Permit the subnet 135.15.134.0/24 in the access list OSPF.

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#ip access-list standard OSPF
R2(config-std-nacl)#permit 135.15.104.0 0.0.0.255
R2(config-std-nacl)#end
R2#

R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#ip access-list standard OSPF
R3(config-std-nacl)#permit 135.15.104.0 0.0.0.255
R3(config-std-nacl)#end
R3#
```

Verify reachability on R9 again:

```
R9#ping 135.15.104.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 135.15.104.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms
R9#
```

Now that we seem to have addressed all of the IPv4 unicast issues, we will test reachability using this simple TCL script: We enter the command **tclsh** and paste in this script. When it is complete, we will have a record of successful and unsuccessful pings. We enter the command **tclquit** to exit the command interpreter.

```
tclsh
foreach address {
135.15.13.1
135.15.12.1
135.15.14.1
135.15.16.1
135.15.101.1
135.15.12.2
135.15.26.2
135.15.25.2
135.15.23.2
135.15.102.1
135.15.13.3
135.15.23.3
135.15.35.3
135.15.34.3
135.15.103.1
135.15.14.4
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
135.15.20.4
135.15.34.4
135.15.104.1
135.15.25.5
135.15.4.1
135.15.5.1
135.15.20.5
135.15.35.5
135.15.105.1
135.15.26.6
135.15.16.6
135.15.106.1
135.15.20.30
135.15.130.1
135.15.20.10
135.15.110.1
135.15.20.20
135.15.120.1
} {ping $address}
```

6. Security Troubleshooting Section

6.1. Symptom: Verifying the reflexive access list.

Analysis and testing:

Let's verify the outbound access list, which, according to the lab requirements, should allow hosts to reach the HTTP, SSL, and Telnet servers. Moreover, the access list should also allow hosts to ping any device and to reply to pings that are received from any device:

```
R4#show ip access-lists
Extended IP access list IN_BOUND
    10 evaluate RETURN
    20 permit ospf any any (19 matches)
    30 deny ip any any log
    40 permit icmp any any
Extended IP access list OUT_BOUND
    10 permit tcp any eq telnet any reflect RETURN
    20 permit tcp any eq 443 any reflect RETURN
    30 permit tcp any eq www any reflect RETURN
    40 permit icmp any any echo reflect RETURN
    50 permit icmp any any echo-reply reflect RETURN
Reflexive IP access list RETURN
```

Likely cause: The access list is incorrect.

We can see that instead of allowing any host to reach these resources, the OUT_BOUND access list is allowing specific resources to reach any host.

Resolution: Modify the OUT_BOUND access list.

```
R4(config)#no ip access-list ex OUT_BOUND
R4(config)#ip access-list ex OUT_BOUND
R4(config-ext-nacl)#permit tcp any any eq 23 reflect RETURN
R4(config-ext-nacl)#permit tcp any any eq 443 reflect RETURN
R4(config-ext-nacl)#permit tcp any any eq 80 reflect RETURN
R4(config-ext-nacl)#permit icmp any any echo reflect RETURN
R4(config-ext-nacl)#permit icmp any any echo-reply reflect RETURN
```


https://t.me/learningnets

6.2. Symptom: Verifying the application of the access list.

Analysis and testing:

Let's verify where the access lists have been applied:

```
R4#show ip interface e0/0 | inc access list
  Outgoing access list is not set
  Inbound  access list is not set
R4#

R4#show ip interface s1/0 | inc access list
  Outgoing access list is OUT_BOUND
  Inbound  access list is IN_BOUND
R4#

R4#show ip interface s1/1 | inc access list
  Outgoing access list is not set
  Inbound  access list is not set
R4#
```

Likely cause: The** access-group **command is missing.

Remember that in the OSPF section, a cost of 65535 had been applied to the interface S1/1.

Therefore, all OSPF routes on R4 currently use the next hop, 135.15.34.3.

If the serial link between R4 and R3 goes down, R4 will begin using the serial link toward R1.

Therefore, the reflexive access list should be applied to both serial interfaces S1/0 and S1/1.

Resolution: Apply access lists to the Serial1/1 interface.

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#interface Serial1/1
R4(config-if)# ip access-group IN_BOUND in
R4(config-if)# ip access-group OUT_BOUND out
R4(config-if)#end
R4#
```

Let's verify our configuration:

```
R4#show ip interface s1/0 | inc access list
  Outgoing access list is OUT_BOUND
  Inbound  access list is IN_BOUND
R4#
```

Note that the reflexive access lists will not evaluate traffic that is sourced from the same router where the reflexive access list is configured. Traffic should be sourced from a different router in your testing of the reflexive lists.

7. IPv6 Troubleshooting Section

7.1. Symptom: R5 reaches subnet FEC2::101:0/125 via OSPF.

Analysis and testing:

This lab requires R5 to reach the FEC2::101:0/125 subnet via IPv6 EIGRP.

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

Verify the IPv6 routing table on R5:

```
R5#show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
OI  FEC2::12:0/125 [110/74]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
OE2 FEC2::13:0/125 [110/20]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
OI  FEC2::23:0/125 [110/74]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
C   FEC2::25:0/125 [0/0]
     via Ethernet0/0.60, directly connected
L   FEC2::25:5/128 [0/0]
     via Ethernet0/0.60, receive
C   FEC2::35:0/125 [0/0]
     via Ethernet0/0.70, directly connected
L   FEC2::35:5/128 [0/0]
     via Ethernet0/0.70, receive
OE2 FEC2::101:0/125 [110/20]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
L   FF00::/8 [0/0]
     via Null0, receive
R5#
```

Note that R5 does not show any EIGRP prefixes, even internal prefixes. Does R5 form the IPv6 EIGRP neighbor relationship with R3?

```
R5#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(100)
R5#
```

R5 does not form the IPv6 EIGRP neighbor relationship with R3.

Is IPv6 EIGRP configured on the E0/0.70 interface on R5?

```
R5#show ipv6 eigrp int
EIGRP-IPv6 Interfaces for AS(100)
                              Xmit  Queue     PeerQ             Mean     Pacing  Time
Multicast    Pending
Interface               Peers  Un/Reliable  Un/Reliable  SRTT    Un/Reliable    Flow
Timer   Routes
Et0/0.70                 0        0/0       0/0           0       0/2          2628
0
R5#
```

IPv6 EIGRP is configured correctly on the E0/0.70 interface on R5. Are the IPv6 EIGRP packets blocked by the IPv6 filter on R5?

Likely cause: R5 does not receive IPv6 EIGRP traffic.

Verify the IPv6 access lists on R5:

```
R5#show ipv6 interface Ethernet 0/0.70 | inc access list
```


https://t.me/learningnets

```
  Inbound access list No99
R5#

R5#show ipv6 access-list
IPv6 access list No99
    deny ipv6 any FEC2::35:98/125 sequence 10
R5#
```

The incorrectly configured IPv6 access list is preventing IPv6 EIGRP updates from reaching R5.

Resolution: Fix the IPv6 access list.

We first correct the access list:

```
R5(config)#no ipv6 access-list No99
R5(config)#ipv6 access-list No99
R5(config-ipv6-acl)# deny ipv6 any FEC2::35:98/125
R5(config-ipv6-acl)# permit ipv6 any any
```

We then apply it outbound on the Et0/0.70 interface, not inbound:

```
R5(config-subif)#int Et0/0.70
R5(config-subif)#ipv6 traffic-filter No99 out
R5(config-subif)#no ipv6 traffic-filter No99 in
```

Verify the IPv6 EIGRP neighbors again:

```
R5#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q  Seq
                                                   (sec)         (ms)       Cnt Num
0   Link-local address:     Et0/0.70                 12 00:02:08    8   100  0  28
    FE80::A8BB:CCFF:FE00:310
R5#
```

Verify the IPv6 routing table on R5 again:

```
R5#show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
OI  FEC2::12:0/125 [110/74]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
D   FEC2::13:0/125 [90/307200]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.70
OI  FEC2::23:0/125 [110/74]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
C   FEC2::25:0/125 [0/0]
     via Ethernet0/0.60, directly connected
L   FEC2::25:5/128 [0/0]
     via Ethernet0/0.60, receive
C   FEC2::35:0/125 [0/0]
     via Ethernet0/0.70, directly connected
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
L   FEC2::35:5/128 [0/0]
     via Ethernet0/0.70, receive
OE2 FEC2::101:0/125 [110/20]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
L   FF00::/8 [0/0]
     via Null0, receive
R5#
```

Note that R5 learns the IPv6 EIGRP prefixes now. But R5 still prefers the FEC2::101:0/125 subnet via OSPFv3.

Configure the OSPFv3 administrative distance 171 to ensure that the IPv6 EIGRP external prefixes are preferred:

```
R5#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R5(config)#ipv6 router ospf 1
R5(config-rtr)#distance ospf external 171
R5(config-rtr)#end
R5#
```

Verify the IPv6 routing table on R5 again:

```
R5#show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
OI  FEC2::12:0/125 [110/74]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
D   FEC2::13:0/125 [90/307200]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.70
OI  FEC2::23:0/125 [110/74]
     via FE80::A8BB:CCFF:FE00:200, Ethernet0/0.60
C   FEC2::25:0/125 [0/0]
     via Ethernet0/0.60, directly connected
L   FEC2::25:5/128 [0/0]
     via Ethernet0/0.60, receive
C   FEC2::35:0/125 [0/0]
     via Ethernet0/0.70, directly connected
L   FEC2::35:5/128 [0/0]
     via Ethernet0/0.70, receive
EX  FEC2::101:0/125 [170/2560051456]
     via FE80::A8BB:CCFF:FE00:310, Ethernet0/0.70
L   FF00::/8 [0/0]
     via Null0, receive
R5#
```

Let's now test reachability using this simple Tcl script: We enter the command **tclsh** and paste in this script. When it is complete, we will have a record of successful and unsuccessful pings. We enter the command **tclquit** to exit the command interpreter.

```
tclsh
foreach address {
FEC2::13:1
```


https://t.me/learningnets

```
FEC2::12:1
FEC2::23:2
FEC2::12:2
FEC2::25:2
FEC2::13:3
FEC2::23:3
FEC2::35:3
FEC2::25:5
FEC2::35:5
} {ping $address}
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter into the engine more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

8. QoS Troubleshooting Section

8.1. Symptom: Verification of the policy map.

Analysis and testing:

There is no symptom that attracts our attention. A good starting point is to verify the configuration of the policy map with either one of these commands:

```
R1#show policy-map interface
 Serial1/0

  Service-policy output: QoS12

    queue stats for all priority classes:
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

    Class-map: HTTP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: protocol http
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0
      bandwidth 25 kbps

    Class-map: SNMP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: protocol snmp
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0
      bandwidth 20 kbps

    Class-map: NTP (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: protocol ntp
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
      bandwidth 5 kbps

    Class-map: VOICE (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: ip rtp 16384 16383
      Priority: 20 kbps, burst bytes 1500, b/w exceed drops: 0


    Class-map: class-default (match-any)
      104 packets, 8478 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any

      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 104/8602
R1#

R1#show policy-map
  Policy Map QoS12
    Class HTTP
      bandwidth 25 (kbps)
    Class SNMP
      bandwidth 20 (kbps)
    Class NTP
      bandwidth 5 (kbps)
    Class VOICE
      priority 20 (kbps)

R1#
```

Likely cause: The configuration of the policy map is incorrect.

We can read that instead of percentage bandwidth configuration, absolute bandwidth has been configured.

Resolution: Modify policy-map.

```
R1(config)#no policy-map QoS12
R1(config)#policy-map QoS12
R1(config-pmap)# class HTTP
R1(config-pmap-c)#  bandwidth percent 25
R1(config-pmap-c)# class SNMP
R1(config-pmap-c)#  bandwidth percent 20
R1(config-pmap-c)# class NTP
R1(config-pmap-c)#  bandwidth percent 5
R1(config-pmap-c)# class VOICE
R1(config-pmap-c)#  priority percent 20
```

Do not forget to apply the policy map to the S1/0 interface, since the QoS policy has been removed from the S1/0 interface as soon as we deleted it.

```
R1(config-pmap-c)#interface s1/0
R1(config-if)#service-policy output QoS12
```

Verify the QoS policy again:

```
R1#show policy-map
  Policy Map QoS12
    Class HTTP
      bandwidth 25 (%)
    Class SNMP
      bandwidth 20 (%)
    Class NTP
```


https://t.me/learningnets

```
      bandwidth 5 (%)
    Class VOICE
      priority 20 (%)

R1#
```

Note	The Mentor Guide engine in the web portal can help you use Cisco IOS Software commands to see a comprehensive view of the configuration for a specific section. With the Mentor Guide engine, you can enter more than 1000 Cisco IOS Software commands as well as a collection of proprietary commands such as **show all**.

9. IP Service Troubleshooting Section

9.1. Symptom: R2 remains unsynchronized.

Analysis and testing:

Verify the NTP configuration on R2:

```
R2#show ntp associations detail
135.15.26.6 configured, ipv4, insane, invalid, unsynced, stratum 16
ref ID .INIT., time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
our mode client, peer mode server, our poll intvl 1024, peer poll intvl 1024
root delay 0.00 msec, root disp 930.63, reach 0, sync dist 16882.99
delay 0.00 msec, offset 0.0000 msec, dispersion 15937.50, jitter 0.00 msec
precision 2**10, version 4
assoc id 6663, assoc name 135.15.26.6
assoc in packets 402, assoc out packets 403, assoc error packets 402
org time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
rec time D54B3383.E353FA40 (04:56:35.888 PST Sat May 25 2013)
xmt time D54B3383.E353FA40 (04:56:35.888 PST Sat May 25 2013)
filtdelay =     0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00
filtoffset =    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00
filterror =  16000.0 16000.0 16000.0 16000.0 16000.0 16000.0 16000.0 16000.0
minpoll = 6, maxpoll = 10

R2#
```

Note that R2 is not synchronized with R6.

Verify the NTP configuration on R6:

```
R6#show ntp associations detail
R6#
R6#show ntp status
Clock is unsynchronized, stratum 16, no reference clock
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
ntp uptime is 6461400 (1/100 of seconds), resolution is 4000
reference time is 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 969.22 msec, peer dispersion is 0.00 msec
loopfilter state is 'FSET' (Drift set from file), drift is 0.000000000 s/s
system poll interval is 8, never updated.
R6#
```

Note that R6 does not show any NTP associations. R6 should be configured as an NTP server and should be synchronized with its own clock.

Likely cause: R6 is not configured as an NTP server.

Verify the NTP configuration on R6:

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
R6#show running-config | inc ntp
ntp authentication-key 1 md5 0632260C69 7
R6#
```

Note that R6 is not configured for any NTP role. According to the lab requirements, R6 should be configured as the NTP master.

Resolution: Configure R6 as the NTP master.

```
R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)#ntp master
R6(config)#end
R6#
```

Verify the NTP configuration on R6 again:

```
R6#show ntp associations detail
127.127.1.1 configured, ipv4, our_master, sane, valid, stratum 7
ref ID .LOCL., time D54B3F1C.AC083300 (05:46:04.672 PST Sat May 25 2013)
our mode active, peer mode passive, our poll intvl 16, peer poll intvl 16
root delay 0.00 msec, root disp 0.00, reach 7, sync dist 1939.55
delay 0.00 msec, offset 0.0000 msec, dispersion 1938.47, jitter 0.97 msec
precision 2**10, version 4
assoc id 49839, assoc name 127.127.1.1
assoc in packets 3, assoc out packets 3, assoc error packets 0
org time D54B3F1C.AC083300 (05:46:04.672 PST Sat May 25 2013)
rec time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
xmt time D54B3F1C.AC083300 (05:46:04.672 PST Sat May 25 2013)
filtdelay =     0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00
filtoffset =    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00
filterror =     0.97    1.21    1.45 16000.0 16000.0 16000.0 16000.0 16000.0
minpoll = 4, maxpoll = 4

R6#
```

Verify the NTP configuration on R2. The NTP synchronization process can be slow, but you can speed it up by reapplying the **ntp server 135.15.26.6** command on R2 to immediately initiate the NTP synchronization from R2:

```
R2#show ntp status
Clock is synchronized, stratum 9, reference is 135.15.26.6
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
ntp uptime is 6543200 (1/100 of seconds), resolution is 4000
reference time is D54B4082.E24DD560 (05:52:02.884 PST Sat May 25 2013)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 6.70 msec, peer dispersion is 2.66 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s
system poll interval is 64, last update was 55 sec ago.
R2#
```

Note the R2 is now synchronized with R6. But the NTP synchronization is not authenticated as required in the lab.

9.2. Symptom: Verifying NTP authentication.

Analysis and testing:

Verify the NTP association on R2:

```
R2#show ntp associations detail
135.15.26.6 configured, ipv4, our_master, sane, valid, stratum 8
ref ID 127.127.1.1    , time D54B414C.AD0E57E0 (05:55:24.676 PST Sat May 25 2013)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 2.19, reach 77, sync dist 7.80
```


https://t.me/learningnets

```
delay 0.00 msec, offset 0.0000 msec, dispersion 4.23, jitter 0.97 msec
precision 2**10, version 4
assoc id 6664, assoc name 135.15.26.6
assoc in packets 11, assoc out packets 11, assoc error packets 0
org time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
rec time D54B414E.E353FA40 (05:55:26.888 PST Sat May 25 2013)
xmt time D54B414E.E353FA40 (05:55:26.888 PST Sat May 25 2013)
filtdelay =     1.00    0.00    1.00    0.00    0.00    0.00    1.00    0.00
filtoffset =    0.50    0.00    0.50    0.00    0.00    0.00    0.50    0.00
filterror =     1.95    2.97    3.99    5.01    6.01    6.85    6.88    6.91
minpoll = 6, maxpoll = 10

R2#
```

Likely cause: NTP authentication on R2 and R6 is misconfigured.

Verify the NTP configuration of R2 and R6:

```
R2#show running-config | inc ntp
ntp authentication-key 1 md5 081565632C 7
ntp trusted-key 1
ntp server 135.15.26.6
R2#

R6#show running-config | inc ntp
ntp authentication-key 1 md5 073B08616B 7
ntp master
R6#
```

Resolution: NTP authentication commands are missing on R2.

Fix the NTP authentication on R2:

```
R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#ntp authenticate
R2(config)#ntp server 135.15.26.6 key 1
R2(config)#end
R2#
```

Fix the NTP authentication on R6:

```
R6#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)# ntp authenticate
R6(config)#ntp trusted-key 1
R6(config)#end
R6#
```

In a few moments, verify the NTP associations on R2 again:

```
R2#show ntp associations detail
135.15.26.6 configured, ipv4, authenticated, our_master, sane, valid, stratum 8
ref ID 127.127.1.1     , time D54B472C.AC49BC38 (06:20:28.673 PST Sat May 25 2013)
our mode client, peer mode server, our poll intvl 64, peer poll intvl 64
root delay 0.00 msec, root disp 2.25, reach 17, sync dist 8.65
delay 0.00 msec, offset 0.0000 msec, dispersion 4.58, jitter 0.97 msec
precision 2**10, version 4
assoc id 6667, assoc name 135.15.26.6
assoc in packets 9, assoc out packets 9, assoc error packets 0
org time 00000000.00000000 (16:00:00.000 PST Wed Dec 31 1899)
```

 © 2013 Cisco Systems, Inc.

https://t.me/learningnets

```
rec time D54B4732.E24DD560 (06:20:34.884 PST Sat May 25 2013)
xmt time D54B4732.E24DD560 (06:20:34.884 PST Sat May 25 2013)
filtdelay =     1.00    1.00    1.00    0.00    1.00    1.00    1.00    0.00
filtoffset =    0.50    0.50    0.50    0.00    0.50    0.50    0.50    0.00
filterror =     1.95    2.98    4.02    4.86    4.89    4.92    4.95    4.98
minpoll = 6, maxpoll = 10

R2#
```

Note	To obtain a comprehensive view of the configuration tasks in this section, access the Mentor Guide engine. You can enter into the engine more than 1000 Cisco IOS Software commands, as well as a collection of proprietary commands such as **show all**.

https://t.me/learningnets