Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

© 2014 Narbik Kocharians. All rights reserved

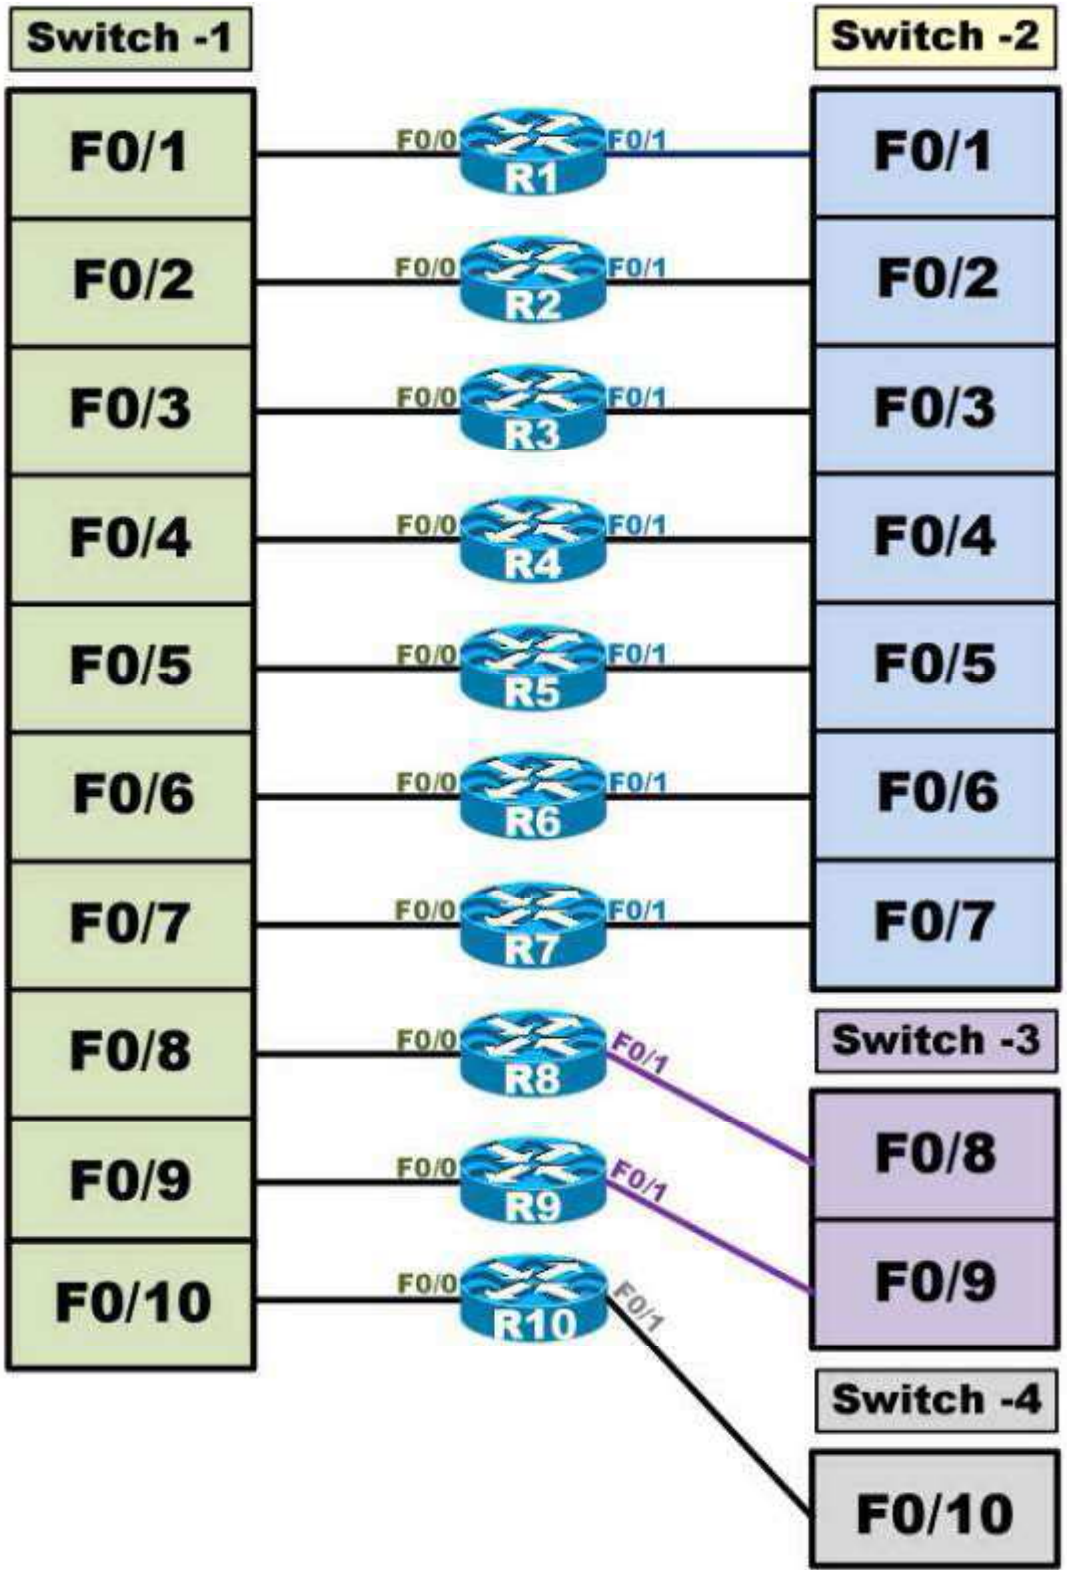
https://t.me/learningnets

Table of content

Subject	Page
Topology	3
Switching	
Lab 1 Troubleshooting Switching Scenario #1	8
Lab 2 Troubleshooting Switching Scenario #2	43
RIPv2	
Lab 1 TS RIPv2 Scenario #1	73
EIGRP	
Lab 1 TS EIGRP Scenario #1	94
OSPF	
Lab 1 TS OSPF Scenario #1	120
BGP	
Lab 1 TS BGP Scenario #1	147
MPLS	
Lab 1 TS MPLS Scenario #1	176
Lab 2 TS MPLS Scenario #2	194
Troubleshooting Mock Labs	
Lab 1 TS-Mock-Lab-1	217
Lab 2 TS-Mock-Lab-2	265
Lab 3 TS-Mock-Lab-3	309
Lab 4 TS-Mock-Lab-4	356
Lab 5 TS-Mock-Lab-5	402

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Serial connections:

Router	Interface	Router	Interface
R1	S1/2	R2	S1/1
R1	S1/3	R3	S1/1
R1	S1/4	R4	S1/1
R1	S1/5	R5	S1/1
R1	S1/6	R6	S1/1
R2	S1/1	R1	S1/2
R2	S1/3	R3	S1/2
R2	S1/4	R4	S1/2
R2	S1/5	R5	S1/2
R2	S1/6	R6	S1/2
R3	S1/1	R1	S1/3
R3	S1/2	R2	S1/3
R3	S1/4	R4	S1/3
R3	S1/5	R5	S1/3
R3	S1/6	R6	S1/3
R4	S1/1	R1	S1/4
R4	S1/2	R2	S1/4
R4	S1/3	R3	S1/4
R4	S1/5	R5	S1/4
R4	S1/6	R6	S1/4
R5	S1/1	R1	S1/5
R5	S1/2	R2	S1/5
R5	S1/3	R3	S1/5
R5	S1/4	R4	S1/5
R5	S1/6	R6	S1/5
R6	S1/1	R1	S1/6
R6	S1/2	R2	S1/6
R6	S1/3	R3	S1/6
R6	S1/4	R4	S1/6
R6	S1/5	R5	S1/6

© 2014 Narbik Kocharians. All rights reserved

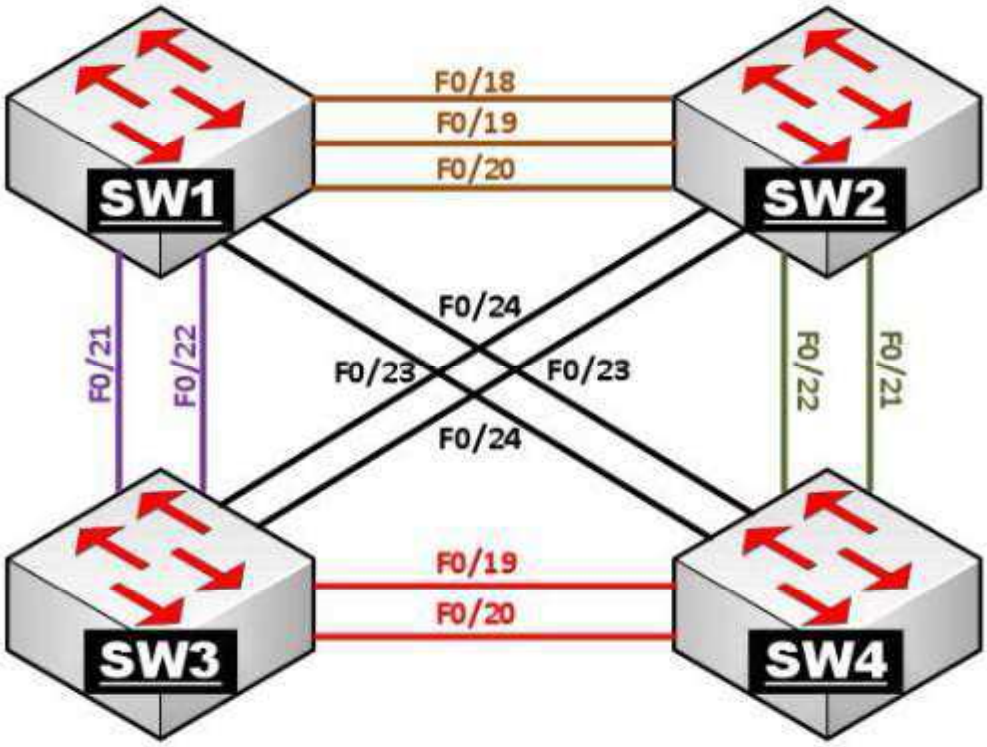
https://t.me/learningnets

R9	S0/0/0	R10	S0/0/0
R9	S0/1/0	R10	S0/1/0
R10	S0/0/0	R9	S0/0/0
R10	S0/1/0	R9	S0/1/0

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Switch to Switch connections



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

Switching

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Troubleshooting Switching
Scenario 1

Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "Switching" → "TS-Switching-Initial-Lab-1".

Lab Rules:

- DO NOT remove any command/s unless otherwise stated
- DO NOT change the VLAN assignment
- You must be VERY specific when resolving these tasks

Ticket 1

R1 and R2 can NOT establish an OSPF session. OSPF is configured on the F0/0 and the loopback 0 interfaces of these two routers.

Let's verify the problem:

On R1:

```
R1#Show ip ospf neighbor
R1#
```

Let's verify reachablility by Pinging R2 from R1 but before we can ping let's find out the IP address of their F0/0 interface:

```
R1#Show run int f0/0 | B interface

interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.0
 duplex auto
 speed auto
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
end
```

On R2:

```
R2#Show run int f0/0 | B interface

interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 duplex auto
 speed auto
end
```

On R1:

```
R1#Ping 10.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Ping is successful. Is OSPF configured on the correct interfaces? Let's verify:

On R1:

```
R1#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/24         1     LOOP  0/0
Fa0/0        1     0               10.1.1.1/24        1     DR    0/0
```

On R2:

```
R2#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               2.2.2.2/24         1     LOOP  0/0
Fa0/0        1     0               10.1.1.2/24        1     DR    0/0
```

It seems like OSPF was configured on the correct interfaces. Let's verify the configuration of OSPF on these two routers:

```
R2#Show run | s router ospf
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
router ospf 1
 network 2.2.2.2 0.0.0.0 area 0
 network 10.1.1.2 0.0.0.0 area 0
```

On R1:

```
R1#Show run | s router ospf

router ospf 1
 network 1.1.1.1 0.0.0.0 area 0
 network 10.1.1.1 0.0.0.0 area 0
```

No filtering whatsoever. Let's check SW1 and see if the VLANs are configured:

On SW1:

```
SW1#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/4, Fa0/6, Fa0/8, Fa0/9
                                                Fa0/10, Fa0/17, Fa0/18, Fa0/20
                                                Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2
21   VLAN0021                         active    Fa0/7
34   VLAN0034                         active    Fa0/3
56   VLAN0056                         active    Fa0/5
186  VLAN0186                         active    Fa0/13
196  VLAN0196                         active    Fa0/11
296  VLAN0296                         active    Fa0/12
339  VLAN0339                         active    Fa0/14
691  VLAN0691                         active    Fa0/15
946  VLAN0946                         active    Fa0/16
```

Let's check the configuration of F0/1 and F0/2:

```
SW1#Show run int f0/1 | B inter

interface FastEthernet0/1
 switchport access vlan 12
 switchport mode access
 spanning-tree portfast
end

SW1#Show run int f0/2 | B inter
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
interface FastEthernet0/2
 switchport access vlan 12
 switchport mode access
 spanning-tree portfast
end
```

Are these routers connected to the correct switch? May be there is a problem with the diagram:

On SW1:

```
SW1#Show cdp neighbors | Inc R1|R2

R2               Fas 0/2             154             R S I      2811       Fas 0/0
R1               Fas 0/1             124             R S I      2811       Fas 0/0
```

Let's see if there is an access-list or an access-map configured on this switch:

```
SW1#Show access-list

Extended IP access list 100
    10 permit icmp any any
```

Sure enough an access-list is configured, but it is ONLY allowing ICMP, let's see what configuration on this switch is calling the access-list:

```
SW1#Show run | inc 100

 match ip address 100
access-list 100 permit icmp any any
```

There must be a route-map or a VLAN Access-map referencing access-list 100, let's verify:

```
SW1#Show run | i route-map
SW1#

SW1#Show run | Inc vlan access-map
vlan access-map tst 10
```

Let's check this vlan access-map:

```
SW1#Show run | B vlan access-map

vlan access-map tst 10
 action forward
 match ip address 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
vlan filter tst vlan-list 1-4094
(The rest of the output is omitted)
```

wow...we can clearly see the problem, the VLAN access-map is referencing access-list 100 which ONLY permits ICMP and forwards it, and since there is no other statement that permits the rest of the traffic, OSPF traffic is dropped. Since we can NOT remove any configuration and we have to be very specific when we resolve this ticket, let's permit OSPF on access-list 100:

```
SW1(config)#Access-list 100 permit ospf any any
```

Let's check R1:

On R1:

Once we connect to R1's console we should see the following message that is confirming OSPF's adjacency:

```
%OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to
FULL, Loading Done

R1#Show ip route ospf

     2.0.0.0/32 is subnetted, 1 subnets
O       2.2.2.2 [110/2] via 10.1.1.2, 00:01:15, FastEthernet0/0
```

Perfect.

Ticket 2

R3 and R4 can NOT Ping each other. These two routers are in VLAN 34, R3 is connected to SW1 and R4 is connected to SW2, the F0/19 interfaces of these two switches are configured as a trunk link. Do not remove any command, configure another trunk or change the VLAN assignment to fix this problem.

Let's confirm the problem:

On R3:

```
R3#Ping 34.1.1.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the ARP table:

```
R3#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  34.1.1.3                -  0007.b35b.8310  ARPA   FastEthernet0/0
Internet  34.1.1.4                0  Incomplete      ARPA
```

Let's check and see if there is an access-list configured on R3 or R4:

On R3:

```
R3#Show access-list
R3#
```

On R4:

```
R4#Show access-list
R4#
```

Let's check the switches (SW1 and SW2):

On SW1:

```
SW1#Show run int f0/3 | B interface

interface FastEthernet0/3
 switchport access vlan 34
 switchport mode access
 spanning-tree portfast
end
```

On SW2:

```
SW2#Show run int f0/4 | B interface

interface FastEthernet0/4
 switchport access vlan 34
 switchport mode access
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 spanning-tree portfast
end
```

Let's check the trunk:

On SW1:

```
SW1#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,12,21,34,56,186,196,296,339,691,946

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,12,21,34,56,186,196,296,339,691,946
```

On SW2:

```
SW2#Show interface F0/19 trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           isl            trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,33,56,152,339,691

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,33,56,152,339,691
```

Why don't we see VLAN 34 in the "VLANs allowed and active in management domain"?

Let's check the VLANs on SW2

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/7
                                                Fa0/8, Fa0/9, Fa0/10, Fa0/13
                                                Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                                Fa0/18, Fa0/20, Fa0/21, Fa0/22
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
21   VLAN0021                         active    Fa0/12
33   VLAN0033                         active    Fa0/11
56   VLAN0056                         active    Fa0/6
152  VLAN0152                         active    Fa0/5
```

Earlier we saw that the F0/4 interface was configured in VLAN 34, but it does not show up in the list of VLANs, let's repeat the same command and see all the VLANs (supported and unsupported):

On SW2:

```
SW2#Show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/16, Fa0/17, Fa0/18
                                                Fa0/20, Fa0/21, Fa0/22, Gi0/1
                                                Gi0/2
33   VLAN0033                         active    Fa0/7
34   VLAN0034                         act/unsup Fa0/4
56   VLAN0056                         active    Fa0/6
152  VLAN0152                         active    Fa0/5
339  VLAN0339                         active    Fa0/14
691  VLAN0691                         active    Fa0/15
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

Sure enough we can see the VLAN but why does it have an "act/unsup" status? Let's verify:

On SW2:

```
SW2#Show vlan id 34

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
34   VLAN0034                         act/unsup Fa0/4, Fa0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
34   fddi  100034     1500  -      -      -        -    -        0      0

Remote SPAN VLAN
----------------
Disabled

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
```

We can clearly see the problem, the media type was changed, let's correct the problem and verify:

On SW2:

```
SW2(config)#Vlan 34
SW2(config-vlan)#media ethernet
SW2(config-vlan)#Exit
```

To verify the configuration:

On SW2:

```
SW2#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/16, Fa0/17, Fa0/18
                                                Fa0/20, Fa0/21, Fa0/22, Gi0/1
                                                Gi0/2
33   VLAN0033                         active    Fa0/7
34   VLAN0034                         active    Fa0/4
56   VLAN0056                         active    Fa0/6
152  VLAN0152                         active    Fa0/5
339  VLAN0339                         active    Fa0/14
691  VLAN0691                         active    Fa0/15
```

This looks much better, let's verify VLAN 34 on the trunk:

```
SW2#Show interface F0/19 trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           isl            trunking      1

Port        Vlans allowed on trunk
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,21,33-34,56,152,339,691

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,21,33-34,56,152,339,691
```

Let's test the solution:

On R3:

```
R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the ARP table:

```
R3#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  34.1.1.3               -   0007.b35b.8310  ARPA   FastEthernet0/0
Internet  34.1.1.4               0   Incomplete      ARPA
```

Why can't we ping? The trunk is up, the VLAN is in act state, what else can be wrong? Let's check the trunks in detail one more time:

On SW1:

```
SW1#Sh interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,12,21,34,56,186,196,296,339,691,946

Port        Vlans in spanning tree forwarding state and not pruned
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/19      1,12,21,34,56,186,196,296,339,691,946
```

Let's summarize the output of the above show command:

- The trunk interface is F0/19
- Mode is ON
- Encapsulation is 802.1q
- It is trunking
- Native VLAN is 1

Let's check the trunk on SW2

On SW2:

```
SW2#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           isl            trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,21,33-34,56,152,339,691

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,21,33-34,56,152,339,691
```

We can see the problem, SW1 is configured with an encapsulation of 802.1q, whereas, SW2 is configured with an encapsulation of ISL. Let's change SW2's encapsulation to 802.1q:

On SW2:

```
SW2(config)#Int F0/19
SW2(config-if)#Swi trunk encapsulation dot1q
```

Let's test the solution again, but we have to wait 30 seconds for STP to converge:

On R3:

```
R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms

R3#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  34.1.1.3                -  0007.b35b.8310  ARPA   FastEthernet0/0
Internet  34.1.1.4                4  0015.fa03.c3a1  ARPA   FastEthernet0/0
```

Perfect..

Ticket 3

R5 and R6 can't establish an Eigrp adjacency. Eigrp AS 100 is configured on the F0/0 of R5 and the F0/1 interface of R6; R5 is advertising networks 5.5.5.0/24 and R6 is advertising 6.6.6.0/24. You must be very specific, do not remove any command to fix this problem.

Let's verify the problem:

On R5:

```
R5#Show ip eigrp neighbor
IP-EIGRP neighbors for process 100
R5#
```

Let's check the interface configuration of R5:

```
R5#Show run int f0/0 | B inter

interface FastEthernet0/0
 ip address 56.1.1.5 255.255.255.0
 duplex auto
 speed auto
end
```

The interface is configured with a an IP address of 56.1.1.5/24 and there is no access-list. Let's verify if Eigrp is configured on the correct interface:

```
R5#Show ip eigrp interface
EIGRP-IPv4 Interfaces for AS(100)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                         Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers   Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo0                0       0/0        0        0/1           0           0
Fa0/0              0       0/0        0        0/1          50           0
```

Can we Ping 56.1.1.6?

```
R5#Ping 56.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 56.1.1.6, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Let's verify the switchport that R5 is connected to:

On SW1:

```
SW1#Sh run int f0/5 | B inter

interface FastEthernet0/5
 switchport access vlan 56
 switchport mode access
end
```

The configuration is correct, let's verify the VLAN on the trunk:

```
SW1#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,12,21,34,56,186,196,296,339,691,946

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,12,21,34,56,186,196,296,339,691,946
```

Everything looks to be correctly configured. Let's check SW2:

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show inter trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1
Fa0/24      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,21,33-34,56,152,339,691
Fa0/23      1,21,33-34,56,152,339,691
Fa0/24      1,21,33-34,56,152,339,691

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,21,33-34,56,152,339,691
Fa0/23      1,21,33-34,56,152,339,691
Fa0/24      21,34,56,152
```

The trunk is configured to carry the VLAN, but we have two additional trunks, let's see which port is forewarding the traffic for VLAN 56:

```
SW2#Show spanning-tree Vlan 56

VLAN0056
  Spanning tree enabled protocol ieee
  Root ID     Priority    32824
              Address     0014.a934.b880
              Cost        19
              Port        21 (FastEthernet0/19)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32824  (priority 32768 sys-id-ext 56)
              Address     001b.2be5.1200
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/6            Desg FWD 19        128.8    Edge P2p
Fa0/19           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.25   P2p
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/24           Desg FWD 19        128.26    P2p
```

The port that R6 is connected to is in FWD state, the root port for this VLAN is F0/19 interface, if this is correct, then, SW1 must be the root for this Vlan. Let's verify:

On SW1:

```
SW1#Show spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ----------
VLAN0001         32769 000c.302d.9980        57    2   20  15  Fa0/19
VLAN0012         32780 0012.7f40.9380         0    2   20  15
VLAN0021         32789 0012.7f40.9380         0    2   20  15
VLAN0034         32802 0012.7f40.9380         0    2   20  15
VLAN0056         32824 0012.7f40.9380         0    2   20  15
VLAN0186         32954 0012.7f40.9380         0    2   20  15
VLAN0196         32964 0012.7f40.9380         0    2   20  15
VLAN0296         33064 0012.7f40.9380         0    2   20  15
VLAN0339         33107 000c.302d.9980        57    2   20  15  Fa0/19
VLAN0691         33459 000c.302d.9980        57    2   20  15  Fa0/19
VLAN0946         33714 0012.7f40.9380         0    2   20  15
```

Perfect, let's verify the configuration of SW2's F0/6 interface:

On SW2:

```
SW2#Show run int f0/6 | B inter

interface FastEthernet0/6
 switchport access vlan 56
 switchport mode access
 storm-control multicast level 0.00
 spanning-tree portfast
end
```

We see the problem, the "Storm-control" command is blocking all Multicast traffic by setting its level to 0.00, but since we can NOT configure any of the switches to resolve this ticket, let's configure the routers to establish an adjacency using unicast:

On R5:

```
R5(config)#Router eigrp 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config-router)#Neighbor 56.1.1.6 F0/0
```

On R6:

```
R6(config)#Router eigrp 100
R6(config-router)#Neighbor 56.1.1.5 F0/1
```

Let's verify the configuration:

On R6:

```
R6#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
R5#
```

What else can be wrong?

The routers are NOT configured with an access-list, Let's enable "debug ip packet detail" and tie an access-list to it and analyze the output:

On R5:

Before configuring any access-list, the "Show access-list" must be used to confirm that no other access-list is using the same number.

```
R5#Show access-list
R5#

R5(config)#access-list 100 permit eigrp host 56.1.1.6 host 56.1.1.5
R5(config)#access-list 100 permit eigrp host 56.1.1.5 host 56.1.1.6

R5#Debug ip packet det 100
IP packet debugging is on (detailed) for access list 100

IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
```

It looks like R5 is sending unicast hellos but not receiving anything from R6. Let's do the same on R6:

On R6:

```
R6#Show access-list
R6#

R6(config)#access-list 100 permit eigrp host 56.1.1.6 host 56.1.1.5
R6(config)#access-list 100 permit eigrp host 56.1.1.5 host 56.1.1.6

R6#Debug ip packet det 100
IP packet debugging is on (detailed) for access list 100

s=56.1.1.6 (local), d=56.1.1.5 (FastEthernet0/1), len 60, sending,
proto=88
s=56.1.1.6 (local), d=56.1.1.5 (FastEthernet0/1), len 60, sending full
packet, proto=88
s=56.1.1.6 (local), d=56.1.1.5 (FastEthernet0/1), len 60, sending,
proto=88
s=56.1.1.6 (local), d=56.1.1.5 (FastEthernet0/1), len 60, sending full
packet, proto=88
s=56.1.1.6 (local), d=56.1.1.5 (FastEthernet0/1), len 60, sending,
proto=88
s=56.1.1.6 (local), d=56.1.1.5 (FastEthernet0/1), len 60, sending full
packet, proto=88
```

It looks like R5 is sending unicast Hellos to R6 but not receiving any, and R6 is sending Unicast Hellos to R5 but it is NOT receiving any from R5.

Let's check the switches for an access-list:

On SW1:

```
SW1#Show access-list

Extended IP access list 100
    10 permit icmp any any
    20 permit ospf any any
```

AH.....This access-list was referenced by a Vlan access-map and it is ONLY allowing ICMP and OSPF and NOT Eigrp. Let's add Eigrp to the access-list and verify again:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1

```
SW1(config)#Access-list 100 permit eigrp any any
```

To verify the configuration:

On R5:

In the output of the debug, you should see that you have established Eigrp adjacency and you are sending and receiving Eigrp packets from R6:

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 56.1.1.6 (FastEthernet0/0) is
up: new adjacency

IP: tableid=0, s=56.1.1.6 (FastEthernet0/0), d=56.1.1.5 (FastEthernet0/0),
routed via RIB

IP: s=56.1.1.6 (FastEthernet0/0), d=56.1.1.5 (FastEthernet0/0), len 60,
rcvd 3, proto=88

IP: s=56.1.1.5 (local), d=56.1.1.6 (FastEthernet0/0), len 60, sending,
proto=88
```

```
R5#Show ip eigrp neighbor
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   56.1.1.6                Fa0/0             11 00:00:49    6   200  0  3
```

```
R5#Show ip route Eigrp

     6.0.0.0/24 is subnetted, 1 subnets
D       6.6.6.0 [90/156160] via 56.1.1.6, 00:01:56, FastEthernet0/0
```

On Both Routers:

```
R6#U all

R5(config)#No access-list 100
```

Perfect.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 4

R7 can NOT ping R8's G0/1.21 interface.

Before the ticket is verified, let's find out the G0/1.21 sub-interface of R8:

On R8:

```
R8#Show run int g0/1.21 | B interface

interface GigabitEthernet0/1.21
 encapsulation dot1Q 21
 ip address 12.1.1.2 255.255.255.0
end
```

On R7:

```
R7#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's verify the ARP table:

```
R7#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  12.1.1.1               -   24e9.b3ab.4b20  ARPA   GigabitEthernet0/0
Internet  12.1.1.2               0   Incomplete      ARPA
Internet  33.3.3.1               -   24e9.b3ab.4b21  ARPA   GigabitEthernet0/1
```

Let's check SW1:

On SW1:

```
SW1#Show interface F0/7 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/7                        connected    21         a-full  a-100 10/100BaseTX

SW1#Show vlan brie | Exc unsup
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/4, Fa0/6, Fa0/8, Fa0/9
                                                Fa0/10, Fa0/17, Fa0/18, Fa0/20
                                                Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2

21   VLAN0021                         active    Fa0/7
34   VLAN0034                         active    Fa0/3
56   VLAN0056                         active    Fa0/5
186  VLAN0186                         active    Fa0/13
196  VLAN0196                         active    Fa0/11
296  VLAN0296                         active    Fa0/12
339  VLAN0339                         active    Fa0/14
691  VLAN0691                         active    Fa0/15
946  VLAN0946                         active    Fa0/16
```

The output of the above show command reveals that VLAN 21 is active and port F0/7 which is connected to R7 is part of this Vlan.

```
SW1#Show run int F0/7 | B interface

interface FastEthernet0/11
 switchport access vlan 21
 switchport mode access
 spanning-tree portfast
end
```

Let's verify the trunk link:

```
SW1#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,12,21,34,56,186,196,296,339,691,946

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,12,21,34,56,186,196,296,339,691,946
```

Interface F0/19 is the trunk that carries traffic. Let's see if STP is using this port:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show spanning-tree Vlan 21

VLAN0021
  Spanning tree enabled protocol ieee
  Root ID     Priority    32789
              Address     0012.7f40.9380
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32789  (priority 32768 sys-id-ext 21)
              Address     0012.7f40.9380
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------
Fa0/7            Desg FWD 19        128.9    Edge P2p
Fa0/19           Desg FWD 19        128.21   P2p
```

Well.....the interface that R7 is connected to is forwarding for VLAN 21, and so is the F0/19 interface. Sicne this interface is connected to SW2, let's go to SW2 and verify this information. Remember that R8's G0/1 interface is connected to SW3's F0/8.

On SW2:

```
SW2#Show inter trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1
Fa0/24      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/23      1-4094
Fa0/24      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,21,33-34,56,152,339,691
Fa0/23      1,21,33-34,56,152,339,691
Fa0/24      1,21,33-34,56,152,339,691

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,21,33-34,56,152,339,691
Fa0/23      1,21,33-34,56,152,339,691
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/24      21,34,56,152
```

This is starting to make sense, SW2 is using interfaces F0/23 and/or F0/24 to reach SW3. Let's see if STP is using these interfaces:

```
SW2#Show spanning-tree Vlan 21 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.25   P2p
Fa0/24           Desg FWD 19        128.26   P2p
```

The traffic for VLAN 21 comes in through F0/19 and uses ports F0/23 and F0/24 to reach SW3. Let's verify SW3's configuration:

On SW3:

```
SW3#Show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
SW4              Fas 0/19          144           S I      WS-C3550-2Fas 0/19
SW2              Fas 0/24          156           S I      WS-C3560-2Fas 0/24
SW2              Fas 0/23          156           S I      WS-C3560-2Fas 0/23
```

We don't even see R8, based on our topology diagram, R8's G0/1 interface should be connected to SW3's F0/8 interface. Let's check R8's configuration:

On R8:

```
R8#Show run int G0/1.21 | B interface

interface GigabitEthernet0/1.21
 encapsulation dot1Q 21
 ip address 12.1.1.2 255.255.255.0
end
```

R8 is configured as a trunk, let's check the interface of SW3 that R8 is connected to:

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3#Show run int f0/8 | B inter

interface FastEthernet0/12
 switchport trunk encapsulation isl
 switchport mode trunk
end
```

The port on SW3 is also configured as a trunk, but it is using an ISL encapsulation and NOT 802.1q, let's change the encapsulation on SW3 to match the router:

On SW3:

```
SW3(config)#int f0/8
SW3(config-if)#Switch trunk encapsulation dot
```

Let's clear the cdp table, wait 60 seconds and verify the connectivity using CDP:

```
SW3#Clear cdp table

SW3#Show cdp neighbor
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
SW4              Fas 0/19          154           S I      WS-C3550-2Fas 0/19
SW2              Fas 0/24          169           S I      WS-C3560-2Fas 0/24
SW2              Fas 0/23          169           S I      WS-C3560-2Fas 0/23
R8               Fas 0/8           156         R B S I    CISCO1921/Gig 0/1
```

Let's verify and test the solution:

On R8:

```
R8#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  12.1.1.1               0   Incomplete      ARPA
Internet  12.1.1.2               -   000e.d774.2c21  ARPA   FastEthernet0/1.21
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Internet  152.1.1.2               -   000e.d774.2c21  ARPA   FastEthernet0/1.152
```

Does Vlan 21 exist on SW3:

On SW3:

```
SW3#Show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------

1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/10
                                                Fa0/17, Fa0/18, Fa0/20, Fa0/21
                                                Fa0/22, Gi0/1, Gi0/2
33   VLAN0033                         active    Fa0/9
186  VLAN0186                         active    Fa0/13
196  VLAN0196                         active    Fa0/11
296  VLAN0296                         active    Fa0/12
339  VLAN0339                         active    Fa0/14
691  VLAN0691                         active    Fa0/15
946  VLAN0946                         active    Fa0/16
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

No, it does not, let's configure Vlan 21 and wait 30 seconds before testing:

```
SW3(config)#Vlan 21
SW3(config-vlan)#Exit
```

If the "Exit" command is NOT used, the Vlan is not created. Let's verify reachability:

On R8:

```
R8#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Perfect.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 5

R9 can NOT ping R7's G0/1 interface

Let's find out the IP address of R7's G0/1 interface:

On R7:

```
R7#Show run int g0/1 | B interface

interface GigabitEthernet0/1
 ip address 33.3.3.1 255.255.255.0
 duplex auto
 speed auto
end
```

To verify the problem

On R9:

```
R9#Ping 33.3.3.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.3.3.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R9#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  33.3.3.3                -   000b.be03.54c1  ARPA   FastEthernet0/1
```

Is this interface up?

```
R9#Show ip int brief F0/1

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/1        33.3.3.3        YES manual up                    down
```

Why is this interface down? Let's check the switchport on SW3 that this router is connected to:

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3#Show run int f0/9 | B inter

interface FastEthernet0/9
 switchport access vlan 33
 switchport mode access
 speed 10
 spanning-tree portfast
end
```

This looks correct, let's check the status of the F0/9 interface on this switch:

```
SW3#Show interface F0/9 Status

Port      Name      Status       Vlan       Duplex  Speed Type
Fa0/13              notconnect   33           auto     10 10/100BaseTX
```

Let's check the G0/1 interface of R9:

On R9:

```
R9#Show run int f0/1 | B inter

interface FastEthernet0/1
 ip address 33.3.3.3 255.255.255.0
 duplex auto
 speed 100
end
```

We can clearly see the problem, R9's F0/1 is configured with a speed of 100 Mbps, whereas, the F0/9 interface on SW3 is configured with 10 Mbps. Let's configure the F0/9 interface of SW3 with a speed of 100 Mbps:

On SW3:

```
SW3(config)#Int F0/9
SW3(config-if)#Speed 100

Interface FastEthernet0/9, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/9, changed
state to up
```

The above console message state that the F0/9 interface of the local switch is now in UP/UP state.

Let's try to confirm reachability:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R9:

```
R9#Ping 33.3.3.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33.3.3.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 6

The Etherchannel between SW2 and SW3 is NOT coming up.

To verify the problem:

On SW2:

```
SW2#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
23     Po23(SD)         PAgP      Fa0/23(I)   Fa0/24(I)
```

The letter "S" in the "Port-Channel" colume tells us that the Port Channel interface is a layer 2 interface and the letter "D" indicates that the status is down.

Let's check Sw3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW3:

```
SW3#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------
23     Po23(SD)         PAgP      Fa0/23(I)   Fa0/24(I)
```

Looking at the output of the above show command, both switches are using PAgP, and the Etherchannel is a layer 2 Etherchannel, BUT it is down (The letter "D" in the SD indicates that it is down). Let's check the configuration of these two switches:

On SW2:

```
SW2#Show run int F0/23 | B interface

interface FastEthernet0/23
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 23 mode auto
end

SW3#Show run int F0/24 | B interface

interface FastEthernet0/24
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 23 mode auto
end
```

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3#Show run int F0/23 | B interface

interface FastEthernet0/23
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 23 mode auto
end

SW3#Show run int F0/24 | B interface

interface FastEthernet0/24
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 23 mode auto
end
```

The output of the above configuration reveals that both switches are configured in "Auto" mode, in "Auto" mode none of the switches initiate the negotiation process. To fix this problem, only one side needs to be changed:

On SW2:

```
SW2(config)#int range f0/23-24
SW2(config-if-range)#Channel-group 23 mode desirable
```

You should see the following console messages:

```
Line protocol on Interface FastEthernet0/23, changed state to down
Line protocol on Interface FastEthernet0/24, changed state to down
Line protocol on Interface FastEthernet0/23, changed state to up
Line protocol on Interface FastEthernet0/24, changed state to up
Interface Port-channel23, changed state to up
Line protocol on Interface Port-channel23, changed state to up
```

To verify the configuration:

On SW2:

```
SW2#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1


Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
23     Po23(SU)         PAgP      Fa0/23(P)   Fa0/24(P)
```

On SW3:

```
SW3#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
23     Po23(SU)         PAgP      Fa0/23(P)   Fa0/24(P)
```

The letters "SU" indicate that the port channel is a layer two and it's in use.

Ticket 7

R8 can NOT ping R5's F0/1 interface.

Let's find out the IP address of R5's F0/1 interface before verifying the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R5:

```
R5#Show run int f0/1 | B interface

interface FastEthernet0/1
 ip address 152.1.1.5 255.255.255.0
 duplex auto
 speed auto
end
```

On R8:

```
R8#Ping 152.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 152.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see R8's G0/1 configuration:

```
R8#Show ip int br | Exc unass

Interface                  IP-Address      OK? Method Status         Protocol
GigabitEthernet0/1.21      12.1.1.2        YES manual up             up
GigabitEthernet0/1.152     152.1.1.2       YES manual up             up
```

Let's check and see if SW3 can see R8:

On SW3:

```
SW3#Show cdp nei | Inc R8

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
SW4              Fas 0/19          175           S I      WS-C3550-2Fas 0/19
SW2              Fas 0/24          130           S I      WS-C3560-2Fas 0/24
SW2              Fas 0/23          130           S I      WS-C3560-2Fas 0/23
R8               Fas 0/8           177         R B S I    CISCO1921/Gig 0/1
R9               Fas 0/9           178          R S I     2811       Fas 0/1
```

Let's see if Vlan 152 is configured on SW3:

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/10
                                                Fa0/17, Fa0/18, Fa0/20, Fa0/21
                                                Fa0/22, Gi0/1, Gi0/2
21   VLAN0021                         active
33   VLAN0033                         active    Fa0/9
186  VLAN0186                         active    Fa0/13
196  VLAN0196                         active    Fa0/11
296  VLAN0296                         active    Fa0/12
339  VLAN0339                         active    Fa0/14
691  VLAN0691                         active    Fa0/15
946  VLAN0946                         active    Fa0/16
```

VLAN 152 is NOT configured on SW3, let's configure this VLAN:

```
SW3(config)#Vlan 152
SW3(config-vlan)#Exit
```

Let's wait for Spanning-tree to converge before testing the configuration:

To verify the configuration:

On R8:

```
R8#Ping 152.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 152.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 8

You must "Shut" and "No Shut" the F0/19 interface of SW3 and SW4 before proceeding, when shutting the f0/19 interface down, you must wait for the interface to transition into down state before issuing the "No Shut" command.



The client keeps on getting the following console messages on SW3:

%UDLD-4-UDLD_PORT_DISABLED: UDLD disabled interface Fa0/19, unidirectional link detected
%PM-4-ERR_DISABLE: udld error detected on Fa0/19, putting Fa0/19 in err-disable state
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/19, changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/19, changed state to down

%PM-4-ERR_RECOVER: Attempting to recover from udld err-disable state on Fa0/19

%LINK-3-UPDOWN: Interface FastEthernet0/19, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/19, changed state to up

You may break one of the rules in this lab's rules stated in the beginning of this lab.

To verify the problem:

On SW3:

```
SW3#Show udld F0/19

Interface Fa0/19
---
Port enable administrative configuration setting: Enabled / in aggressive
mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Unidirectional
Current operational state: Disabled port
Message interval: 7
Time out interval: 5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
No neighbor cache information stored
```

The output of the above show command states that a Unidirectional link is what is detected and the console messages verify this information. But Why?

Let's check the interface configuration of SW3 and SW4; since we can NOT physically inspect or touch the pods in this lab or in the actual CCIE lab, how do we fix this problem?

```
SW3#Show run int f0/19 | B inter

interface FastEthernet0/19
 switchport trunk encapsulation dot1q
 switchport mode trunk
 udld port aggressive
end
```

On SW4:

```
SW4#Show run int f0/19 | B inter

interface FastEthernet0/19
 switchport trunk encapsulation dot1q
 switchport mode trunk
 udld port aggressive

 mac access-group tst in
end
```

Hmmm.....let's check the mac access-list called "tst":

```
SW4#Show run | B mac access-list

mac access-list extended tst
 deny   any host 0100.0ccc.cccc
 permit any any
```

WOW.... The MAC that is being filtered on ingress F0/19 interface of SW3 is the destination Mac address used by UDLD. In order to fix this problem, we should just remove the mac access-list all together.

On SW4

```
SW4(config)#No mac access-list extended tst
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW3

```
SW3#Show udld f0/19

Interface Fa0/19
---
Port enable administrative configuration setting: Enabled / in aggressive
mode

Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Bidirectional

Current operational state: Advertisement - Single neighbor detected
Message interval: 7
Time out interval: 5
    Entry 1
    ---
    Expiration time: 43
    Cache Device index: 1
    Current neighbor state: Bidirectional
    Device ID: CAT0629X0DG
    Port ID: Fa0/19
    Neighbor echo 1 device: CAT0628Z0LB
    Neighbor echo 1 port: Fa0/19

    Message interval: 15
    Time out interval: 5
    CDP Device name: SW4
```

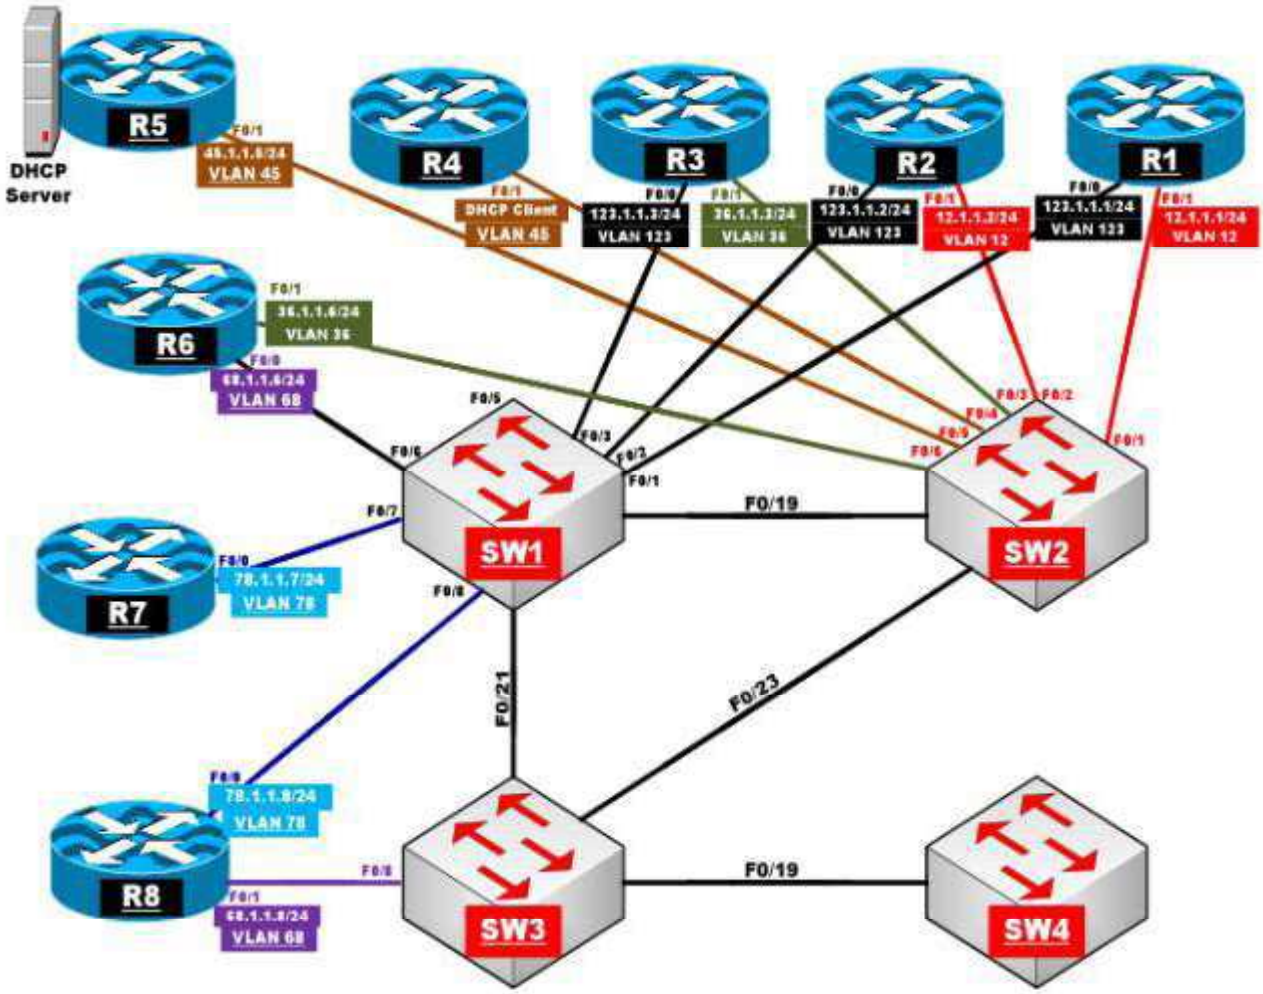
Ticket 9

Erase the startup configuration and the "vlan.dat" and reload the devices before proceeding to the next scenario.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Troubleshooting Switching
Scenario 2



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "Switching" →"TS-Switching-Initial-Lab-2".

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Rules:

- DO NOT configure or change the VTP domain on any of the four switches.
- DO NOT change VLAN assignments

Ticket 1

R6 and R8 can not establish an Eigrp neighbor adjacency. DO NOT change the VLAN assignment, or remove a command by negating it to fix this problem.

Let's verify the problem:

On R6:

```
R6#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(1)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   68.1.1.8                Fa0/0             14 00:04:31    2   200  0  6
```

It looks like the two routers have established a neighbor adjacency. Let's ping 68.1.1.8 to verify reachability.

```
R6#Ping 68.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 68.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Are these two routers on the same VLAN? Let's check SW1 and SW3:

On SW1:

```
SW1#Show Vlan br | Exc unsup

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Fa0/4, Fa0/5, Fa0/9, Fa0/10
                                                 Fa0/11, Fa0/18, Fa0/20, Fa0/22
                                                 Fa0/23, Fa0/24, Gi0/1, Gi0/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
68   VLAN0068                         active    Fa0/6
212  VLAN0212                         active    Fa0/12
213  VLAN0213                         active    Fa0/13
214  VLAN0214                         active    Fa0/14
215  VLAN0215                         active    Fa0/15
216  VLAN0216                         active    Fa0/16
217  VLAN0217                         active    Fa0/17
```

On SW3:

```
SW3#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/9
                                                Fa0/10, Fa0/11, Fa0/18, Fa0/22
                                                Fa0/24, Gi0/1, Gi0/2
68   VLAN0068                         active    Fa0/8
212  VLAN0212                         active    Fa0/12
213  VLAN0213                         active    Fa0/13
214  VLAN0214                         active    Fa0/14
215  VLAN0215                         active    Fa0/15
216  VLAN0216                         active    Fa0/16
217  VLAN0217                         active    Fa0/17
```

R6 and R8 are in VLAN 68, let's check the trunk and see if the trunk is allowing the traffic for this VLAN to traverse through:

On SW1:

```
SW1#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/21      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/21      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,68,212-217
Fa0/21      1,68,212-217

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,68,212-217
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/21      1,68,212-217
```

On SW3:

```
SW3#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan

Fa0/19      desirable    n-802.1q       trunking      1
Fa0/20      desirable    n-isl          trunking      1

Fa0/21      on           802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/20      1-4094
Fa0/21      1-4094
Fa0/23      1-4094

Port        Vlans allowed and active in management domain
Fa0/19      1,68,212-217
Fa0/20      1,68,212-217
Fa0/21      1,68,212-217
Fa0/23      1,68,212-217

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1,68,212-217
Fa0/20      68
Fa0/21      1,68,212-217

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      1,68,212-217
```

So far so good, let's verify the path from SW1 to SW3:

On SW1:

```
SW1#Show spanning-tree vlan 68

VLAN0068
  Spanning tree enabled protocol ieee
  Root ID     Priority    24644
              Address     000c.858b.7a00
              Cost        38
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
             Port        21 (FastEthernet0/19)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32836  (priority 32768 sys-id-ext 68)
             Address     0012.7f40.9380
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ------
Fa0/6            Desg FWD 19        128.8    P2p
Fa0/19           Root FWD 19        128.21   P2p
Fa0/21           Desg FWD 19        128.23   P2p
```

The Fa0/6 interface of SW1 connects to R6, and from SW1's perspective, both F0/19 and F0/21 are in forwarding state, but F0/19 is the root port, and the local switch has a cost of 38, which means that it has to traverse two 100M links (The default cost of a Fast Ethernet interface is 19), this means that SW3 is the root. Let's verify this by repeating the same show command on SW2:

On SW2:

```
SW2#Show spanning-tree vlan 68

VLAN0068
  Spanning tree enabled protocol ieee
  Root ID     Priority    24644
              Address     000c.858b.7a00
              Cost        19
              Port        25 (FastEthernet0/23)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32836  (priority 32768 sys-id-ext 68)
              Address     001d.e5d6.0000
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/23           Root FWD 19        128.25   P2p
```

Let's verify this information on SW3:

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3#Show spanning-tree Vlan 68

VLAN0068
  Spanning tree enabled protocol ieee
  Root ID     Priority    24644
              Address     000c.858b.7a00
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24644  (priority 24576 sys-id-ext 68)
              Address     000c.858b.7a00

              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ------
Fa0/8            Desg FWD 19        128.8    P2p
Fa0/19           Desg FWD 19        128.19   P2p
Fa0/20           Desg FWD 19        128.20   P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.23   P2p
```

Why is SW1 going through SW2 to get to SW3, when the link (The F0/21) between SW1 and SW3 are trunking and in "FWD" state? This can mean that we have a spanning-tree loop. Let's check the F0/21 interfaces of SW1 and SW3:

On SW1:

```
SW1#Show run int f0/21 | B interface

interface FastEthernet0/21
 switchport trunk encapsulation dot1q
 switchport mode trunk
end
```

On SW3:

```
SW3#Show run int f0/21 | B interface

interface FastEthernet0/21
 switchport trunk encapsulation dot1q
 switchport mode trunk
 spanning-tree bpdufilter enable
end
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

We can see the problem, BPDUFilter is configured on the F0/21 interface of SW3, this does mean that we have a spanning-tree loop. Let's check the CPU utilization on SW3:

```
SW3#Show processes cpu | Inc five

CPU utilization for five seconds: 99%/93%; one minute: 99%; five minutes:92%
```

let's remove this and verify this configuration; once the BPDUFilter is removed, SW1 should use its F0/21 interface as its root port and either SW1 or SW2 will transition their F0/19 interface into "BLK" state for this VLAN:

On SW3:

```
SW3(config)#Int f0/21
SW3(config-if)#spanning-tree bpdufilter disable
```

On SW1:

```
SW1#Show spanning-tree vlan 68 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/6            Desg FWD 19        128.8    P2p
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/21           Root FWD 19        128.23   P2p
```

Let's check SW2, it has probably transitioned its F0/19 interface to blocking:

On SW2:

```
SW2#Show spanning-tree vlan 68 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ------
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/23           Root FWD 19        128.25   P2p
```

Let's check R6 and see if has established an Eigrp neighbor adjacency:

On R6:

```
R6#Show ip eigrp neighbors

IP-EIGRP neighbors for process 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
H   Address                 Interface       Hold Uptime    SRTT    RTO  Q  Seq
                                            (sec)          (ms)        Cnt Num
0   68.1.1.8                Fa0/0             12 00:04:49     3    200  0  15

R6#Ping 68.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 68.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On SW3:

```
SW3#Show processes cpu | Inc five

CPU utilization for five seconds: 0%/0%; one minute: 0%; five minutes: 8%
```

Ticket 2

SHUT and NO Shut the F0/1 interface on R4.

In VLAN 45, R5 is configured as a DHCP server, and R4 as a DHCP client. R4 can not acquire an IP address from this DHCP server. Fix this problem without removing any commands.

Let's verify the problem:

On R4:

```
R4#Show ip int brief

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        unassigned      YES unset  administratively down down
Serial0/0              unassigned      YES unset  administratively down down
FastEthernet0/1        unassigned      YES DHCP   up                    up
Serial0/1              unassigned      YES unset  administratively down down
```

Let's verify the configuration on the F0/1 interface:

```
R4#Show run int f0/1 | B interface
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
interface FastEthernet0/1
 ip address dhcp
 duplex auto
 speed auto
end
```

The F0/1 interface of R4 is configured correctly, let's verify the configuration of the DHCP server (R5):

On R5:

```
R5#Show run | S dhcp pool

ip dhcp pool tst
   network 45.1.1.0 255.255.255.0

R5#Show run | Inc dhcp

ip dhcp excluded-address 45.1.1.5
ip dhcp pool tst

R5#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/1        45.1.1.5        YES manual up                    up
```

It seems like the configuration was performed correctly. Let's verify the VLAN configuration on SW2:

On SW2:

```
SW2#Show vlan br | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/18, Fa0/20, Fa0/21, Fa0/22
                                                Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2
36   VLAN0036                         active    Fa0/3, Fa0/6
45   VLAN0045                         active    Fa0/4, Fa0/5
68   VLAN0068                         active
212  VLAN0212                         active    Fa0/12
213  VLAN0213                         active    Fa0/13
214  VLAN0214                         active
215  VLAN0215                         active
216  VLAN0216                         active
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
217  VLAN0217                         active    Fa0/17
```

The VLAN is configured properly, the routers are configured properly, could there be a Vlan access-map or an access-list blocking this traffic?
Let's verify:

```
SW2#Show vlan access-map
SW2#

SW2#Show access-list
SW2#
```

Let's check for any configuration that has 45 in it:

```
SW2#Show run | Inc 45

ip dhcp snooping vlan 45
 switchport access vlan 45
 switchport access vlan 45
```

WOW...DHCP Snooping is enabled for VLAN 45, let's verify DHCP Snooping configuration:

```
SW2#Show ip dhcp snooping

Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
45
Insertion of option 82 is enabled
   circuit-id format: vlan-mod-port
    remote-id format: MAC
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled

Interface                  Trusted     Rate limit (pps)
------------------------   -------     ----------------
```

The output of the above command reveals that DHCP snooping is enabled globally, and it's applied to VLAN 45, and option 82 is also enabled. If option 82 is enabled, the host's DHCP Discover message will include the "circuit-id" which is the MAC address of SW2, and the "port-id" which is the port to which the router is connected to, since these added information is NOT understood by the DHCP server, it ignores the discover messages received by R4. Let's disable option 82:

```
SW2(config)#No ip dhcp snooping information option
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's verify this configuration on R4:

On R4:

```
R4#Show ip int br
```

Interface	IP-Address	OK?	Method	Status	Protocol
FastEthernet0/0	unassigned	YES	unset	up	up
Serial0/0	unassigned	YES	unset	administratively down	down
FastEthernet0/1	**unassigned**	**YES**	**DHCP**	**up**	**up**
Serial0/1	unassigned	YES	unset	administratively down	down

Let's enable "Debug ip packet det" on R4:

```
R4#Debug ip packet detail
IP packet debugging is on (detailed)

IP: s=0.0.0.0 (local), d=255.255.255.255 (FastEthernet0/1), len 604,
sending broad/multicast
    UDP src=68, dst=67

IP: s=0.0.0.0 (local), d=255.255.255.255 (FastEthernet0/1), len 604,
sending broad/multicast
    UDP src=68, dst=67

IP: s=0.0.0.0 (local), d=255.255.255.255 (FastEthernet0/1), len 604,
sending broad/multicast
    UDP src=68, dst=67
```

We can see that the local router is sending the DHCP discover messages but it's NOT receiving any offers.

Let's go through what we have verified so far:

R4 is configured as a DHCP client
R5 is configured as a DHCP Server
SW2 has these two routers in VLAN 45
SW2 is configured with "IP DHCP Snooping" and "IP dhcp snooping vlan 45".
Option 82 is also disabled

Since the dhcp snooping option is enabled and applied to VLAN 45, all ports in VLAN 45 will transition into untrusted state. When a port is in untrusted state, the switch will discard the DHCP messages like Offer, ACK, NACK coming from all ports, this includes the port that R5 is connected to. This is designed that way to mitigate rogue DHCP servers. Let's see if this port is trusted:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2#Show run int f0/5 | B interface

interface FastEthernet0/5
 switchport access vlan 45
 switchport mode access
 spanning-tree portfast
end
```

OK.....we can see the problem. Let's trust this port:

```
SW2(config)#int f0/5
SW2(config-if)#ip dhcp snooping trust
```

To verify the configuration:

On R4:

```
R4(config)#int f0/1
R4(config-if)#Shut
R4(config-if)#No shut
```

You should see the following debug output:

```
IP: s=0.0.0.0 (local), d=255.255.255.255 (FastEthernet0/1), len 604,
sending broad/multicast
    UDP src=68, dst=67

IP: s=45.1.1.5 (FastEthernet0/1), d=255.255.255.255, len 328, rcvd 2
    UDP src=67, dst=68

IP: s=0.0.0.0 (local), d=255.255.255.255 (FastEthernet0/1), len 604,
sending broad/multicast
    UDP src=68, dst=67

IP: s=45.1.1.5 (FastEthernet0/1), d=255.255.255.255, len 328, rcvd 2
    UDP src=67, dst=68

%DHCP-6-ADDRESS_ASSIGN: Interface FastEthernet0/1 assigned DHCP address
45.1.1.1, mask 255.255.255.0, hostname R4
```

You may get a different IP address from the same range.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/1        45.1.1.1        YES DHCP   up                    up
```

Let's disable the debug:

```
R4#Undebug all
All possible debugging has been turned off
```

Ticket 3

VLAN 123 is configured with static IP Source Guard but it is not working.

By looking at the diagram, we can see that the F0/0 interfaces of R1, R2 and R3 are configured in VLAN 123. Let's verify:

On SW1:

```
SW1#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/4, Fa0/5, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/18, Fa0/20, Fa0/22
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active
36   VLAN0036                         active
45   VLAN0045                         active
68   VLAN0068                         active    Fa0/6
114  VLAN0114                         active
121  VLAN0121                         active
123  VLAN0123                         active    Fa0/1, Fa0/2, Fa0/3
212  VLAN0212                         active    Fa0/12
213  VLAN0213                         active    Fa0/13
214  VLAN0214                         active    Fa0/14
215  VLAN0215                         active    Fa0/15
216  VLAN0216                         active    Fa0/16
217  VLAN0217                         active    Fa0/17
232  VLAN0232                         active
```

Let's verify the IP Source Guard configuration on SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1#Show run int f0/1 | B interface

interface FastEthernet0/1
 switchport access vlan 123
 switchport mode access
 spanning-tree portfast
 ip verify source
end

SW1#Show run int f0/2 | B interface

interface FastEthernet0/2
 switchport access vlan 123
 switchport mode access
 spanning-tree portfast
 ip verify source
end

SW1#Show run int f0/3 | B interface

interface FastEthernet0/3
 switchport access vlan 123
 switchport mode access
 spanning-tree portfast
 ip verify source
end
```

Let's verify the source binding table, we should see three bindings, one for R1, one for R2, and another for R3 :

On SW1:

```
SW1#Show ip source binding | Inc 123.1.1.1

00:00:22:22:22:22   123.1.1.2       infinite    static      123    FastEthernet0/2
00:00:11:11:11:11   123.1.1.1       infinite    static      123    FastEthernet0/1
00:00:33:33:33:33   123.1.1.3       infinite    static      123    FastEthernet0/3
```

The table looked correct, let's check the configuration of the F0/0 interfaces of R1, R2 and R3:

On R1:

```
R1#Show int f0/0 | Inc bia|address
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Hardware is Gt96k FE, address is 0000.1111.1111 (bia 0016.c898.c6e0)
  Internet address is 123.1.1.1/24
```

On R2:

```
R2#Show int f0/0 | Inc bia|address

  Hardware is Gt96k FE, address is 0000.2222.2222 (bia 0011.9279.1f70)
  Internet address is 123.1.1.2/24
```

On R3:

```
R3#Show int f0/0 | Inc bia|address

  Hardware is Gt96k FE, address is 0000.3333.3333 (bia 0012.d927.d710)
  Internet address is 123.1.1.3/24
```

This is configured correctly. Let's verify DHCP Snooping configuration:

On SW1:

```
SW1#Show ip dhcp snooping

Switch DHCP snooping is disabled
DHCP snooping is configured on following VLANs:
123
Insertion of option 82 is enabled
   circuit-id format: vlan-mod-port
    remote-id format: MAC
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Interface                  Trusted    Rate limit (pps)
------------------------   -------    ----------------
```

The output of the above show command reveals that the DHCP Snooping feature is NOT enabled, let's verify this by looking at the running configuration:

On SW1:

```
SW1#Show run | Inc ip dhcp snooping

ip dhcp snooping vlan 123
```

Yes, we can see that it is applied to VLAN 123, but it is NOT enabled globally, let's enable DHCP Snooping

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

globally:

On SW1:

```
SW1(config)#IP Dhcp snooping
```

To verify the configuration:

On SW1

```
SW1#Show ip dhcp snooping

Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
123
Insertion of option 82 is enabled
   circuit-id format: vlan-mod-port
    remote-id format: MAC
Option 82 on untrusted port is not allowed

Verification of hwaddr field is enabled
Interface                  Trusted     Rate limit (pps)
------------------------   -------     ----------------
```

Perfect....let's test the configuration by changing the IP address of R1 to 123.1.1.10 or any IP address that is not configured in the source binding table and test the configuration:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 123.1.1.10 255.255.255.0

R1#Ping 123.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's change the IP address of R1 back to 123.1.1.1/24 and test again:

On R1

```
R1(config)#int f0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#ip addr 123.1.1.1 255.255.255.0

R1#Ping 123.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

On SW1:

```
SW1#Sh ip source binding static

MacAddress          IpAddress        Lease(sec)  Type        VLAN  Interface
------------------  ---------------  ----------  ----------  ----  ----------
00:00:22:22:22:22   123.1.1.2        infinite    static      123   FastEthernet0/2
00:00:11:11:11:11   123.1.1.1        infinite    static      123   FastEthernet0/1
00:00:33:33:33:33   123.1.1.3        infinite    static      123   FastEthernet0/3

Total number of bindings: 3
```

Ticket 4

The F0/19 interface of SW4 is configured as a trunk link. The user has setup port security on this interface and set the maximum number of MAC addresses to 2, but for some reason this is NOT working.

Let's verify the configuration on SW4:

On SW4:

```
SW4#Show run int f0/19 | B interface

interface FastEthernet0/19
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport port-security maximum 2
 switchport port-security aging time 5
 switchport port-security aging type inactivity
```

We can see that the maximum number of MAC addresses allowed is 2, and the aging type and aging

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

time are also set. But the port security is NOT enabled. Let's verify this information:

```
SW4#Show port-security interface F0/19

Port Security              : Disabled
Port Status                : Secure-down
Violation Mode             : Shutdown
Aging Time                 : 5 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 0000.0000.0000:0
Security Violation Count   : 0
```

Let's enable port-security:

```
SW4(config)#int f0/19
SW4(config-if)#Switchport port-security
```

Once the port-security is enabled, you should see the following console messages:

```
%PM-4-ERR_DISABLE: psecure-violation error detected on Fa0/19, putting
Fa0/19 in err-disable state

%PORT_SECURITY-2-PSECURE_VIOLATION: Security violation occurred, caused
by MAC address 000a.f402.6d93 on port FastEthernet0/19.
```

We can clearly see that port security is working, once the feature was enabled, the interface transitioned into err-disable state. To verify this information:

```
SW4#Show interface f0/19 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/19                       err-disabled 1            auto   auto 10/100BaseTX
```

Ticket 5

In VLAN 12, R1 can NOT ping R2's F0/1 interface

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's verify the problem:

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R1#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  12.1.1.1                -   0000.2121.2121  ARPA   FastEthernet0/1
Internet  12.1.1.2                0   Incomplete      ARPA
Internet  123.1.1.1               -   0000.1111.1111  ARPA   FastEthernet0/0
Internet  123.1.1.2              11   0000.2222.2222  ARPA   FastEthernet0/0
```

We can see that ARP was "incomplete", this can mean that we have a layer 2 problem, before checking layer 2, let's see the configuration of the F0/1 interfaces of R1 and R2:

On R1:

```
R1#Show run int f0/1 | B interface

interface FastEthernet0/1
 mac-address 0000.2121.2121
 ip address 12.1.1.1 255.255.255.0
 duplex auto
 speed auto
end
```

On R2:

```
R2#Show run int f0/1 | B interface

interface FastEthernet0/1
 mac-address 0000.1212.1212
 ip address 12.1.1.2 255.255.255.0
 duplex auto
 speed auto
end
```

Well……the interfaces are configured correctly, let's verify the switch configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2#Show vlan br | Exc unsup

VLAN Name                                Status    Ports
---- ----------------------------------- --------- -------------------------------
1    default                             active    Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                   Fa0/11, Fa0/14, Fa0/15, Fa0/16
                                                   Fa0/18, Fa0/20, Fa0/21, Fa0/22
                                                   Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                            active    Fa0/1, Fa0/2
36   VLAN0036                            active    Fa0/3, Fa0/6
45   VLAN0045                            active    Fa0/4, Fa0/5
68   VLAN0068                            active
114  VLAN0114                            active
121  VLAN0121                            active
123  VLAN0123                            active
212  VLAN0212                            active    Fa0/12
213  VLAN0213                            active    Fa0/13
214  VLAN0214                            active
215  VLAN0215                            active
216  VLAN0216                            active
217  VLAN0217                            active    Fa0/17
232  VLAN0232                            active
```

Let's check the VLAN:

```
SW2#Show vlan id 12

VLAN Name                                Status    Ports
---- ----------------------------------- --------- -------------------------------
12   VLAN0012                            active    Fa0/1, Fa0/2, Fa0/19, Fa0/23

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
12   enet  100012     1500  -      -      -        -    -        0      0

Remote SPAN VLAN
----------------
Disabled

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------
```

This is also correct, let's verify the interface configuration, maybe an access-list is configured to block the communication between the two routers:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show run int f0/1 | B interface

interface FastEthernet0/1
 switchport access vlan 12
 switchport mode access
 spanning-tree portfast

SW2#Show run int f0/2 | B interface

interface FastEthernet0/2
 switchport access vlan 12
 switchport mode access
 spanning-tree portfast
```

Let's verify the MAC address of these two routers on the switch:

```
SW2#Show mac address-table dynamic interface f0/1 vlan 12

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  12    0000.2121.2121    DYNAMIC     Fa0/1
Total Mac Addresses for this criterion: 1

SW2#Show mac address-table dynamic interface f0/2 vlan 12

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
```

The switch does not see R2, this could be that R2 did not see the ping generated by R1, let's ping from R2 and verify the result:

On R2:

```
R2#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2#Show mac address-table dynamic interface f0/2 vlan 12

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
```

WOW…may be something is filtering/blocking the MAC address of R2, let's verify:

```
SW2#Show run | Inc 0000.1212.1212

mac-address-table static 0000.1212.1212 vlan 12 drop
```

We can clearly see the problem, this is called Unicast MAC address filtering. When this feature is enabled, the switch drops packets with the configured MAC address as source or destination.

The switch is configured to drop any traffic sourcing or destined to "0000.1212.1212" MAC address in VLAN 12. Let's remove this filter and try again:

```
SW2(config)#No mac address-table static 0000.1212.1212 vlan 12 drop
```

To verify the configuration:

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

On SW2:

```
SW2#Show mac address-table dynamic interface f0/2 vlan 12

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  12    0000.1212.1212    DYNAMIC     Fa0/2
Total Mac Addresses for this criterion: 1
```

Ticket 6

In VLAN 36, R6 can not ping R3.

Let's verify the problem:

On R3:

```
R3#Ping 36.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 36.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  36.1.1.3               -   0000.3636.3636  ARPA   FastEthernet0/1
Internet  36.1.1.6               0   Incomplete      ARPA
Internet  123.1.1.3              -   0000.3333.3333  ARPA   FastEthernet0/0
```

It looks like the problem is with SW2, let's verify SW2's configuration:

On SW2:

```
SW2#Show vlan br | exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/18, Fa0/20, Fa0/21, Fa0/22
                                                Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2
36   VLAN0036                         active    Fa0/3, Fa0/6
45   VLAN0045                         active    Fa0/4, Fa0/5
68   VLAN0068                         active
114  VLAN0114                         active
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
121  VLAN0121                         active
123  VLAN0123                         active
212  VLAN0212                         active    Fa0/12
213  VLAN0213                         active    Fa0/13
214  VLAN0214                         active
215  VLAN0215                         active
216  VLAN0216                         active
217  VLAN0217                         active    Fa0/17
232  VLAN0232                         active

SW2#Show mac address-table dynamic interface F0/3 vlan 36

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  36    0000.3636.3636    DYNAMIC     Fa0/3

Total Mac Addresses for this criterion: 1

SW2#Show mac address-table dynamic interface F0/6 vlan 36

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
```

The switch does not see the MAC address of R6, let's check the configuration of the F0/3, and F0/6 interfaces on the switch, and then check for a unicast Mac address filtering:

```
SW2#Show run int f0/3 | B inter

interface FastEthernet0/3
 switchport access vlan 36
 switchport mode access
 spanning-tree portfast

SW2#Show run int f0/6 | B inter

interface FastEthernet0/6
 switchport access vlan 36
 switchport mode access
 spanning-tree portfast
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's find out the MAC address of R6:

On R6:

```
R6#Show int f0/1 | inc bia

  Hardware is MV96340 Ethernet, address is 0000.6363.6363 (bia 0017.5aad.52ab)

SW2#Show run | Inc 0000.6363.6363
SW2#
```

What else could be the problem? May be some type of filtering is configured based on the interface? Let's verify:

On SW2

```
SW2#Show run | Inc Fa0/6

monitor session 1 destination interface Fa0/6
```

It looks like SPAN was configured on SW2 and the destination port or the monitor port was configured to be F0/6 interface. When a port is configured as monitor port, the port does not transmit any traffic except that required for the SPAN session and incoming traffic is never learned or forwarded on a destination port.

Let's remove the command and verify the communication:

```
SW2(config)#No monitor session 1 destination interface Fa0/6
```

On R3

```
R3#Ping 36.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 36.1.1.6, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 7

In VLAN 78, R7 can NOT ping R8's F0/0 interface.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's verify the problem:

On R7:

```
R7#Ping 78.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 78.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R7#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  78.1.1.7               -   0000.7777.7777  ARPA   GigabitEthernet0/0
Internet  78.1.1.8               0   Incomplete      ARPA
```

We should not waste time, we should go right to the SW1 and check its configuration:

On SW1:

WOW.....you should see the following console messages on SW1:

```
%SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa0/7, vlan
78.([0000.7777.7777/78.1.1.7/0000.0000.0000/78.1.1.8/02:24:37 UTC Mon Mar
1 1993])

%SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa0/7, vlan
78.([0000.7777.7777/78.1.1.7/0000.0000.0000/78.1.1.8/02:24:37 UTC Mon Mar
1 1993])

%SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa0/7, vlan
78.([0000.7777.7777/78.1.1.7/0000.0000.0000/78.1.1.8/02:24:37 UTC Mon Mar
1 1993])

%SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa0/7, vlan
78.([0000.7777.7777/78.1.1.7/0000.0000.0000/78.1.1.8/02:24:37 UTC Mon Mar
1 1993])

%SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa0/7, vlan
78.([0000.7777.7777/78.1.1.7/0000.0000.0000/78.1.1.8/02:24:37 UTC Mon Mar
1 1993])
```

From the above console messages we can see that we have a problem with DAI (Dynamic Arp Inspection).

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The above console message states the following:
The Switch (SW1) received an invalid ARP request (Req) on Fa0/7 in VLAN 78.
The ARP request was from a host with a MAC address of 0000.7777.7777 and an IP address of 78.1.1.7 for 78.1.1.8 (R8).

This means that R7 is NOT allowed to send ARP request and this is why we are getting this error message. Let's verify the configuration of DAI.

On SW1:

```
SW1#Show run | Inc ip arp

ip arp inspection vlan 78
```

We can see that the "IP ARP inspection" command is NOT referencing an ARP access-list, and obviously "IP DHCP Snooping" is also enabled, or else Dynamic ARP Inspection would not have worked. Let's verify the entries in the DHCP Snooping DB:

```
SW1#Show ip dhcp snooping binding

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
00:00:88:88:88:88   78.1.1.8         9217        dhcp-snooping  78    FastEthernet0/8
Total number of bindings: 1
```

We can clearly see the problem, when the F0/7 interface of SW1 receives the ARP request from R7, it consults the DHCP snooping DB for an entry for R7, and since it does not find one, it drops the packet and sends a console message. Let's add an entry for R7 in the Snooping DB and try pinging again. Let's find out the MAC address of R8:

On R7:

```
R7#Show int g0/0 | Inc bia

Hardware is CN Gigabit Ethernet, address is 0000.7777.7777 (bia 24e9.b3ab.4b20)
```

On SW1:

```
SW1#IP Dhcp snooping bindin 0000.7777.7777 vlan 78 78.1.1.7 int F0/7 exp 1000
```

To verify the configuration:

On R7:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7#Ping 78.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 78.1.1.8, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Ticket 8

Your client is complaining that he can not save the running configuration on SW4 by using "Wr" and everytime he tries to save using the "WR" command he gets the following console error message:
Error reading config file "flash:/"startup-config file open failed (Not enough space)
OR
startup-config file open failed (Not enough space)

Let's verify the problem:

On SW4:

```
SW4#wr

startup-config file open failed (Not enough space)
```

WOW...what can cause that? Let's verify the boot parameters:

```
SW4#Show boot

BOOT path-list:      c3550-ipservicesk9-mz.122-25.SEE4.bin
Config file:         flash:/
Private Config file: flash:/private-config.text
Enable Break:        no
Manual Boot:         no
HELPER path-list:
NVRAM/Config file
      buffer size:   393216
```

We can see the problem, the "Config file" is instructed to be saved in the flash, but a file name has not been specified. Let's fix this problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW4(config)#Boot config-file flash:/config.text
```

You must reload the switch for this change to be implemented:

To verify the solution:

On SW4

```
SW4#reload
```

Let's try to save the running configuration when the switch comes up:

```
SW4>en
SW4#wr
Building configuration...
[OK]
```

Ticket 9

Erase the startup configuration and the "vlan.dat" and reload the devices before proceeding to the next scenario.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

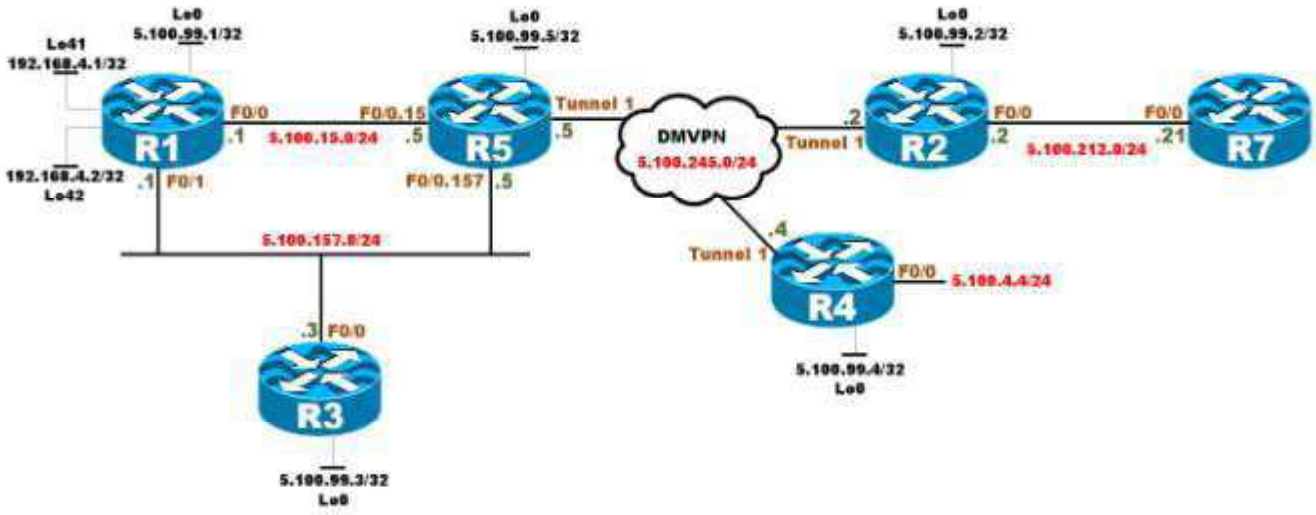
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

RIPv2

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Troubleshooting RIPv2 Scenario - I



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "RIPv2" → "TS-RIPv2-Lab-1".

Ticket 1

R1 is configured to filter its Lo41, however, this interface is still reachable from R5.

Let's verify the problem:

On R5:

```
R5#Ping 192.168.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.1, timeout is 2 seconds:
.!!!!
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Let's verify R5's routing table entry for R1's Lo41:

```
R5#Show ip route 192.168.4.1

Routing entry for 192.168.4.0/24
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 5.100.157.1 on FastEthernet0/0.157, 00:00:03 ago
  Routing Descriptor Blocks:
    5.100.157.1, from 5.100.157.1, 00:00:03 ago, via FastEthernet0/0.157
      Route metric is 1, traffic share count is 1
  * 5.100.15.1, from 5.100.15.1, 00:00:18 ago, via FastEthernet0/0.15
      Route metric is 1, traffic share count is 1
```

We can see that R1's Lo41 is reachable via both F0/0.157 and the F0/0.15 interfaces. The output of the above show command also reveals that R1's Lo41 is reachable through a /24 route.

Shouldn't R1 send a /32 route for its Loopback41?
Let's check RIP's database on R1.

On R1:

```
R1#Show ip rip database

5.0.0.0/8    auto-summary
5.100.15.0/24    directly connected, FastEthernet0/0
5.100.99.1/32    directly connected, Loopback0
5.100.99.5/32
    [1] via 5.100.157.5, 00:00:23, FastEthernet0/1
    [1] via 5.100.15.5, 00:00:19, FastEthernet0/0
5.100.157.0/24    directly connected, FastEthernet0/1
5.100.245.0/24
    [1] via 5.100.157.5, 00:00:23, FastEthernet0/1
    [1] via 5.100.15.5, 00:00:19, FastEthernet0/0
192.168.4.0/24    auto-summary
192.168.4.1/32    directly connected, Loopback41
192.168.4.2/32    directly connected, Loopback42
```

We can see the problem, R1 is auto-summarizing its routes therefore, it summarized the Lo41 and Lo42 IP addresses to a single classful /24 route. Let's verify this in the configuration:

```
R1#Show run | s router rip
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
router rip
 version 2
 network 5.0.0.0
 network 192.168.4.0
 distribute-list toR5 out FastEthernet0/0
```

We can see that the "No auto-summary" command is missing. Let's correct this and verify.

```
R1(config)#Router rip
R1(config-router)#No auto-summary
```

To verify the configuration:

On R1:

```
R1#Show ip rip database

5.0.0.0/8    auto-summary
5.100.15.0/24    directly connected, FastEthernet0/0
5.100.99.1/32    directly connected, Loopback0
5.100.99.5/32
    [1] via 5.100.157.5, 00:00:16, FastEthernet0/1
    [1] via 5.100.15.5, 00:00:16, FastEthernet0/0
5.100.157.0/24    directly connected, FastEthernet0/1
5.100.245.0/24
    [1] via 5.100.157.5, 00:00:16, FastEthernet0/1
    [1] via 5.100.15.5, 00:00:16, FastEthernet0/0
192.168.4.0/24    auto-summary
192.168.4.1/32    directly connected, Loopback41
192.168.4.2/32    directly connected, Loopback42
```

This did not work. The entries are still summarized. Let's check the routing table of R5:

On R5:

```
R5#Show ip route 192.168.4.1

Routing entry for 192.168.4.1/32
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 5.100.157.1 on FastEthernet0/0.157, 00:00:22 ago
  Routing Descriptor Blocks:
  * 5.100.157.1, from 5.100.157.1, 00:00:22 ago, via FastEthernet0/0.157
      Route metric is 1, traffic share count is 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The above show command reveals that R1's Lo41 is ONLY reachable through R5's F0/0.157 interface, can R5 Ping this route? Let's try:

```
R5#Ping 192.168.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Sure can, let's verify the configured filter on R1:

On R1:

```
R1#Show ip protocol
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
    FastEthernet0/0 filtered by toR5 (per-user), default is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 27 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0       2     2
    FastEthernet0/1       2     2                   TST
    Loopback0             2     2
    Loopback41            2     2
    Loopback42            2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    5.0.0.0
    192.168.4.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    Gateway         Distance      Last Update
    5.100.15.5           120      00:00:16
    5.100.157.5          120      00:00:02
  Distance: (default is 120)

R1#Show run | S router rip
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
router rip
 version 2
 network 5.0.0.0
 network 192.168.4.0
 distribute-list toR5 out FastEthernet0/0
 no auto-summary
```

The access-list is ONLY applied to the F0/0 interface, whereas, it should have been applied to both F0/0 and the F0/1 interfaces.

This can be solved in two ways:

Solution #1
Remove the "Distribute-list" command and apply it without referencing the interface, in which case it will be applied to all interfaces.

Solution #2
Add another access-list to the F0/1 interface.

In this case the second method is used:

On R1:

```
R1(config)#Router rip
R1(config-router)#distribute-list toR5 out F0/1
```

To verify the configuration:

On R1 and R5:

```
Rx#Clear ip route *

R5#Ping 192.168.4.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Show ip route rip B Gate
Gateway of last resort is not set

      5.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
R        5.100.99.1/32 [120/1] via 5.100.157.1, 00:00:05, FastEthernet0/0.157
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                     [120/1] via 5.100.15.1, 00:00:06, FastEthernet0/0.15
      192.168.4.0/32 is subnetted, 1 subnets
R        192.168.4.2 [120/1] via 5.100.157.1, 00:00:05, FastEthernet0/0.157
                     [120/1] via 5.100.15.1, 00:00:06, FastEthernet0/0.15
```

Ticket 2

R7 is configued to filter all even third octet networks of 192.168.2xx.0/24, however, this affected the odd and the even routes and none of them are reachable from R2.

Let's verify the problem:

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is not set

      5.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
R        5.100.15.0/24 [120/1] via 5.100.245.5, 00:00:12, Tunnel1
R        5.100.99.1/32 [120/2] via 5.100.245.5, 00:00:12, Tunnel1
R        5.100.99.5/32 [120/1] via 5.100.245.5, 00:00:12, Tunnel1
R        5.100.157.0/24 [120/1] via 5.100.245.5, 00:00:12, Tunnel1
      192.168.4.0/32 is subnetted, 1 subnets
R        192.168.4.2 [120/2] via 5.100.245.5, 00:00:12, Tunnel1
```

None of the routes from R7 are in the routing table of R2. Let's see if R2's F0/0 interface is configured to receive RIPv2 routes:

```
R2#Show ip protocol
*** IP Routing is NSF aware ***

Routing Protocol is "nhrp"
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 250)

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is (prefix-list) plIn
  Incoming update filter list for all interfaces is not set
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Sending updates every 30 seconds, next due in 19 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface              Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0        2     2
    Loopback0              2     2
    Tunnel1                2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    5.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    5.100.245.5           120      00:00:00
  Distance: (default is 120)
```

Yes it is configured correctly. Let's debug ip rip:

```
R2#Debug ip rip
RIP protocol debugging is on

RIP: ignored v1 packet from 5.100.212.21 (illegal version)
```

I think we have seen the problem. Let's disbale debugging:

```
R2#Undebug all
```

On R7:

```
R7(config)#Router rip
R7(config-router)#Version 2

R7#Clear ip route *
```

To verify the configuration:

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is not set

      5.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
R        5.100.15.0/24 [120/1] via 5.100.245.5, 00:00:20, Tunnel1
R        5.100.99.1/32 [120/2] via 5.100.245.5, 00:00:20, Tunnel1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        5.100.99.5/32 [120/1] via 5.100.245.5, 00:00:20, Tunnel1
R        5.100.157.0/24 [120/1] via 5.100.245.5, 00:00:20, Tunnel1
      192.168.4.0/32 is subnetted, 1 subnets
R        192.168.4.2 [120/2] via 5.100.245.5, 00:00:20, Tunnel1
R     192.168.211.0/24 [120/1] via 5.100.212.21, 00:00:02, FastEthernet0/0
R     192.168.213.0/24 [120/1] via 5.100.212.21, 00:00:02, FastEthernet0/0
R     192.168.215.0/24 [120/1] via 5.100.212.21, 00:00:02, FastEthernet0/0
R     192.168.217.0/24 [120/1] via 5.100.212.21, 00:00:02, FastEthernet0/0
R     192.168.219.0/24 [120/1] via 5.100.212.21, 00:00:02, FastEthernet0/0
```

Ticket 3

R4 can't reach R1's Lo42 using its Lo0 as the source.

Let's verify the problem:

On R4:

```
R4#Ping 192.168.4.2 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.2, timeout is 2 seconds:
Packet sent with a source address of 5.100.99.4
.....
Success rate is 0 percent (0/5)
```

Let's see if R4 can reach R1's Lo42 interface using its tunnel interface as the source:

```
R4#Ping 192.168.4.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/52 ms
```

Yes it can. Let's verify R4's routing table:

```
R4#Show ip route 192.168.4.2

Routing entry for 192.168.4.2/32
  Known via "rip", distance 120, metric 2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Redistributing via rip
  Last update from 5.100.245.5 on Tunnel1, 00:00:09 ago
  Routing Descriptor Blocks:
  * 5.100.245.5, from 5.100.245.5, 00:00:09 ago, via Tunnel1
      Route metric is 2, traffic share count is 1
```

Well, the route is in R4's routing table. There are two basic possiblities: The packets are filtered, or R1 can't reach R4's Lo0. Let's verify:

On R1:

```
R1#Show ip route 5.100.99.4
% Subnet not in table
```

We can se that R1 does NOT have a route back to R4's Lo0. Let's move closer to R4 by verifying the routing table of R5:

On R5:

```
R5#Show ip route 5.100.99.4
% Subnet not in table
```

R5 doesn't have a route either. Let's check and see if R5 receives a route for R4's Lo0:

```
R5#Debug ip rip
RIP protocol debugging is on

RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0 (5.100.15.5)
RIP: build update entries
        5.100.99.5/32 via 0.0.0.0, metric 1, tag 0
        5.100.157.0/24 via 0.0.0.0, metric 1, tag 0
        5.100.245.0/24 via 0.0.0.0, metric 1, tag 0

RIP: received packet with MD5 authentication
RIP: received v2 update from 5.100.157.1 on FastEthernet0/1
     5.100.15.0/24 via 0.0.0.0 in 1 hops
     5.100.99.1/32 via 0.0.0.0 in 1 hops
     192.168.4.2/32 via 0.0.0.0 in 1 hops
RIP: received v2 update from 5.100.15.1 on FastEthernet0/0
     5.100.99.1/32 via 0.0.0.0 in 1 hops
     5.100.157.0/24 via 0.0.0.0 in 1 hops
     192.168.4.2/32 via 0.0.0.0 in 1 hops

RIP: received packet with MD5 authentication
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
RIP: ignored v2 packet from 5.100.157.3 (invalid authentication)
R5#
RIP: sending v2 update to 224.0.0.9 via Serial0/0/0 (5.100.245.5)
RIP: build update entries
        5.100.15.0/24 via 0.0.0.0, metric 1, tag 0
        5.100.99.1/32 via 0.0.0.0, metric 2, tag 0
        5.100.99.5/32 via 0.0.0.0, metric 1, tag 0
        5.100.157.0/24 via 0.0.0.0, metric 1, tag 0
        5.100.245.0/24 via 0.0.0.0, metric 1, tag 0
        192.168.4.2/32 via 0.0.0.0, metric 2, tag 0
RIP: sending v2 update to 224.0.0.9 via FastEthernet0/1 (5.100.157.5)
RIP: build update entries
        5.100.15.0/24 via 0.0.0.0, metric 1, tag 0
        5.100.99.5/32 via 0.0.0.0, metric 1, tag 0
        5.100.245.0/24 via 0.0.0.0, metric 1, tag 0
```

We don't see any routing update for that route. Let's disable the debugging:

```
R5#Undebug all
All possible debugging has been turned off
```

Let's verify R4's RIP configuration:

On R4:

```
R4#Show run | S router rip

router rip
 version 2
 passive-interface default
 network 5.0.0.0
 no auto-summary
```

R4 is configured NOT to send any updates out of any of its interfaces and that's the reason R1 is not getting the route. Let's fix this problem:

```
R4(config)#Router rip
R4(config-router)#No passive-interface Tunnel 1

R4#Ping 192.168.4.2 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.2, timeout is 2 seconds:
Packet sent with a source address of 5.100.99.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

Ticket 4

R2 is configured to filter its F0/0 interface with an IP address of 5.100.212.2/24, but R5 can't reach R2's Lo0.

Let's verify the problem:

On R5:

```
R5#Ping 5.100.99.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check and see if R5 has an entry for R2's Lo0 in its routing table:

On R5:

```
R5#Show ip route 5.100.99.2
% Subnet not in table
```

R5 does not have a route for that Loopback interface in its routing table, let's verify R2's configuration:

On R2:

```
R2#Show ip protocol
*** IP Routing is NSF aware ***

Routing Protocol is "nhrp"
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 250)

Routing Protocol is "rip"
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Outgoing update filter list for all interfaces is (prefix-list) plIn
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 28 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0       2     2
    Loopback0             2     2
    Tunnel1               2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    5.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    5.100.245.5           120      00:00:01
    5.100.212.21          120      00:00:17
  Distance: (default is 120)
```

Let's check the prefix-list:

```
R2#Show ip prefix-list

ip prefix-list plIn: 2 entries
   seq 5 deny 5.100.212.0/24
   seq 10 permit 0.0.0.0/0
```

The prefix-list plIn is filtering outgoing RIP updates. Can you see the problem in the prefix-list?

The prefix-list denies R2's F0/0 interface and allows 0.0.0.0/0 which is a default route instead of a catch-all entry (Permit any) which should have been configued as 0.0.0.0/0 Le 32. Let's correct this problem:

```
R2(config)#No ip prefix-list PlIn seq 10 permit 0.0.0.0/0
R2(config)#ip prefix-list PlIn seq 10 permit 0.0.0.0/0 LE 32
```

Let's verify and see if R5 can reach R2's Lo0:

On R5:

```
R5#Show ip route 5.100.99.2

Routing entry for 5.100.99.2/32
  Known via "rip", distance 120, metric 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Redistributing via rip
  Last update from 5.100.245.2 on Tunnel1, 00:00:06 ago
  Routing Descriptor Blocks:
  * 5.100.245.2, from 5.100.245.2, 00:00:06 ago, via Tunnel1
      Route metric is 1, traffic share count is 1

R5#Ping 5.100.99.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/54/56 ms
```

Ticket 5

R4 can't reach R2's Lo0 interface.

Let's verify the problem:

On R4:

```
R4#Ping 5.100.99.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the routing table of R4 for this prefix:

```
R4#Show ip route 5.100.99.2

% Subnet not in table
```

Does R4 have reachability to R2's Tunnel IP address?

```
R4#Ping 5.100.245.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.245.2, timeout is 2 seconds:
.....
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

Let's check the tunnel configuration of R4 and R2:

```
R4#Show run int tunnel 1 | B inter

interface Tunnel1
 ip address 5.100.245.4 255.255.255.0
 no ip redirects
 ip nhrp map multicast 192.1.5.5
 ip nhrp map 5.100.245.5 192.1.5.5
 ip nhrp network-id 444
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

This looks ok. Let's check R2's tunnel interface:

On R2:

```
R2#Show run int tunnel 1 | B inter

interface Tunnel1
 ip address 5.100.245.2 255.255.255.0
 no ip redirects
 ip nhrp map multicast 192.1.5.5
 ip nhrp map 5.100.245.5 192.1.5.5
 ip nhrp network-id 222
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

This also looks good, let's enable debugging while pinging R2's tunnel IP address from R4:

On R4:

```
R4#Debug condition interface tunnel 1
Condition 1 set

R4#Debug ip packet
IP packet debugging is on

IP: s=5.100.245.4 (local), d=5.100.245.2 (Tunnel1), len 100, sending
IP: s=5.100.245.4 (local), d=5.100.245.2 (Tunnel1), len 100,
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
encapsulation failed.
```

We can see the problem, R4 does NOT have a mapping for R2's Tunnel IP address and this is why it is failing with "encapsulation failure". Let's disable the debug and correct the problem and verify:

```
R4#Undebug all
All possible debugging has been turned off

R4(config)#Int tunnel 1
R4(config-if)#ip nhrp map 5.100.245.2 192.1.2.2
```

On R2:

```
R2(config)#int tunnel 1
R2(config-if)#ip nhrp map 5.100.245.4 192.1.4.4
```

On R4:

```
R4#Ping 5.100.99.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Now that R2 and R4 have NHRP mapping, we should be able to summarize 5.100.99.0/24 on R5 and send a summary route to both R2 and R4:

```
R5(config)#Int tunnel 1
R5(config-if)#IP summary-address rip 5.100.99.0 255.255.255.0
```

To verify the configuration:

On R4:

```
R4#Ping 5.100.99.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Remember the mapping is checked before routing.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 6

R3 can't ping R1's Lo0.

Let's verify the configuration:

On R3:

```
R3#Ping 5.100.99.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check R3's routing table:

```
R3#Show ip route 5.100.99.1
% Subnet not in table
```

Let's enable Debug ip rip:

```
R3#Debug ip rip
RIP protocol debugging is on

RIP: received packet with MD5 authentication
RIP: ignored v2 packet from 5.100.157.1 (invalid authentication)
```

The problem is obvious. Let's disable the deugging

```
R3#Undebug all
All possible debugging has been turned off
```

The only way to see this is to compare the configuration on all routers involved:

On R1:

```
R1#Show run | s key chain

key chain TST
 key 1
   key-string ccie

R1#Show run int F0/1 | B interface
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
interface FastEthernet0/1
 ip address 5.100.157.1 255.255.255.0
 ip rip authentication mode md5
 ip rip authentication key-chain TST
 duplex auto
 speed auto

R1#Show run | Inc _ccie$

   key-string ccie
```

The last command is entered to check for spaces after the password.

On R3:

```
R3#Show run | s key chain

key chain TST
 key 1
   key-string ccie

R3#Show run int F0/0 | B interface

interface FastEthernet0/0
 ip address 5.100.157.3 255.255.255.0
 ip rip authentication mode md5
 ip rip authentication key-chain TST
 duplex auto
 speed auto
end

R3#Show run | Inc _ccie$
R3#
```

OK, we can see the problem, let's correct this problem:

```
R3(config)#key chain TST
R3(config-keychain)#key 1
R3(config-keychain-key)#key-string ccie
```

To verify the problem:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show run | Inc _ccie$
   key-string ccie
```

To test the configuration:

On R3:

```
R3#Ping 5.100.99.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.100.99.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Ticket 7

R3 is configured to filter all routes received from R5, however, the routes are still there.
DO NOT use another method to fix this problem, correct the existing problem.

Let's verify the problem:

On R3:

```
R3#Show ip route rip | B Gate
Gateway of last resort is not set

      5.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
R        5.100.4.0/24 [120/2] via 5.100.157.5, 00:00:07, FastEthernet0/0
R        5.100.15.0/24 [120/1] via 5.100.157.5, 00:00:07, FastEthernet0/0
                       [120/1] via 5.100.157.1, 00:00:27, FastEthernet0/0
R        5.100.99.1/32 [120/1] via 5.100.157.1, 00:00:27, FastEthernet0/0
R        5.100.99.2/32 [120/2] via 5.100.157.5, 00:00:07, FastEthernet0/0
R        5.100.99.4/32 [120/2] via 5.100.157.5, 00:00:07, FastEthernet0/0
R        5.100.99.5/32 [120/1] via 5.100.157.5, 00:00:07, FastEthernet0/0
R        5.100.245.0/24 [120/1] via 5.100.157.5, 00:00:07,
FastEthernet0/0
      192.168.4.0/32 is subnetted, 1 subnets
R        192.168.4.2 [120/1] via 5.100.157.1, 00:00:27, FastEthernet0/0
R     192.168.211.0/24 [120/3] via 5.100.157.5, 00:00:07, FastEthernet0/0
R     192.168.213.0/24 [120/3] via 5.100.157.5, 00:00:07, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R     192.168.215.0/24 [120/3] via 5.100.157.5, 00:00:07, FastEthernet0/0
R     192.168.217.0/24 [120/3] via 5.100.157.5, 00:00:07, FastEthernet0/0
R     192.168.219.0/24 [120/3] via 5.100.157.5, 00:00:07, FastEthernet0/0
```

Sure enough we can see bunch of routes from R5, let's see how R3 was configured to accomplish this filtering:

```
R3#Show ip protocol

*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is 157
  Sending updates every 30 seconds, next due in 27 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0       2     2                   KC157
    Loopback0             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    5.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    5.100.157.5           120      00:00:24
    5.100.157.1           120      00:00:14
  Distance: (default is 120)
```

Let's verify the access-list:

```
R3#Show access-list 157

Extended IP access list 157

    10 deny ip any host 5.100.157.5
    20 permit ip any any (182 matches)
```

We can see a brave attempt to filter those routes using an extended ACL, but the ACL was misconfigured. When using an extended ACL for filtering, the "SOURCE" should be configured to reference the advertising router, and the "DESTINATION" should be configured to reference the actual route/s.

In R3's configured ACL, the advertising router states "ANY" and the actual route is configured to be

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

5.100.157.5, these need to be swapped.
Let's configure R3 to correct this mistake:

On R3

```
R3(config)#No access-list 157
R3(config)#Access-list 157 deny ip host 5.100.157.5 any
R3(config)#Access-list 157 permit ip any any

R3#Clear ip route *
```

To verify the configuration:

On R3:

```
R3#Show ip route | B Gate
Gateway of last resort is not set

      5.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R        5.100.15.0/24 [120/1] via 5.100.157.1, 00:00:00, FastEthernet0/0
R        5.100.99.1/32 [120/1] via 5.100.157.1, 00:00:00, FastEthernet0/0
      192.168.4.0/32 is subnetted, 1 subnets
R        192.168.4.2 [120/1] via 5.100.157.1, 00:00:00, FastEthernet0/0
```

Much better.

Ticket 8

Erase the startup configuration and reload the devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

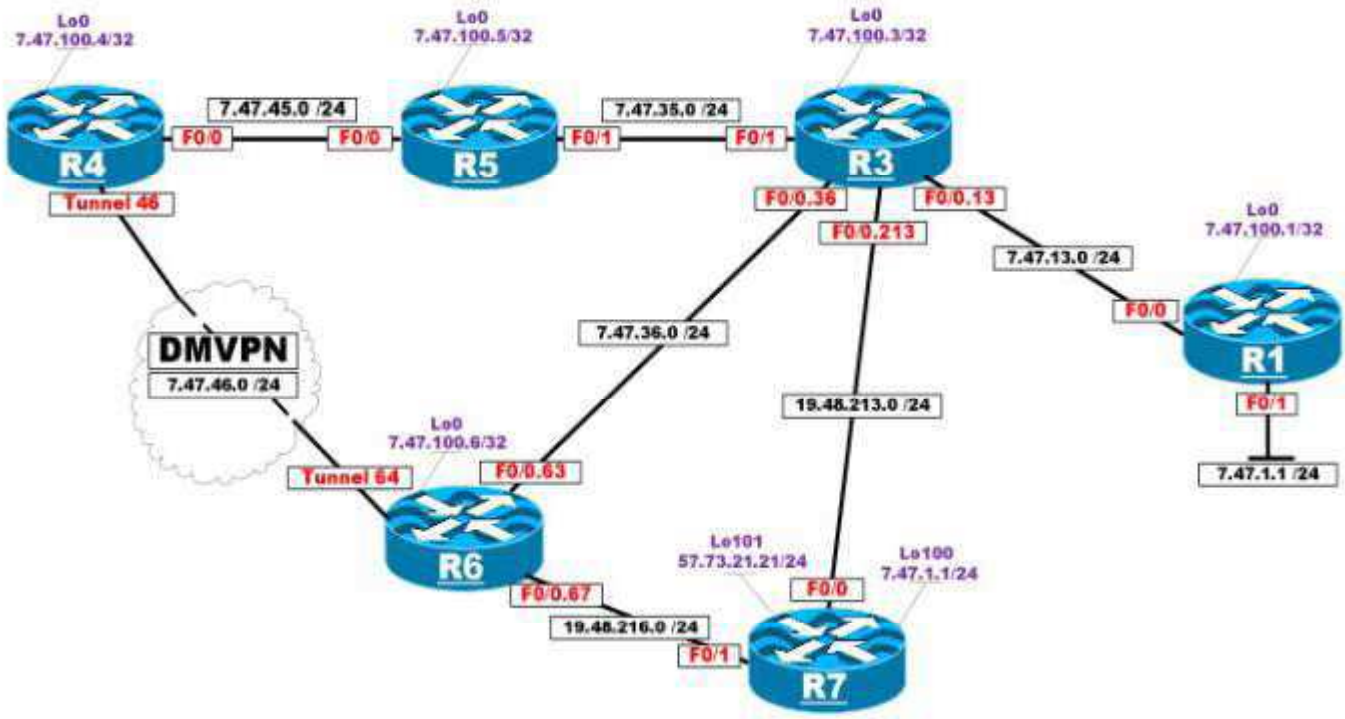
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

EIGRP

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Troubleshooting EIGRP Scenario - I



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "EIGRP" → "TS-EIGRP-TS-Lab-1".

NOTE: DO NOT access R7 at all.

Ticket 1

R1 can't reach R3's Lo0. You must configure R1 to fix the problem.

Let's verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Ping 7.47.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see if we have a neighbor adjacency with R3:

On R1:

```
R1#Show ip eigrp neighbors
IP-EIGRP neighbors for process 10
```

No adjacencies at all. Are we running EIGRP on R3's F0/0.13 sub-interface?

On R3:

```
R3#Show ip eigrp interfaces
IP-EIGRP interfaces for process 1
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/0.13	0	0/0	0	0/1	0	0
Fa0/1	1	0/0	2	0/1	50	0
Fa0/0.36	0	0/0	0	0/1	0	0
Lo0	0	0/0	0	0/1	0	0
Fa0/0.213	1	0/0	1	0/1	50	0

EIGRP is running on the F0/0.13 sub-interface of R3. Is there a connectivity between R1 and R3? Let's verify:

On R1:

```
R1#Ping 7.47.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.13.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/6/24 ms
```

We can see the connectivity is there. Why aren't they establishing an adjacency? Let's enable EIGRP debugging and verify:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Debug ip eigrp
IP-EIGRP Route Events debugging is on
```

No output at all. Let's enable IP packet debugging, this should NOT be used in a production network, but since we are in a small lab it should be fine.

```
R1#undebug all
All possible debugging has been turned off

R1#Debug ip packet
IP packet debugging is on

IP: s=7.47.100.1 (local), d=224.0.0.10 (Loopback0), len 60, sending
broad/multicast
IP: s=7.47.100.1 (Loopback0), d=224.0.0.10, len 60, rcvd 2
```

We can see that R1 is Multicasting EIGRP packets, let's disable the debug and look at the EIGRP configuration of these two routers:

On R1:

```
R1#undebug all
All possible debugging has been turned off

R1#Show run | S router eigrp

router eigrp 10
 network 7.0.0.0
```

On R3:

```
R3#Show run | S router eigrp

router eigrp 1
 variance 2
 network 7.47.13.0 0.0.0.255
 network 7.47.35.0 0.0.0.255
 network 7.47.36.0 0.0.0.255
 network 7.47.100.3 0.0.0.0
 network 19.0.0.0
```

OK, we can see the problem, R1 is running in EIGRP AS 1, whereas, R3 is running in EIGRP AS 10, let's fix this problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#No router eigrp 10

R1(config)#Router eigrp 1
R1(config-router)#network 7.0.0.0
```

We should immediately see the following comnsole messags:

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 7.47.13.3 (FastEthernet0/0) is
up: new adjacency
```

To test the configuration:

On R1:

```
R1#Ping 7.47.100.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 2

R6 does NOT have a stable EIGRP adjacency to R4. DO NOT use an EIGRP command to fix this ticket.

Let's verify the problem:

On R4:

```
R4#Show ip eigrp neighbors
IP-EIGRP neighbors for process 1

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   7.47.45.5               Fa0/0             11 00:29:35    1   200  0  23
```

Let's verify the debug output on R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#Access-list 100 permit ip host 7.47.46.6 host 224.0.0.10

R4#Debug ip packet 100
IP packet debugging is on for access list 100
```

Nothing....let's check their configuration:

```
R4#Show run | S router eigrp

router eigrp 1
 network 7.0.0.0
 metric maximum-hops 2
```

On R6:

We should see the following console messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 7.47.36.3 (FastEthernet0/0.63)
is up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 19.48.216.21
(FastEthernet0/0.67) is up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 7.47.46.4 (Tunnel64) is down:
retry limit exceeded

%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 7.47.46.4 (Tunnel64) is up: new
adjacency
```

We can see that the djacency is flapping.

```
R6#Show run | S router eigrp

router eigrp 1
 network 7.0.0.0
 network 19.0.0.0
```

The configuration look fine, let's check the configuration of the tunnel interfaces:

On R4:

```
R4#Show run int tunnel 46 | B interface

interface Tunnel46
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 ip address 7.47.46.4 255.255.255.0
 no ip redirects
 ip nhrp map 7.47.46.6 192.1.6.6
 ip nhrp map multicast 192.1.6.6
 ip nhrp network-id 444
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

On R6:

```
R6#Show run int tunnel 64 | B interface

interface Tunnel64
 ip address 7.47.46.6 255.255.255.0
 no ip redirects
 ip nhrp map 7.47.46.4 192.1.4.4
 ip nhrp network-id 666
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

OK, we can see the problem, R6 is not configured for Multicast capabilities. The task states that we should not fix the problem using an EIGRP command. This problem can be fixed by enabling Multicast capabilities on R6, let configure this and verify:

```
R6(config)#int tunnel 64
R6(config-if)#ip nhrp map multicast 192.1.4.4
```

To verify the configuration:

On R6:

```
R6#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1)

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   7.47.46.4               Tu64              11 00:00:50    6  1434  0  11
2   19.48.216.21            Fa0/0.67          12 00:04:29    1   200  0  20
1   7.47.36.3               Fa0/0.63          11 00:04:51    1   200  0  23

R6#Show ip rou eigrp | Inc 7.47.46.4

D        7.47.100.4/32 [90/27008000] via 7.47.46.4, 00:01:35, Tunnel64
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 3

When R3's F0/0 is down, R3 can't reach R4's Lo0.

Let's verify the configuration:

On R3:

```
R3#Ping 7.47.100.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

Let's shutdown the F0/0 interface of R3 and test again:

```
R3(config)#Int F0/0
R3(config-if)#Shut

R3#Ping 7.47.100.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check R3's routing table:

```
R3#Show ip route eigrp

     7.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D       7.47.45.0/24 [90/30720] via 7.47.35.5, 00:06:15, FastEthernet0/1
D       7.47.100.5/32 [90/156160] via 7.47.35.5, 00:06:15, FastEthernet0/1
```

We can ONLY see R5's Loopback 0 and R5's connection to R4. Let's check R5's routing table may be it does not have a route to R4's Lo0 interface:

On R5:

```
R5#Show ip route 7.47.100.4

Routing entry for 7.47.100.4/32
  Known via "eigrp 1", distance 90, metric 156160, type internal
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Redistributing via eigrp 1
  Last update from 7.47.45.4 on FastEthernet0/0, 00:04:51 ago
  Routing Descriptor Blocks:
  * 7.47.45.4, from 7.47.45.4, 00:04:51 ago, via FastEthernet0/0
      Route metric is 156160, traffic share count is 1
      Total delay is 5100 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

R5 does have a route for R4's Lo0, let see if R5 is configured with a filter:

```
R5#Show ip protocol
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(1)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 7.47.100.5
    Stub, connected, summary
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    7.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    7.47.45.4             90      00:02:18
    7.47.35.3             90      00:02:18
  Distance: internal 90 external 170
```

R5 is NOT configured with any filtering but we can see that it is configued with "EIGRP Stub", and this is why it does NOT pass R4's Lo0's network. Let's remove that configuration and verify the ticket:

```
R5#Show run | S router eigrp
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
router eigrp 1
 network 7.0.0.0
 no auto-summary
 eigrp stub connected summary
```

```
R5(config)#Router eigrp 1
R5(config-router)#No eigrp stub connected summary
```

To test the configuration:

On R3:

```
R3#Show ip route eigrp | B Gate
Gateway of last resort is not set

      7.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D        7.47.1.0/24 [90/26887936] via 7.47.35.5, 00:00:18, FastEthernet0/1
D        7.47.36.0/24 [90/26887680] via 7.47.35.5, 00:00:18, FastEthernet0/1
D        7.47.45.0/24 [90/30720] via 7.47.35.5, 00:00:20, FastEthernet0/1
D        7.47.46.0/24 [90/26885120] via 7.47.35.5, 00:00:18, FastEthernet0/1
D        7.47.100.4/32 [90/158720] via 7.47.35.5, 00:00:18, FastEthernet0/1
D        7.47.100.5/32 [90/156160] via 7.47.35.5, 00:00:20, FastEthernet0/1
D        7.47.100.6/32 [90/27013120] via 7.47.35.5, 00:00:18, FastEthernet0/1
      19.0.0.0/24 is subnetted, 2 subnets
D        19.48.213.0 [90/26890240] via 7.47.35.5, 00:00:18, FastEthernet0/1
D        19.48.216.0 [90/26887680] via 7.47.35.5, 00:00:18, FastEthernet0/1

R3#Ping 7.47.100.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R3(config)#Int F0/0
R3(config-if)#No shut
```

You should see the adjacency before proceeding to the next ticket, you should see the following console messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 7.47.36.6 (FastEthernet0/0.36)
is up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 7.47.13.1 (FastEthernet0/0.13)
is up: new adjacency
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 19.48.213.21
(FastEthernet0/0.213) is up: new adjacency
```

Ticket 4

R1's Lo0 should always have connectivity to R4's Lo0 and F0/0 interfaces, however, this is NOT the case. You should fix this problem without configuring R1 or R4. YOU should NOT remove any commands to resolve this ticket.

Let's verify the ticket by sourcing the Pings from the F0/0 interface of R1 and then the Lo0 interface of R1, if the ping is unsuccessful, please "No shut" the F0/0 interface of R3 and wait till t establishes an adjacency before performing the following ping.

On R1:

```
R1#Ping 7.47.100.4 source f0/0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
Packet sent with a source address of 7.47.13.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 7.47.100.4 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
Packet sent with a source address of 7.47.100.1
.....
Success rate is 0 percent (0/5)
```

May be R4 does not have a route back to R1's Lo0, let's verify:

On R4:

```
R4#Show ip route 7.47.100.1
% Subnet not in table
```

Does R5 have a route to R1's Lo0? Let's verify:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R5:

```
R5#Show ip route 7.47.100.1

Routing entry for 7.47.100.1/32
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via eigrp 1
  Last update from 7.47.35.3 on FastEthernet0/1, 00:07:33 ago
  Routing Descriptor Blocks:
  * 7.47.35.3, from 7.47.35.3, 00:07:33 ago, via FastEthernet0/1
      Route metric is 158720, traffic share count is 1
      Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

R5#Ping 7.47.100.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Is R4 filtering?

On R4:

```
R4#Show ip protocol
*** IP Routing is NSF aware ***

Routing Protocol is "nhrp"
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 250)

Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(1)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 7.47.100.4
    Topology : 0 (base)
      Active Timer: 3 min
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    Distance: internal 90 external 170
    Maximum path: 4
    Maximum hopcount 2
    Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    7.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    7.47.46.6             90      02:05:03
    7.47.45.5             90      02:05:02
  Distance: internal 90 external 170
```

R4 is NOT filtering using the traditional way but the EIGRPs maximum hop is set to 2, well, this is one kind of filtering. But the ticket states that we can NOT configure R1 or R4 to resolve this ticket, what else can we do? Summarization is the answer.

On R5:

```
R5(config)#Int F0/0
R5(config-if)#IP summary-address eigrp 1 7.47.0.0 255.255.0.0
```

Let's test the configuration:

On R1:

```
R1#Ping 7.47.100.4 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
Packet sent with a source address of 7.47.100.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 7.47.100.4 source f0/0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.47.100.4, timeout is 2 seconds:
Packet sent with a source address of 7.47.13.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 24/26/28 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 5

R3 is configured to use both paths to R4's Lo0, however, it's ONLY using one of the paths.

Let's verify the problem:

On R3:

```
R3#Show ip route 7.47.100.4

Routing entry for 7.47.100.4/32
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via eigrp 1
  Last update from 7.47.35.5 on FastEthernet0/1, 00:19:29 ago
  Routing Descriptor Blocks:
  * 7.47.35.5, from 7.47.35.5, 00:19:29 ago, via FastEthernet0/1
      Route metric is 158720, traffic share count is 1
      Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

We can ONLY see the path through R5, let's check R3's EIGRP topology table:

```
R3#Show ip eigrp topology 7.47.100.4/32

IP-EIGRP (AS 1): Topology entry for 7.47.100.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 158720
  Routing Descriptor Blocks:
  7.47.35.5 (FastEthernet0/1), from 7.47.35.5, Send flag is 0x0
      Composite metric is (158720/156160), Route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5200 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
```

We don't see R6 as an alternate route. Let's check R6's routing table:

On R6:

```
R6#Show ip route 7.47.100.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 7.47.100.4/32
  Known via "eigrp 1", distance 90, metric 161280, type internal
  Redistributing via eigrp 1
  Last update from 7.47.36.3 on FastEthernet0/0.63, 02:10:53 ago
  Routing Descriptor Blocks:
  * 7.47.36.3, from 7.47.36.3, 02:10:53 ago, via FastEthernet0/0.63
      Route metric is 161280, traffic share count is 1
      Total delay is 5300 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 3
```

R6 is using the path through R3 and NOT from its tunnel 64 interface. Let's see why R6 prefers the route throught R3, let's check R6's EIGRP's topology table:

```
R6#Show ip eigrp topology 7.47.100.4/32

EIGRP-IPv4 Topology Entry for AS(1)/ID(7.47.100.6) for 7.47.100.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 161280
  Descriptor Blocks:
  7.47.36.3 (FastEthernet0/0.63), from 7.47.36.3, Send flag is 0x0
      Composite metric is (161280/158720), route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5300 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 3
        Originating router is 7.47.100.4
  19.48.216.21 (FastEthernet0/0.67), from 19.48.216.21, Send flag is 0x0
      Composite metric is (163840/161280), route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5400 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 4
  7.47.46.4 (Tunnel64), from 7.47.46.4, Send flag is 0x0
      Composite metric is (2297856/128256), route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25000 microseconds
        Reliability is 255/255
        Load is 1/255
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
        Minimum MTU is 1476
        Hop count is 1
        Originating router is 7.47.100.4
```

Now we can see why, the delay and minimum bandwidth are better on the path through R3. Let's resolve the problem by changing R6's tunnel64 interface's bandwidth and delay; a higher value should be configured for the R6's tunnel 64 interface.

By default, EIGRP calculates the composite metric by adding the delay to the bandwidth and then multiplying it by 256.

The delay is the sum of all the delays along the path to the Loopback0 interface of R4; R4 is advertising a delay of 5000 for its Lo0 interface; the delay of R6's tunnel 64 interface is 20,000, therefore, delay through the tunnel is 25,000, but the delay through R3 is 5300, so the delay of the tunnel64 interface of R6 should be set to 300, this will make the total delay through the tunnel and then to the Lo0 interface of R4 to 5300. Remember that delay value is displayed in tens of microseconds. To make this path a little more attractive, the bandwidth is set to 100,001.

```
R6(config)#Int tunnel 64
R6(config-if)#bandwidth 100001
R6(config-if)#delay 30
```

To verify the configuration:

On R6:

```
R6#Show ip eigrp topology 7.47.100.4/32

EIGRP-IPv4 Topology Entry for AS(1)/ID(7.47.100.6) for 7.47.100.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 161024
  Descriptor Blocks:
  7.47.46.4 (Tunnel64), from 7.47.46.4, Send flag is 0x0
      Composite metric is (161024/128256), route is Internal
      Vector metric:
        Minimum bandwidth is 101010 Kbit
        Total delay is 5300 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1476
        Hop count is 1
        Originating router is 7.47.100.4
  7.47.36.3 (FastEthernet0/0.63), from 7.47.36.3, Send flag is 0x0
      Composite metric is (161280/158720), route is Internal
      Vector metric:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
        Minimum bandwidth is 100000 Kbit
        Total delay is 5300 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 3
        Originating router is 7.47.100.4
  19.48.216.21 (FastEthernet0/0.67), from 19.48.216.21, Send flag is 0x0
      Composite metric is (163840/161280), route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5400 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 4
```

Let's verify the configuration:

On R6:

```
R6#Show ip route 7.47.100.4

Routing entry for 7.47.100.4/32
  Known via "eigrp 1", distance 90, metric 161024, type internal
  Redistributing via eigrp 1
  Last update from 7.47.46.4 on Tunnel64, 00:02:33 ago
  Routing Descriptor Blocks:
  * 7.47.46.4, from 7.47.46.4, 00:02:33 ago, via Tunnel64
      Route metric is 161024, traffic share count is 1
      Total delay is 5300 microseconds, minimum bandwidth is 101010 Kbit
      Reliability 255/255, minimum MTU 1476 bytes
      Loading 1/255, Hops 1
```

Let's test the configuration:

On R3:

```
R3#Show ip eigrp topology 7.47.100.4/32

EIGRP-IPv4 Topology Entry for AS(1)/ID(7.47.100.3) for 7.47.100.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 158720
  Descriptor Blocks:
  7.47.35.5 (FastEthernet0/1), from 7.47.35.5, Send flag is 0x0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      Composite metric is (158720/156160), route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5200 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 7.47.100.4
  7.47.36.6 (FastEthernet0/0.36), from 7.47.36.6, Send flag is 0x0
      Composite metric is (163840/161024), route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5400 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1476
        Hop count is 2
        Originating router is 7.47.100.4

R3#Show ip route eigrp | B Gate
Gateway of last resort is not set

      7.0.0.0/8 is variably subnetted, 15 subnets, 3 masks
D        7.47.0.0/16 [90/40960] via 7.47.36.6, 00:05:11, FastEthernet0/0.36
D        7.47.1.0/24
           [90/28416] via 19.48.213.21, 00:18:02, FastEthernet0/0.213
           [90/30720] via 7.47.13.1, 00:18:02, FastEthernet0/0.13
D        7.47.45.0/24 [90/30720] via 7.47.35.5, 02:45:26, FastEthernet0/1
D        7.47.46.0/24 [90/35840] via 7.47.36.6, 00:05:21, FastEthernet0/0.36
D        7.47.100.1/32 [90/156160] via 7.47.13.1, 02:45:32, FastEthernet0/0.13
D        7.47.100.4/32 [90/158720] via 7.47.35.5, 00:05:11, FastEthernet0/1
D        7.47.100.5/32 [90/156160] via 7.47.35.5, 02:45:26, FastEthernet0/1
D        7.47.100.6/32 [90/156160] via 7.47.36.6, 00:18:02, FastEthernet0/0.36
      19.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D        19.48.216.0/24
           [90/30720] via 19.48.213.21, 00:18:04, FastEthernet0/0.213
           [90/30720] via 7.47.36.6, 00:18:04, FastEthernet0/0.36
      57.0.0.0/16 is subnetted, 1 subnets
D EX     57.73.0.0
           [170/2560002816] via 19.48.213.21, 02:45:31, FastEthernet0/0.213
```

This is a little better, but R3 is still using a single route, let's check R3's protocol configuration:

```
R3#Show ip protocol
*** IP Routing is NSF aware ***
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(1)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 7.47.100.3
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 2

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    7.47.13.0/24
    7.47.35.0/24
    7.47.36.0/24
    7.47.100.3/32
    19.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    7.47.36.6             90      00:07:14
    7.47.35.5             90      00:07:15
    7.47.13.1             90      00:20:05
    19.48.213.21          90      00:07:14
  Distance: internal 90 external 170
```

The variance is set to two, which should be more than enough, let's have another look at R3's EIGRP topology table.

```
R3#Show ip eigrp topology 7.47.100.4/32 | I Composite
      Composite metric is (158720/156160), Route is Internal
      Composite metric is (163840/161024), Route is Internal
```

The metric of the best route **Reported distance**

Now we can see the problem, the reported metric from R6 to R3 is higher than the calculated metric for the best route. this means that feasibility condition was not met. The delay of R6's tunnel64 interface should be decreased by a value such that the reported metric from R6 is lower than R3's best path.

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6(config)#int tunnel 64
R6(config-if)#delay 20
```

To verify the configuration:

On R3:

```
R3#Show ip eigrp topology 7.47.100.4/32 | I Composite

      Composite metric is (158720/156160), Route is Internal
      Composite metric is (161280/158464), Route is Internal
```

This looks much better, let's check R3's routing table:

```
R3#Show ip route 7.47.100.4

Routing entry for 7.47.100.4/32
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via eigrp 1
  Last update from 7.47.36.6 on FastEthernet0/0.36, 00:00:24 ago
  Routing Descriptor Blocks:
    7.47.36.6, from 7.47.36.6, 00:00:24 ago, via FastEthernet0/0.36
      Route metric is 161280, traffic share count is 59
      Total delay is 5300 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1476 bytes
      Loading 1/255, Hops 2
  * 7.47.35.5, from 7.47.35.5, 00:00:24 ago, via FastEthernet0/1
      Route metric is 158720, traffic share count is 60
      Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2

R3#Show ip route eigrp | B Gate
Gateway of last resort is not set

      7.0.0.0/8 is variably subnetted, 15 subnets, 3 masks
D        7.47.0.0/16 [90/38400] via 7.47.36.6, 00:01:37, FastEthernet0/0.36
D        7.47.1.0/24
           [90/28416] via 19.48.213.21, 00:24:57, FastEthernet0/0.213
           [90/30720] via 7.47.13.1, 00:24:57, FastEthernet0/0.13
D        7.47.45.0/24 [90/30720] via 7.47.35.5, 02:52:21, FastEthernet0/1
D        7.47.46.0/24 [90/33280] via 7.47.36.6, 00:01:47, FastEthernet0/0.36
D        7.47.100.1/32 [90/156160] via 7.47.13.1, 02:52:27, FastEthernet0/0.13
D        7.47.100.4/32 [90/161280] via 7.47.36.6, 00:01:37, FastEthernet0/0.36
                       [90/158720] via 7.47.35.5, 00:01:37, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D        7.47.100.5/32 [90/156160] via 7.47.35.5, 02:52:21, FastEthernet0/1
D        7.47.100.6/32 [90/156160] via 7.47.36.6, 00:24:57, FastEthernet0/0.36
      19.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D        19.48.216.0/24
           [90/30720] via 19.48.213.21, 00:24:59, FastEthernet0/0.213
           [90/30720] via 7.47.36.6, 00:24:59, FastEthernet0/0.36
      57.0.0.0/16 is subnetted, 1 subnets
D EX     57.73.0.0
           [170/2560002816] via 19.48.213.21, 02:52:26, FastEthernet0/0.213
```

Ticket 6

R6 can't Telnet to R1's F0/1 interface. Fix this problem in the shortest possible way.

Let's verify the problem:

On R6:

```
R6#Telnet 7.47.1.1
Trying 7.47.1.1 ...
% Connection refused by remote host
```

let's see the routing table of R6 for 7.47.1.1:

```
R6#Show ip route 7.47.1.1

Routing entry for 7.47.1.0/24
  Known via "eigrp 1", distance 90, metric 28416, type internal
  Redistributing via eigrp 1
  Last update from 19.48.216.21 on FastEthernet0/0.67, 00:26:46 ago
  Routing Descriptor Blocks:
  * 19.48.216.21, from 19.48.216.21, 00:26:46 ago, via FastEthernet0/0.67
      Route metric is 28416, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

R6 is going through R7, but why is it NOT going through R3? Let's check R3's routing table:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip route 7.47.1.1

Routing entry for 7.47.1.0/24
  Known via "eigrp 1", distance 90, metric 28416, type internal
  Redistributing via eigrp 1
  Last update from 7.47.13.1 on FastEthernet0/0.13, 04:27:14 ago

  Routing Descriptor Blocks:
  * 19.48.213.21, from 19.48.213.21, 04:27:14 ago, via FastEthernet0/0.213
      Route metric is 28416, traffic share count is 40
      Total delay is 110 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
    7.47.13.1, from 7.47.13.1, 04:27:14 ago, via FastEthernet0/0.13
      Route metric is 30720, traffic share count is 37
      Total delay is 200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

R3 has two routes, it received one from R1 and the other from R7, which means that both routers (R1 and R7) are advertising the same route.

So R7 is also advertising the same route, what solution are we going to use?
An Offset-list command can be used to add a metric to all routes advertised by R7. Let's configure this option and test:

Only configuring R6 may not work, because it depends on how R3 load share the TCP connection from R6 to R1.

On R6:

```
R6(config)#Router eigrp 1
R6(config-router)#offset-list 0 in 1000000 F0/0.67
```

On R3:

```
R3(config)#Router eigrp 1
R3(config-router)#offset-list 0 in 1000000 F0/0.213
```

To verify the configuration:

On R6:

```
R6#Show ip route 7.47.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 7.47.1.0/24
  Known via "eigrp 1", distance 90, metric 33280, type internal
  Redistributing via eigrp 1
  Last update from 7.47.36.3 on FastEthernet0/0.63, 00:00:16 ago
  Routing Descriptor Blocks:
  * 7.47.36.3, from 7.47.36.3, 00:00:16 ago, via FastEthernet0/0.63
      Route metric is 33280, traffic share count is 1
      Total delay is 300 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

To test the configuration:

On R6:

```
R6#Telnet 7.47.1.1
Trying 7.47.1.1 ... Open

R1>
R1>quit

[Connection to 7.47.1.1 closed by foreign host]
```

Ticket 7

R6 can't reach R7's Lo101.

Let's verify the problem:

On R6:

```
R6#Ping 57.73.21.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 57.73.21.21, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check R6's routing table:

```
R6#Show ip route 57.73.21.21
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
% Network not in table
```

Let's see if R3 can Ping 57.73.21.21:

On R3:

```
R3#Ping 57.73.21.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 57.73.21.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Let's check R6 and see if it's configured to filter this route:

On R6:

```
R6#Show ip protocol
*** IP Routing is NSF aware ***

Routing Protocol is "nhrp"
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 250)

Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(1)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 7.47.100.6
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    7.0.0.0
    19.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    7.47.46.4             90      00:03:11
    7.47.36.3             90      00:02:36
    19.48.216.21          90      00:02:36
  Distance: internal 90 external 170
```

Let's check R3's routing table:

On R3:

```
R3#Show ip route 57.73.21.21

Routing entry for 57.73.0.0/16
  Known via "eigrp 1", distance 170, metric 2561002816, type external
  Redistributing via eigrp 1
  Last update from 19.48.213.21 on FastEthernet0/0.213, 00:26:00 ago
  Routing Descriptor Blocks:
  * 19.48.213.21, from 19.48.213.21, 00:26:00 ago, via FastEthernet0/0.213
      Route metric is 2561002816, traffic share count is 1
      Total delay is 39172 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

Let's check the topology table:

```
R3#Show ip eigrp topology 57.73.0.0/16

EIGRP-IPv4 Topology Entry for AS(1)/ID(7.47.100.3) for 57.73.0.0/16
  State is Passive, Query origin flag is 1, 0 Successor(s), FD is
4294967295
  Descriptor Blocks:
  19.48.213.21 (FastEthernet0/0.213), from 19.48.213.21, Send flag is 0x0
      Composite metric is (2561002816/2561000256), route is External
      Vector metric:
        Minimum bandwidth is 1 Kbit
        Total delay is 39172 microseconds
        Reliability is 1/255
        Load is 1/255
        Minimum MTU is 1
        Hop count is 1
      External data:
        Originating router is 7.47.100.6
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
        AS number of route is 0
        External protocol is Static, external metric is 0
        Administrator tag is 0 (0x00000000)
```

We can see that the route is coming from R7. R7 advertises a route for 75.73.0.0/16 and we can also see that it's an external route, which means that R7 is redistributing the 57.73.21.21 IP address. We also see that the originating router is R6.

Well, we can guess the problem, R7's router-id is configured to match R6's router-id, and since the route is external, R6 is ignoring the update for that route.

Let's change R6's router-id:

On R6:

```
R6(config)#Router eigrp 1
R6(config-router)#Eigrp router-id 6.6.6.6
```

To test the configuration:

On R6:

```
R6#Ping 57.73.21.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 57.73.21.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 8

Erase the startup configuration and reload the devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

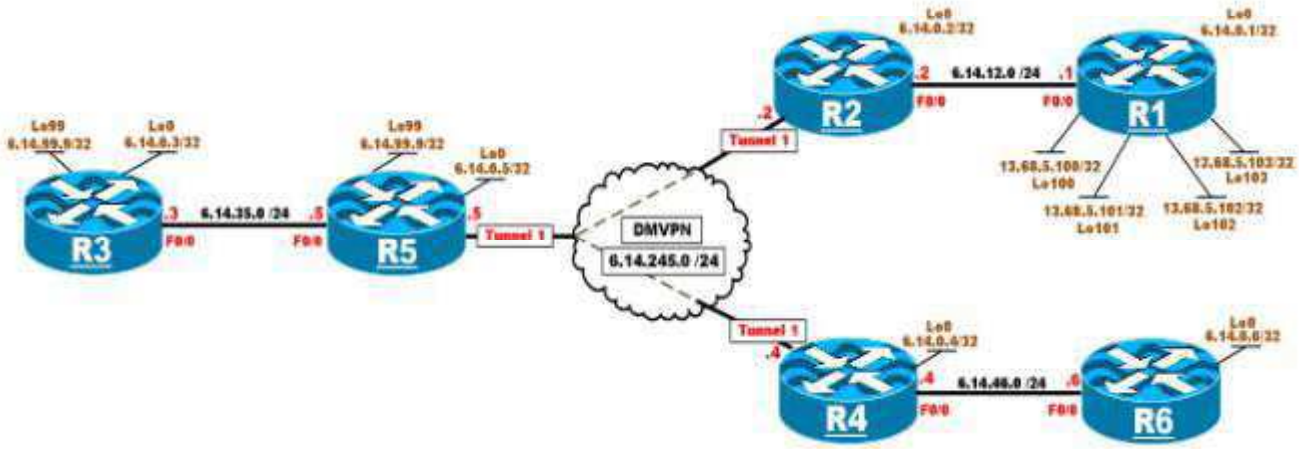
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

OSPF

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Troubleshooting OSPF Scenario - I



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "OSPF" → "TS-OSPF-initial-Lab-1".

Ticket 1

R5 can't Ping R3's Lo0.

Let's verify the problem:

On R5:

```
R5#Ping 6.14.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Is the route in the routing table?

```
R5#Show ip route 6.14.0.3
% Subnet not in table
```

Do we have a neighbor adjacency with R3?

```
R5#Show ip ospf neighbor
R5#
```

Can we ping R3?

```
R5#Ping 6.14.35.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.35.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Let's see if OSPF is running on the F0/0 interface of R5:

```
R5#Show ip ospf int brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0                6.14.0.5/32       1     LOOP  0/0
Tu1          1     0                6.14.245.5/24     64    DR    0/0
Fa0/0        1     35               6.14.35.5/24      1     DR    0/0
Tu56         1     56               6.14.56.5/24      50    DOWN  0/0
```

Let's see how logging is configured, because we don't see any console messages:

```
R5#Show logging

Syslog logging: enabled (0 messages dropped, 3 messages rate-limited,
                0 flushes, 0 overruns, xml disabled, filtering disabled)

No Active Message Discriminator.
No Inactive Message Discriminator.

    Console logging: level emergencies, 29 messages logged, xml disabled,
                     filtering disabled
    Monitor logging: level debugging, 0 messages logged, xml disabled,
                     filtering disabled
    Buffer logging:  disabled, xml disabled,
                     filtering disabled
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    Exception Logging: size (4096 bytes)
    Count and timestamp logging messages: disabled
    Persistent logging: disabled

No active filter modules.

    Trap logging: level informational, 68 message lines logged
        Logging Source-Interface:       VRF Name:
```

The console is configured to log on emergencies, let's reconfigure console logging for all messages:

```
R5(config)#Logging console debugging
```

You should see the following console message within 30 seconds:

```
%OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 6.14.99.9 from
6.14.35.3 on interface FastEthernet0/0
```

OK we can see the problem, R3 and R5 have identical router IDs, let's fix this problem by configuring the router id of these two routers based on the IP address of their Loopback0 interface:

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 6.14.0.3

R3#Clear ip ospf process
Reset ALL OSPF processes? [no]: Y
```

On R5:

```
R5(config)#Router ospf 1
R5(config-router)#router-id 6.14.0.5

R5#Clear ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

You should see the following console message on R5:

```
%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.3 on FastEthernet0/0 from FULL to
DOWN, Neighbor Down: Interface down or detached

%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.3 on FastEthernet0/0 from LOADING
to FULL, Loading Done
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R5:

```
R5#Ping 6.14.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 2

R2 can't reach R5's Lo0.

Let's verify the problem:

On R2:

```
R2#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Is the route in the routing table?

```
R2#Show ip route 6.14.0.5
% Subnet not in table
```

Do we have a neighbor adjacency with R5?

```
R2#Show ip ospf neighbor
R2#
```

Let's see if OSPF is running on the tunnel interface of R2:

```
R2#Show ip ospf inter brief
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Interface     PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0           1     0               6.14.0.2/32        1     LOOP  0/0
Tu1           1     0               6.14.245.2/24      64    P2P   0/0
Fa0/0         1     12              6.14.12.2/24       1     DR    0/0
```

Can R2 reach R5?

```
R2#Ping 6.14.245.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.245.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms
```

Let's enable Conditional debuging on R2's tunnel interface for OSPF hellos:

```
R2#Debug condition interface tunnel 1
Condition 1 set

R2#Debug ip ospf hello
OSPF hello events debugging is on

OSPF-1 HELLO Tu1: Send hello to 224.0.0.5 area 0 from 6.14.245.2
OSPF-1 HELLO Tu1: Send hello to 224.0.0.5 area 0 from 6.14.245.2
```

We are not seeing any output, that means R2 can NOT send OSPF hello messages to R5, even though these routers have reachability to each other. Let's check the network type of the tunnel interfaces of R2 and R5:

```
R2#Show ip ospf int tunne 1 | Inc Network

  Internet Address 6.14.245.2/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 6.14.0.2, Network Type NON_BROADCAST, Cost: 64
```

On R5:

```
R5#Show ip ospf int tunne 1 | Inc Network

  Internet Address 6.14.245.5/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 6.14.0.5, Network Type BROADCAST, Cost: 64
```

We can see that there is a network type mismatch, let's change R2's network type to broadcast:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#Int tunnel 1
R2(config-if)#ip ospf network broadcast
```

Once the network is changed, you should see that R2 is sending hello messages but not receiving any:

```
OSPF-1 HELLO Tu1: Send hello to 224.0.0.5 area 0 from 6.14.245.2
```

Let's see if they have established an OSPF neighbor adjacency:

```
R2#Show ip ospf neighbor
R2#
```

Let's enable conditional debugging for OSPF on R5's tunnel interface as well and see if it's sending or receiving hello messages:

On R5:

```
R5#Debug condition int tunnel 1
Condition 1 set

R5#Debug ip ospf hello
OSPF hello events debugging is on

OSPF: Send hello to 224.0.0.5 area 0 on Tunnel 1 from 6.14.245.5
```

The above output reveals that R5 is also sending the hello messages but not receiving any hello messags from R2, let's check the DMVPN configuration of these two routers:

```
R5#Show run int tunnel 1 | B interface

interface Tunnel1
 bandwidth 1544
 ip address 6.14.245.5 255.255.255.0
 no ip redirects
 ip nhrp map multicast 192.1.4.4
 ip nhrp map 6.14.245.4 192.1.4.4
 ip nhrp map 6.14.245.2 192.1.2.2
 ip nhrp network-id 555
 ip ospf network broadcast
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

We can see the problem, the NHRP map Multicast for R2's connection is missing and this is why the hello

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

messags are not getting to R2, let's fix this problem:

On R5:

```
R5(config)#Int Tunnel 1
R5(config-if)#ip nhrp map multicast 192.1.2.2
```

The output of the following debug shows that once the above configuration was performed the neighbor adjacency came up but a little later it went back down. Let's check the DMVPN configuration of R2:

```
%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.2 on Tunnel1 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.2 on Tunnel1 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

On R2:

Once we console into R2 we see that R2 is receiving the hello messages from R5, but we know that R5 was not receving hello messages from R2.

```
R2#Show run int tunnel 1 | B interface

interface Tunnel1
 bandwidth 1544
 ip address 6.14.245.2 255.255.255.0
 no ip redirects
 ip nhrp map 6.14.245.5 192.1.5.5
 ip nhrp network-id 222
 ip ospf network broadcast
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

Now we can see the full picture, the reason R5 was not receiving any hello messages from R2 was because R2's NHRP map Multicast statement was missing, let's fix this problem:

```
R2(config)#Int tunnel 1
R2(config-if)#ip nhrp map multicast 192.1.5.5
```

Let's verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms
```

Let's disable the conditional debugging on both routers:

On R2 and R5:

```
RX#Undebug all
All possible debugging has been turned off
```

On R2:

```
R2#Undebug condition interface tunnel 1

This condition is the last interface condition set.
Removing all conditions may cause a flood of debugging
messages to result, unless specific debugging flags
are first removed.

Proceed with removal? [yes/no]: Y
Condition 1 has been removed
```

On R5:

```
R5#Undebug condition inter tunnel 1

This condition is the last interface condition set.
Removing all conditions may cause a flood of debugging
messages to result, unless specific debugging flags
are first removed.

Proceed with removal? [yes/no]: Y
Condition 1 has been removed
```

Ticket 3

R1 can't ping R5's Lo0. You MUST use an OSPF command to fix this problem.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's verify the ticket:

On R1:

```
R1#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the routing table:

```
R1#Show ip route 6.14.0.5
% Subnet not in table
```

Let's check the neighbor table:

```
R1#Show ip ospf neighbor
R1#
```

Are we running OSPF on R1's F0/0 interface?

```
R1#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Fa0/0        1     12              6.14.12.1/23       1     DR    0/0
Lo0          1     12              6.14.0.1/32        1     LOOP  0/0
```

Let's debug OSPF hello messages:

```
R1#Debug ip ospf hello
OSPF hello events debugging is on

OSPF-1 HELLO Fa0/0: Rcv hello from 6.14.0.2 area 12 6.14.12.2
OSPF-1 HELLO Fa0/0: Mismatched hello parameters from 6.14.12.2
OSPF-1 HELLO Fa0/0: Dead R 40 C 40, Hello R 10 C 10 Mask R 255.255.255.0
C 255.255.254.0
```

We can see the problem, the output of the above debug message states that the received mask is 255.255.255.0, whereas, the mask for the F0/0 interface on the local router is set to 255.255.254.0. Let's disable the debug and verify this information with the "**Show run interface F0/0**" command on both routers:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show run int f0/0 | Inc ip address

 ip address 6.14.12.1 255.255.254.0
```

On R2:

```
R2#Show run int F0/0 | Inc ip address

 ip address 6.14.12.2 255.255.255.0
```

But the ticket states that an OSPF command must be used to fix this problem. In OSPF if the network type is changed to point-to-point, the mask is ignored. Let's try this:

On R1 and R2:

```
RX(config)#Int F0/0
RX(config-if)#ip ospf network point-to-point
```

You should see the following console messge:

```
%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.2 on FastEthernet0/0 from LOADING
to FULL, Loading Done
```

Sure enough the second the command was entered the neighbor adjacency was established.

To verify the configuration:

On R1:

```
R1#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 52/53/56 ms
```

Ticket 4

R4 is configured to filter R3's Lo0 from reaching R6. R6 should have reachability to all Loopback interfaces on all routers in this topology except R3's Lo0. However, this is NOT the case.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's verify the ticket:

On R6:

```
R6#Ping 6.14.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Well, R6 can not ping the lo0 interface of R3 which is good, but let's see if R6 can ping the loopback 0 interfaces of the other routers starting with R1's:

```
R6#Ping 6.14.0.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 6.14.0.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 6.14.0.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R6#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Only R4's Lo0 is reachable. Let's check the routing table of R6:

```
R6#Show ip route ospf | B Gate
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Gateway of last resort is not set

      6.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O IA     6.14.0.4/32 [110/2] via 6.14.46.4, 00:36:17, FastEthernet0/0
O IA     6.14.245.0/24 [110/65] via 6.14.46.4, 00:36:17, FastEthernet0/0
```

R6 can see R4's Lo0 but not the other ones, let's move closer and go to R4 and check R4's routing table:

On R4:

```
R4#Show ip route ospf | B Gate
Gateway of last resort is not set

      6.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O        6.14.0.6/32 [110/2] via 6.14.46.6, 00:37:11, FastEthernet0/0
```

Let's enable debugging for OSPF hello messages on R4 and see why this router does not have the loopback interfaces in its routing table:

```
R4#Debug ip ospf hello
OSPF hello events debugging is on

OSPF-1 HELLO Tu1: Rcv hello from 6.14.0.5 area 0 6.14.245.5
OSPF-1 HELLO Tu1: Mismatched hello parameters from 6.14.245.5
OSPF-1 HELLO Tu1: Dead R 40 C 120, Hello R 10 C 30 Mask R 255.255.255.0 C
255.255.255.0
```

We can see the problem, the dead interval of received hello messages is set to 40, but the locally configured one is set to 120, and there is also a mismatch in the hello intervals as well, let's disable the debugging and verify the problem by using the "**Show ip ospf inter tunnel 1**" command:

```
R4#Undebug all
All possible debugging has been turned off

R4#Sh ip ospf int tunne 1 | Inc Network|Timer

  Process ID 1, Router ID 6.14.0.4, Network Type NON_BROADCAST, Cost: 64
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
```

On R5:

```
R5#Sh ip ospf int tunne 1 | Inc Network|Timer

  Internet Address 6.14.245.5/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 6.14.0.5, Network Type BROADCAST, Cost: 64
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

Now we can see the problem, there is a network type mismatch which results into a mismatch on hellos and dead interval. Since the network type on R4 is "Non-broadcast", R4 should have a neighbor command to R5, let's verify R4's Tunnel configuration:

On R4:

```
R4#Show run int tunnel 1 | B interface

interface Tunnel1
 bandwidth 1544
 ip address 6.14.245.4 255.255.255.0
 no ip redirects
 ip nhrp map 6.14.245.4 192.1.5.5
 ip nhrp network-id 444
 ip ospf network non-broadcast
 ip ospf 1 area 0
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

Let's change the network type which in turn will fix the hello and dead interval between R4 and R5, but remember that the NHRP map Multicast command must be configured on R5.

```
R4(config)#Int Tunnel 1
R4(config-if)#ip ospf network broadcast
R4(config-if)#ip nhrp map multicast 192.1.5.5
```

You should see the following console message on R4:

```
%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.5 on Tunnel1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R6:

```
R6#Ping 6.14.0.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Still not working, let's traceroute to R1's Lo0 interface:

```
R6#Traceroute 6.14.0.1 Numeric

Type escape sequence to abort.
Tracing the route to 6.14.0.1

  1 6.14.46.4 4 msec 0 msec 0 msec
  2  *  *
```

The packet gets to R4 and then nothing, let's ping from the Lo0 interface of R1 from R4:

On R4:

```
R4#Ping 6.14.0.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the routing table of R4 for this IP address:

```
R4#Show ip route 6.14.0.1

Routing entry for 6.14.0.1/32
  Known via "ospf 1", distance 110, metric 66, type inter area
  Last update from 6.14.245.2 on Tunnel1, 00:01:54 ago
  Routing Descriptor Blocks:
  * 6.14.245.2, from 6.14.0.2, 00:01:54 ago, via Tunnel1
      Route metric is 66, traffic share count is 1
```

The next hop is R2, let's check and see if we have an NHRP map to 6.14.245.2:

On R4:

```
R4#Show run int tunnel 1 | B inter
interface Tunnel1
 bandwidth 1544
 ip address 6.14.245.4 255.255.255.0
 no ip redirects
 ip nhrp map 6.14.245.5 192.1.5.5
 ip nhrp map multicast 192.1.5.5
 ip nhrp network-id 444
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 ip ospf network broadcast
 ip ospf 1 area 0
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

There is no NHRP map to R2's Tunnel IP address, in order for this to work, both R2 and R4 must have NHRP maps to each other's IP address:

```
R4(config)#int Tunnel 1
R4(config-if)#IP nhrp map 6.14.245.2 192.1.2.2
```

Let's check R2's Tunnel configuration:

On R2:

```
R2#Show run int tunnel 1 | B inter

interface Tunnel1
 bandwidth 1544
 ip address 6.14.245.2 255.255.255.0
 no ip redirects
 ip nhrp map multicast 192.1.5.5
 ip nhrp map 6.14.245.5 192.1.5.5
 ip nhrp network-id 222
 ip ospf network broadcast
 tunnel source FastEthernet0/1
 tunnel mode gre multipoint
end
```

R2 does not have an NHRP map back to R4, let's configure one:

On R2

```
R2(config)#int Tunnel 1
R2(config-if)#IP nhrp map 6.14.245.4 192.1.4.4
```

To verify the configuration:

On R6:

```
R6#Ping 6.14.0.1

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 6.14.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 104/104/104 ms

R6#Ping 6.14.0.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/103/104 ms

R6#Ping 6.14.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/52 ms

R6#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/52 ms
```

This is good but R4 was suppose to filter R3's Lo0 from reaching R6. Let's check R6's routing table for R3's Lo0 interface:

```
R6#Show ip route 6.14.0.3

Routing entry for 6.14.0.3/32
  Known via "ospf 1", distance 110, metric 67, type inter area
  Last update from 6.14.46.4 on FastEthernet0/0, 00:00:14 ago
  Routing Descriptor Blocks:
  * 6.14.46.4, from 6.14.0.4, 00:00:14 ago, via FastEthernet0/0
      Route metric is 67, traffic share count is 1
```

There is a route in R6's routing table for the Lo0 interface of R3, let's check R4's OSPF configuration and see how R4 is filtering this route.

<u>On R4:</u>

```
R4#Show run | s router ospf

router ospf 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
area 0 range 6.14.0.3 255.255.255.255 not-advertise
```

We can see the problem, the "Area range" command is for filtering intra-area routes and NOT inter-area routes, let's re-configure this filter the correct way:

On R4:

```
R4(config)#ip prefix-list NET deny 6.14.0.3/32
R4(config)#ip prefix-list NET permit 0.0.0.0/0 le 32

R4(config)#router ospf 1
R4(config-router)#area 0 filter-list prefix NET out
```

To verify the configuration:

On R6:

```
R6#Ping 6.14.0.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/102/104 ms

R6#Ping 6.14.0.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 100/102/104 ms

R6#Ping 6.14.0.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R6#Ping 6.14.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/52 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Ping 6.14.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Ticket 5

OSPF on Tunnel 56 is NOT stable. Please test the previous ticket and make sure that it is working.

We can verify the ticket by looking at all the messages on R5 and R6, let's see if we can find out the problem. On R5 it shows how the tunnel is going up and down:

On R5:

```
%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.6 on Tunnel56 from DOWN to DOWN,
Neighbor Down: Ignore timer expired

%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.6 on Tunnel56 from EXSTART to DOWN,
Neighbor Down: Too many retransmissions

%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.6 on Tunnel56 from DOWN to DOWN,
Neighbor Down: Ignore timer expired
```

On R6:

We can see the same on R6. Let's check the tunnel configuration on R6:

```
R6#Show run int tunnel 56 | B inter

interface Tunnel56
 bandwidth 2000
 ip address 6.14.56.6 255.255.255.0
 ip ospf 1 area 56
 tunnel source Loopback0
 tunnel destination 6.14.0.5
```

The above output reveals that the tunnel destination is R5's Lo0. R6 can learn this loopback interface through R4 or from R6 via the tunnel.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's view OSPF's database for this IP address:

```
R6#Show ip ospf database summary 6.14.0.5

            OSPF Router with ID (6.14.0.6) (Process ID 1)

                Summary Net Link States (Area 46)

  LS age: 1184
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 6.14.0.5 (summary Network Number)
  Advertising Router: 6.14.0.4
  LS Seq Number: 80000006
  Checksum: 0x24A1
  Length: 28
  Network Mask: /32
        MTID: 0         Metric: 65

                Summary Net Link States (Area 56)

  Routing Bit Set on this LSA in topology Base with MTID 0
  LS age: 109
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 6.14.0.5 (summary Network Number)
  Advertising Router: 6.14.0.5
  LS Seq Number: 80000008
  Checksum: 0x976B
  Length: 28
  Network Mask: /32
        MTID: 0         Metric: 1
```

Why is the cost 65? Let's go to R4 and verify its cost to R5:

```
R4#Show ip ospf int tunnel 1 | Inc Cost

  Process ID 1, Router ID 6.14.0.4, Network Type BROADCAST, Cost: 64
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

Let's see R6's cost by going through Tunnel 56:

```
R6#Show ip ospf inter tunne 56 | Inc Cost
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Process ID 1, Router ID 6.14.0.6, Network Type POINT_TO_POINT, Cost: 1
Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

We can see that there are two possible routes to get to the tunnel destination, one through R4 with a cost of 65 and the second one from R5 with a metric of 1.
Why is the metric of the tunnel 56 which is going through R4 lower than R4's metric?

Now we can see why R6 preferred the tunnel interface over the path via R4. Basically R6 is going through the tunnel to reach the tunnel's destination.

How can we fix this with one command?

On R6:

```
R6(config)#Int tunnel 56
R6(config-if)#ip ospf cost 66
```

You should see the following console message:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel56, changed state
to up

%OSPF-5-ADJCHG: Process 1, Nbr 6.14.0.5 on Tunnel56 from LOADING to FULL,
Loading Done
```

```
R6#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
6.14.0.4          1   FULL/DR         00:00:34    6.14.46.4       FastEthernet0/0
6.14.0.5          0   FULL/  -        00:00:35    6.14.56.5       Tunnel56
```

```
R6#Show ip ospf inter brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     46              6.14.0.6/32        1     LOOP  0/0
Fa0/0        1     46              6.14.46.6/24       1     BDR   1/1
Tu56         1     56              6.14.56.6/24       66    P2P   1/1
```

Even though we fixed this ticket, we broke the previous one, let's fix this problem, let's check:

On R6:

```
R6#Ping 6.14.0.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/62/64 ms
```

To fix this problem:

On R5:

```
R5(config)#ip prefix-list NET deny 6.14.0.3/32
R5(config)#ip prefix-list NET permit 0.0.0.0/0 le 32

R5(config)#router ospf 1
R5(config-router)#area 56 filter-list prefix NET in
```

To verify the configuration:

On R6:

```
R6#Ping 6.14.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.14.0.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Ticket 6

R5 is configured to inject a default route if R3's Lo0 is reachable. However, R5 is NOT injecting a default route.

Let's verify the problem:

On R3:

```
R3#Show ip route 0.0.0.0 0.0.0.0

% Network not in table
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check R5's OSPF database:

On R5:

```
R5#Show ip ospf da external 0.0.0.0

            OSPF Router with ID (6.14.0.5) (Process ID 1)

R5#Show ip ospf da external self-originate

            OSPF Router with ID (6.14.0.5) (Process ID 1)
```

Let's see if R5 has R3's Lo0 in its routing table:

```
R5#Show ip route 6.14.0.3

Routing entry for 6.14.0.3/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 6.14.35.3 on FastEthernet0/0, 04:04:20 ago
  Routing Descriptor Blocks:
  * 6.14.35.3, from 6.14.0.3, 04:04:20 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```

Let's check R5's configuration:

```
R5#Show run | S router ospf

router ospf 1
 router-id 6.14.0.5
 log-adjacency-changes
 area 56 filter-list prefix NET in
 network 6.14.0.0 0.0.0.0 area 0
 network 6.14.35.0 0.0.0.255 area 35
 network 6.14.245.0 0.0.0.255 area 0
 default-information originate route-map rmDefault
```

Let's verify the route-map:

```
R5#Show route-map rmDefault

route-map rmDefault, deny, sequence 10
  Match clauses:
    ip address (access-lists): 5
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R5 was configured to reference the "rmDefault" route-map for conditional default route injection, but the route-map was configured as deny. Let's fix this problem:

```
R5(config)#No route-map rmDefault
R5(config)#route-map rmDefault permit 10
R5(config-route-map)#match ip addr 5
```

To verify the configuration:

On R5:

```
R5#Show ip ospf database external 0.0.0.0

            OSPF Router with ID (6.14.0.5) (Process ID 1)

                Type-5 AS External Link States

  LS age: 41
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 0.0.0.0 (External Network Number )
  Advertising Router: 6.14.0.5
  LS Seq Number: 80000001
  Checksum: 0x7623
  Length: 36
  Network Mask: /0
        Metric Type: 2 (Larger than any link state path)
        TOS: 0
        Metric: 1
        Forward Address: 0.0.0.0
        External Route Tag: 1
```

On R3:

```
R3#Show ip route 0.0.0.0 0.0.0.0

Routing entry for 0.0.0.0/0, supernet
  Known via "ospf 1", distance 110, metric 1, candidate default path
  Tag 1, type extern 2, forward metric 1
  Last update from 6.14.35.5 on FastEthernet0/0, 00:01:28 ago
  Routing Descriptor Blocks:
  * 6.14.35.5, from 6.14.0.5, 00:01:28 ago, via FastEthernet0/0
      Route metric is 1, traffic share count is 1
      Route tag 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Now that the default route is injected, let's test the condition:

On R3:

```
R3(config)#Int lo0
R3(config-if)#Shut

R3#Show ip route 0.0.0.0 0.0.0.0

% Network not in table
```

Let's "No shut" the interface:

```
R3(config)#Int lo0
R3(config-if)#No shut

R3#Show ip route 0.0.0.0 0.0.0.0

Routing entry for 0.0.0.0/0, supernet
  Known via "ospf 1", distance 110, metric 1, candidate default path
  Tag 1, type extern 2, forward metric 1
  Last update from 6.14.35.5 on FastEthernet0/0, 00:00:30 ago
  Routing Descriptor Blocks:
  * 6.14.35.5, from 6.14.0.5, 00:00:30 ago, via FastEthernet0/0
      Route metric is 1, traffic share count is 1
      Route tag 1
```

Ticket 7

R2 is configured to summarize R1's Loopback interfaces (Lo100 – Lo103). However R5 can see all the specific routes of the summary in its routing table.

Let's verify the ticket:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      6.0.0.0/8 is variably subnetted, 16 subnets, 3 masks
O IA     6.14.0.1/32 [110/66] via 6.14.245.2, 03:42:50, Tunnel1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O        6.14.0.2/32 [110/65] via 6.14.245.2, 03:49:50, Tunnel1
O        6.14.0.3/32 [110/2] via 6.14.35.3, 00:00:56, FastEthernet0/0
O        6.14.0.4/32 [110/65] via 6.14.245.4, 03:23:38, Tunnel1
O IA     6.14.0.6/32 [110/66] via 6.14.245.4, 00:09:58, Tunnel1
O IA     6.14.12.0/23 [110/66] via 6.14.245.2, 03:42:45, Tunnel1
O IA     6.14.12.0/24 [110/65] via 6.14.245.2, 03:49:50, Tunnel1
O IA     6.14.46.0/24 [110/65] via 6.14.245.4, 03:23:38, Tunnel1
      13.0.0.0/32 is subnetted, 4 subnets
O E2     13.68.5.100 [110/20] via 6.14.245.2, 03:42:45, Tunnel1
O E2     13.68.5.101 [110/20] via 6.14.245.2, 03:42:45, Tunnel1
O E2     13.68.5.102 [110/20] via 6.14.245.2, 03:42:45, Tunnel1
O E2     13.68.5.103 [110/20] via 6.14.245.2, 03:42:45, Tunnel1
```

Let's se how R2 was configuerd to summarize these routes:

On R2:

```
R2#Show run | S router ospf

router ospf 1
 log-adjacency-changes
 summary-address 13.68.5.0 255.255.255.0
 network 6.14.0.0 0.0.0.255 area 0
 network 6.14.12.0 0.0.0.255 area 12
 network 6.14.245.0 0.0.0.255 area 0
```

The "Summary" was configured correctly, but since R2 is NOT the router that originated the routes the summary will NOT work. We have two options to fix this problem:

1. R1 can perform the summary
2. Force R2 to originate the specific routes, this can be done by converting the area into a NSSA.

In this case the second option is configured:

On R1 and R2:

```
RX(config)#router ospf 1
RX(config-router)#Area 12 nssa
```

To verify the configuration:

On R5:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      6.0.0.0/8 is variably subnetted, 16 subnets, 3 masks
O IA     6.14.0.1/32 [110/66] via 6.14.245.2, 00:00:17, Tunnel1
O        6.14.0.2/32 [110/65] via 6.14.245.2, 03:51:44, Tunnel1
O        6.14.0.3/32 [110/2] via 6.14.35.3, 00:02:50, FastEthernet0/0
O        6.14.0.4/32 [110/65] via 6.14.245.4, 03:25:32, Tunnel1
O IA     6.14.0.6/32 [110/66] via 6.14.245.4, 00:11:52, Tunnel1
O IA     6.14.12.0/23 [110/66] via 6.14.245.2, 00:00:17, Tunnel1
O IA     6.14.12.0/24 [110/65] via 6.14.245.2, 03:51:44, Tunnel1
O IA     6.14.46.0/24 [110/65] via 6.14.245.4, 03:25:32, Tunnel1
      13.0.0.0/24 is subnetted, 1 subnets
O E2     13.68.5.0 [110/20] via 6.14.245.2, 00:00:12, Tunnel1
```

Ticket 8

Erase the startup configuration and reload the devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

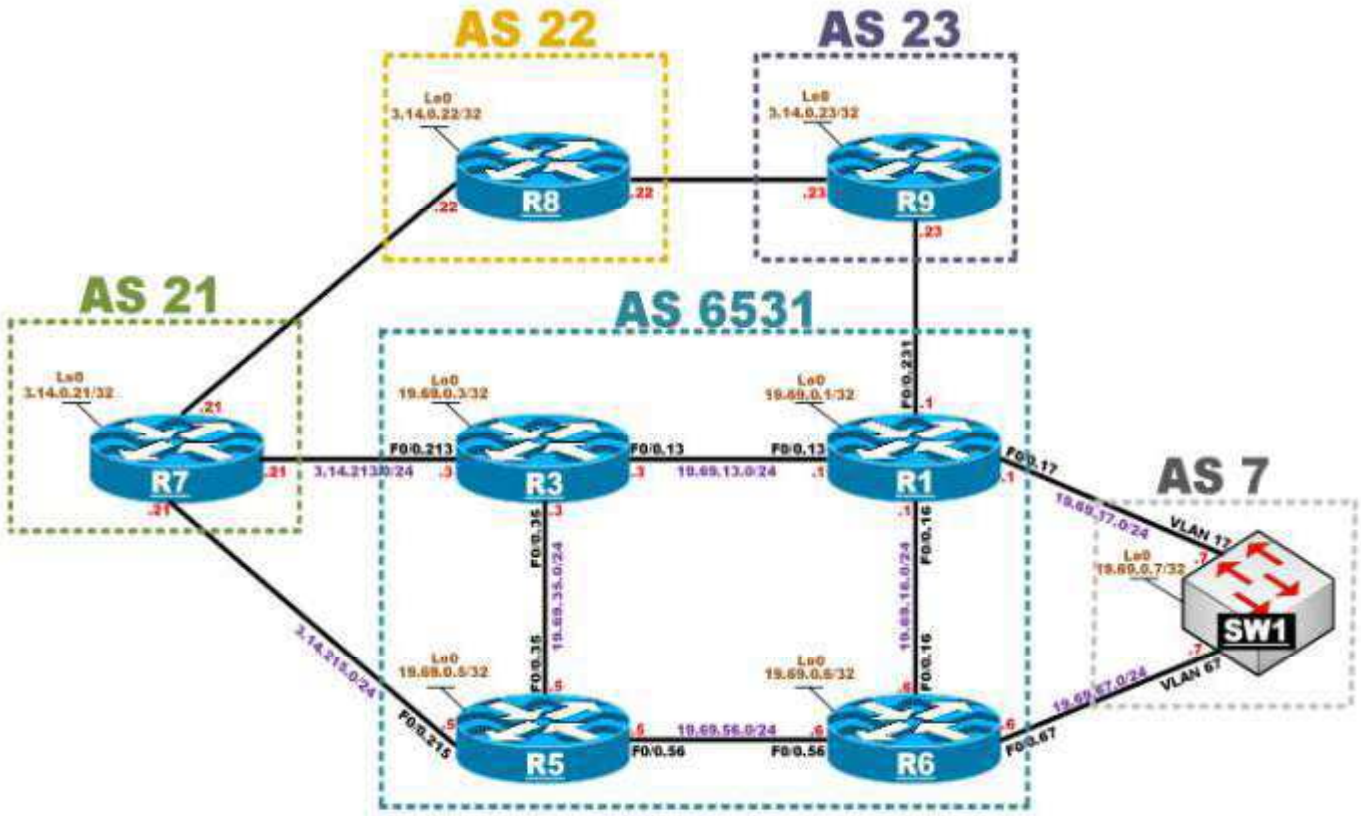
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

BGP

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Troubleshooting BGP Scenario - I



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "BGP" → "TS-BGP-Initial-Lab-1".

Ticket 1

R5 is not establishing a BGP peer session with R7.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the ticket:

On R5:

```
R5#Show ip bgp summary | B Neighbor

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.14.215.21     4           21       0       0        0    0    0 never    Active
19.69.35.3      4         6531      40      40       11    0    0 00:32:46        4
19.69.56.6      4         6531      39      39       11    0    0 00:32:12        5
```

Let's try the "Show ip bgp neighbor 3.14.215.21" command it will provide more details:

```
R5#Show ip bgp nei 3.14.215.21

BGP neighbor is 3.14.215.21,  remote AS 21, external link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Active
  Last read 00:48:01, last write 00:48:01, hold time is 180, keepalive
interval is 60 seconds
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  0          0
    Notifications:          0          0
    Updates:                0          0
    Keepalives:             0          0
    Route Refresh:          0          0
    Total:                  0          0
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 11, neighbor version 0/0
  Output queue size : 0
  Index 2, Offset 0, Mask 0x4
  2 update-group member
  Outbound path policy configured
  Outgoing update AS path filter list is 1
  Route map for outgoing advertisements is rmR7
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               0          0
    Prefixes Total:                 0          0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

                                  Outbound     Inbound
  Local Policy Denied Prefixes:   --------     -------
    Total:                               0           0
  Number of NLRIs in the update sent: max 0, min 0

  Address tracking is enabled, the RIB does have a route to 3.14.215.21

  Connections established 0; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  TCP session must be opened passively
  No active TCP connection
```

We can clearly see the problem, R5 and R7 are both configured to be passive. To verify this:

```
R5#Show run | Inc passive

 neighbor 3.14.215.21 transport connection-mode passive
```

On R7:

```
R7>Sh ip bgp neighbor 3.14.215.5 | Inc pass

  TCP session must be opened passively
```

When two routers are trying to establish a BGP peer session at least one router has to be configured as active, the active router will initiate the session, and since it initiates the session it will become the client and the router that is configured as passive will be the server, they can't be both servers, the server passively waits for an active client to establish this session, therefore, if they are both configured as "passive" none of them are initiating the session. Let's change R5 to be the client:

On R5

```
R5(config)#router bgp 6531
R5(config-router)#No neigh 3.14.215.21 transport connection-mode passive
R5(config-router)#neigh 3.14.215.21 transport connection-mode active
```

You can see the neighbor adjacency being established once the command is removed, but in troubleshooting we should fix the problem with the configuration and remove command/s as the

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

absolute last resort.

```
R5#Clear ip bgp *
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 3.14.215.21 Up
```

```
R5#Show ip bgp summary | B Neigh
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.14.215.21	**4**	**21**	**7**	**5**	**15**	**0**	**0**	**00:02:54**	**1**
19.69.35.3	4	6531	71	70	15	0	0	00:03:08	4
19.69.56.6	4	6531	70	69	15	0	0	00:03:08	5

Ticket 2

R3 can't reach R7's Lo0 (3.14.0.21).

Let's verify the ticket:

On R3

```
R3#Ping 3.14.0.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see if we can ping the directly connected interface of R7:

```
R3#Ping 3.14.213.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.213.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Let's check and see if R3 has established a neighbor adjacency with R7:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip bgp summ | B Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.14.213.21     4    21       0       0        1    0    0 never    Idle
19.69.13.1      4  6531      31      29       12    0    0 00:19:41        7
19.69.35.5      4  6531      14      13       12    0    0 00:05:37        5
```

Since the R7 router is reachable, let's check R3's BGP configuration for R7:

```
R3#Show run | S router bgp

router bgp 6531
 no synchronization
 bgp log-neighbor-changes
 network 19.69.0.1 mask 255.255.255.255
 network 19.69.0.3 mask 255.255.255.255
 network 19.69.0.5 mask 255.255.255.255
 network 19.69.0.6 mask 255.255.255.255
 neighbor 3.14.213.21 remote-as 21
 neighbor 3.14.213.21 filter-list 1 out
 neighbor 19.69.13.1 remote-as 6531
 neighbor 19.69.35.5 remote-as 6531
```

The configuration looks fine, let's enable debugging using an access-list:

```
R3(config)#Access-list 100 permit tcp any host 3.14.213.21
R3(config)#Access-list 100 permit tcp host 3.14.213.21 any

R3#Debug ip packet det 100 dump

IP: tableid=0, s=3.14.213.21 (FastEthernet0/0.213), d=3.14.213.3
(FastEthernet0/0.213), routed via RIB

IP: s=3.14.213.21 (FastEthernet0/0.213), d=3.14.213.3
(FastEthernet0/0.213), len 85, rcvd 3

TCP src=33454, dst=179, seq=2480203895, ack=1044490382, win=16384 ACK
0F7FE920:                         0011 92A03A10            ... :.
0F7FE930: 000D6584 A8800800 45C00055 86B10000  ..e.(...E@.U.1..
0F7FE940: FE0684FC 030ED515 030ED503 82AE00B3  ~..|..U...U....3
0F7FE950: 93D4E877 3E41A88E 50104000 B34E0000  .Thw>A(.P.@.3N..
0F7FE960: FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF  ................
0F7FE970: 002D0104 001500B4 030E0015 10020601  .-.....4........
0F7FE980: 04000100 01020280 00020202 00        .............
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's disable the debug:

```
R3#Undebug all
All possible debugging has been turned off
```

We can see the problem, the highlighted areas should be our focus, these are incoming packets from R7, the highlighted "**FE**"indicates the TTL, and if this Hexadecimal value is converted to decimal you should see that the TTL is set to 254.

There are two reasons as to why R7 is sending a TTL of 254:

1. R7 was configured with the "Ebgp-multihop 254".
2. R7 was configured with the "TTL-security hops 1".

If R7 is configured using the "Ebgp-multihop 254", then this configuration should have worked, but if the R7 router is configured with the "TTL-security hops 1", then it would accept BGP packets with TTL of 254 or higher.

Let's try to configure R3 with "TTL-security hops 1" and see the result:

```
R3(config)#router bgp 6531
R3(config-router)#Neighbor 3.14.213.21 ttl-security hops 1
```

You should see the following console message stating that the neighbor adjacency is now up.

```
%BGP-5-ADJCHANGE: neighbor 3.14.213.21 Up
```

Ticket 3

R1 can't reach R8's Lo0 from its Lo0 interface. All routers in AS6531 are having the same problem. DO NOT change the BGP design to fix this problem.

Let's verify the ticket:

On R1

```
R1#Ping 3.14.0.22 source loopback 0

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 3.14.0.22, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.1
.....
Success rate is 0 percent (0/5)
```

Let's see the routing table of R1:

```
R1#Show ip route 3.14.0.22

Routing entry for 3.14.0.22/32
  Known via "bgp 6531", distance 20, metric 0
  Tag 23, type external
  Last update from 3.14.231.23 00:28:12 ago
  Routing Descriptor Blocks:
  * 3.14.231.23, from 3.14.231.23, 00:28:12 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 23
      MPLS label: none
```

R1 goes through R9 to reach R8's Lo0. Let's traceroute:

```
R1#Traceroute 3.14.0.22

Type escape sequence to abort.
Tracing the route to 3.14.0.22

  1 3.14.231.23 0 msec 4 msec 0 msec
  2  *  *  *
  3  *  *  *
  4  *
```

It is going to R9 and it's stopping there. Let's see if we can Ping R9's Lo0:

```
R1#Ping 3.14.0.23 source loopback 0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.23, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.1
.....
Success rate is 0 percent (0/5)
```

Let's check our routing table for R9's loopback0 interface:

```
R1#Show ip route 3.14.0.23
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 3.14.0.23/32
  Known via "bgp 6531", distance 20, metric 0
  Tag 23, type external
  Last update from 3.14.231.23 00:30:40 ago
  Routing Descriptor Blocks:
  * 3.14.231.23, from 3.14.231.23, 00:30:40 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 23
      MPLS label: none
```

We do have a route to R9's Lo0 but we can't Ping, may be packets get there but R9 does not have a route back to our Lo0's IP address. Let's see what routes are advertised to R9:

```
R1#Show ip bgp neighbor 3.14.231.23 advertised-routes

Total number of prefixes 0
```

Let's check the BGP table for R1's Loopback interface:

```
R1#Show ip bgp 19.69.0.1

BGP routing table entry for 19.69.0.1/32, version 2
Paths: (3 available, best #3, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1    2
Local
 19.69.13.3 from 19.69.13.3 (19.69.0.3)
  Origin IGP, metric 11, localpref 100, valid, internal
Local
 19.69.16.6 from 19.69.16.6 (19.69.0.6)
  Origin IGP, metric 11, localpref 100, valid, internal
Local
 0.0.0.0 from 0.0.0.0 (19.69.0.1)
  Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
```

Let's check and see the policies that are configured for R9 on our router (R1):

```
R1#Show ip bgp neighbors 3.14.231.23 policy detail

 Neighbor: 3.14.231.23, Address-Family: IPv4 Unicast
 Locally configured policies:
  filter-list 1 out
 Neighbor: 3.14.231.23, Address-Family: IPv4 Unicast <detail>
 Locally configured policies:
  filter-list 1 out
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
AS path access list 1
    permit ^6531$
```

Now we can see the problem. R1 is configured with an as-path filter that ONLY allows the routes that originated in its own AS to be advertised to other ASes, this is a typical way to make an AS a non-transit AS.

However, the "^6531$" regular expression is NOT correct. In R1's BGP table, the 19.69.0.1/32 prefix has an empty AS path, it ONLY adds its own AS when that prefix is advertised to an EBGP peer, therefore, R1 should have been configured to match an empty AS-Path list "^$" and NOT "^6531$".

Let's correct this problem, but before correcting the problem let's have an example with an incorrect regular expression:

```
R1#Show ip bgp regexp ^6531$
R1#
```

Now, let's verify the correct regular expression before applying it as a filter:

```
R1#Show ip bgp regexp ^$

BGP table version is 12, local router ID is 19.69.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i19.69.0.0/16     19.69.16.6               0    100      0 i
*>                  0.0.0.0                               32768 i
* i19.69.0.1/32     19.69.13.3              11    100      0 i
* i                 19.69.16.6              11    100      0 i
*>                  0.0.0.0                  0             32768 i
* i19.69.0.3/32     19.69.13.3               0    100      0 i
* i                 19.69.16.6              21    100      0 i
*>                  19.69.13.3              11             32768 i
* i19.69.0.5/32     19.69.35.5              11    100      0 i
*>                  19.69.13.3              21             32768 i
* i                 19.69.56.5              11    100      0 i
* i19.69.0.6/32     19.69.13.3              21    100      0 i
*>                  19.69.16.6              11             32768 i
* i                 19.69.16.6               0    100      0 i
```

Let's implement the correct regular expression, this MUST be done on R1, R3 and R5 because these routers are connecting to the R7 and R9 routers:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1, R3 and R5:

```
RX(config)#No ip as-path access-list 1
RX(config)#IP as-path access-list 1 permit ^$
```

On R1:

```
R1#Clear ip bgp 23 out
```

On R3 and R5:

```
RX#Clear ip bgp 21 out
```

Now, let's see if the "IP as-path access-list" is applied:

```
RX#Show run | Inc filter-list

 neighbor 3.14.215.21 filter-list 1 out
```

To test the configuration:

On R1:

```
R1#Ping 3.14.0.22 source loopback 0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.22, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 4

R6 can't reach R7's Lo0.

Let's verify the ticket:

On R6:

```
R6#Ping 3.14.0.21
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the routing table:

```
R6#Show ip route 3.14.0.21

% Network not in table
```

Let's see if the other routers in this AS (AS 6531) have a route to R7's Lo0 interface:

On R1:

```
R1#Show ip route 3.14.0.21
% Subnet not in table
```

On R3:

```
R3#Show ip route 3.14.0.21

Routing entry for 3.14.0.21/32
  Known via "bgp 6531", distance 20, metric 0
  Tag 21, type external
  Last update from 3.14.213.21 00:48:27 ago
  Routing Descriptor Blocks:
  * 3.14.213.21, from 3.14.213.21, 00:48:27 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 21
      MPLS label: none

R3#Ping 3.14.0.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

We can see that R3 has reachability.

On R5:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Sh ip route 3.14.0.21
% Subnet not in table
```

Let's see the BGP entry for this route on all routers in this AS:

On R1:

```
R1#Show ip bgp 3.14.0.21

BGP routing table entry for 3.14.0.21/32, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  21
    3.14.213.21 (inaccessible) from 19.69.13.3 (19.69.0.3)
      Origin IGP, metric 0, localpref 100, valid, internal
```

On R3:

```
R3#Show ip bgp 3.14.0.21

BGP routing table entry for 3.14.0.21/32, version 11
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1
  21
    3.14.213.21 from 3.14.213.21 (3.14.0.21)
      Origin IGP, metric 0, localpref 100, valid, external, best
```

On R5:

```
R5#Show ip bgp 3.14.0.21

BGP routing table entry for 3.14.0.21/32, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  21
    3.14.213.21 (inaccessible) from 19.69.35.3 (19.69.0.3)
      Origin IGP, metric 0, localpref 100, valid, internal
```

On R6:

```
R6#Sh ip bgp 3.14.0.21

% Network not in table
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Well, when R3 receives prefix 3.14.0.21/32 from R7, it advertises to its iBGP peers but the next hop remains unchanged, we can redistribute R3's F0/0.213 into OSPF so the routers in AS 6531 can have reachability to it. Let's configure this and verify:

On R3:

```
R3(config)#Route-map tst per 10
R3(config-route-map)#match inter f0/0.213

R3(config)#Router ospf 1
R3(config-router)#Redistr connected route-map tst sub
```

To verify the configuration:

On R1:

```
R1#Ping 3.14.0.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R5:

```
R5#Ping 3.14.0.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.5
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R6:

```
R6#Ping 3.14.0.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.6
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R1 and R5 have reachability to the R7's Lo0, but not R6. Let's see R6's neighbors:

On R6

```
R6#Show ip bgp summary | B Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
19.69.16.1	4	6531	127	126	8	0	0	01:48:39	6
19.69.56.5	4	6531	113	112	8	0	0	01:34:56	4
19.69.67.7	4	7	120	127	8	0	0	01:48:43	0

R6 does NOT have an adjacency to R3, it only has an adjacency to R1 and R5. Let's check the adjacency on R1:

On R1:

```
R1#Show ip bgp summary | B Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.14.231.23	4	23	129	126	11	0	0	01:48:15	2
19.69.13.3	4	6531	132	133	11	0	0	01:51:34	7
19.69.16.6	4	6531	129	130	11	0	0	01:51:14	5
19.69.17.7	4	7	123	132	11	0	0	01:51:13	0

R1 does not have an adjacency to R5. Let's check R5:

On R5:

```
R5#Show ip bgp summary | B Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.14.215.21	4	21	113	112	9	0	0	01:38:31	1
19.69.35.3	4	6531	116	115	9	0	0	01:38:31	6
19.69.56.6	4	6531	116	117	9	0	0	01:38:31	4

R5 dos not have an adjacency to R1. Let's check R3:

On R3:

```
R3#Show ip bgp summary | B Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.14.213.21	4	21	103	99	15	0	0	01:23:26	3
19.69.13.1	4	6531	135	134	15	0	0	01:53:37	7
19.69.35.5	4	6531	116	118	15	0	0	01:39:33	5

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

WOW…R3 does not have an adjacency to R6. Let's fix this problem:

On R3:

```
R3(config)#router bgp 6531
R3(config-router)#neighbor 19.69.0.6 remote-as 6531
R3(config-router)#neighbor 19.69.0.6 update-source Lo0
```

On R1:

```
R1(config)#router bgp 6531
R1(config-router)#neighbor 19.69.0.5 remote-as 6531
R1(config-router)#neighbor 19.69.0.5 update-source Lo0
```

On R5:

```
R5(config)#router bgp 6531
R5(config-router)#neighbor 19.69.0.1 remote-as 6531
R5(config-router)#neighbor 19.69.0.1 update-source Lo0
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 19.69.0.1 Up
```

On R6:

```
R6(config)#router bgp 6531
R6(config-router)#neighbor 19.69.0.3 remote-as 6531
R6(config-router)#neighbor 19.69.0.3 update-source Lo0
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 19.69.0.3 Up
```

To test the configuration:

On R6:

```
R6#Ping 3.14.0.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.6
!!!!!
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Ticket 5

AS 6531 is configured such that R7 in AS 21 prefers R5 as the next hop to reach the Lo0 interface of R6, but this is not working. You should use an interface configuration mode command to fix this problem.

Since we have no access to the BB routers one way to test this ticket is to configure an access-list with a "log" option and then check for the number of matches. This access-list must be configured on both R3 and R5 since these are the two border routers:

On R3:

```
R3(config)#Access-list 101 permit icmp host 3.14.0.21 host 19.69.0.6
R3(config)#Access-list 101 permit ip any any

R3(config)#int f0/0.213
R3(config-subif)#ip access-group 101 in
```

On R5:

```
R5(config)#Access-list 101 permit icmp host 3.14.0.21 host 19.69.0.6
R5(config)#Access-list 101 permit ip any any

R5(config)#Int f0/0.215
R5(config-subif)#ip access-group 101 in
```

Now, let's ping from R6's Lo0 to R7's Lo0 and check for matches:

On R6:

```
R6#Ping 3.14.0.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.6
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's go to R3 and see the matches:

On R3:

```
R3#Show access-list 101

Extended IP access list 101
    10 permit icmp host 3.14.0.21 host 19.69.0.6 (5 matches)
    20 permit ip any any (3 matches)
```

On R5:

```
R5#Show access-list 100

Extended IP access list 101
    10 permit icmp host 3.14.0.21 host 19.69.0.6
    20 permit ip any any (5 matches)
```

We can see the matches on R3 and NOT R5. Let's see how AS 6531 was configured to influence the returning path from R7 to R6's Lo0, this verification should be done on R3 and R5:

On R3:

```
R3#Show ip bgp neighbors  3.14.213.21 policy

 Neighbor: 3.14.213.21, Address-Family: IPv4 Unicast
 Locally configured policies:
  filter-list 1 out
```

On R5:

```
R5#Show ip bgp neighbors  3.14.215.21 policy

 Neighbor: 3.14.215.21, Address-Family: IPv4 Unicast
 Locally configured policies:
  filter-list 1 out
  route-map rmR7 out
```

We can see the "filter-list" from Ticket 3 on both R3 and R5, but we only see a route-map called "rmR7" on R5. Let's verify the route-map:

```
R5#Show run | s route-map rmR7

 neighbor 3.14.215.21 route-map rmR7 out
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
route-map rmR7 permit 10
 match ip address prefix-list p1R6Lo0
 set metric 100
```

Let's verify the ip prefix-list "p1R6Lo0":

```
R5#Show ip prefix-list p1R6Lo0

ip prefix-list p1R6Lo0: 1 entries
   seq 5 permit 19.69.0.6/32
```

We can see that R5 was configured to change the MED value to 100 when sending the prefix to R7. Let's check the MED value for the same prefix on R3:

On R3:

```
R3#Show ip bgp 19.69.0.6

BGP routing table entry for 19.69.0.6/32, version 8
Paths: (4 available, best #4, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1     2
  Local
    19.69.0.6 (metric 21) from 19.69.0.6 (19.69.0.6)
      Origin IGP, metric 0, localpref 100, valid, internal
  Local
    19.69.56.6 (metric 20) from 19.69.35.5 (19.69.0.5)
      Origin IGP, metric 11, localpref 100, valid, internal
  Local
    19.69.16.6 (metric 20) from 19.69.13.1 (19.69.0.1)
      Origin IGP, metric 11, localpref 100, valid, internal
  Local
    19.69.13.1 from 0.0.0.0 (19.69.0.3)
      Origin IGP, metric 21, localpref 100, weight 32768, valid, sourced, local,
best
```

R3 is using a MED value of 21 when sending this prefix to R7 and since it's lower than the MED value advertised by R5 (Which is 100), R7 is preferring the route via R3 to R6's Lo0. But how can we use an interface command to resolve this problem?

To answer this question we need to find out how is R3 assigning a MED value to this prefix when we did not see any policy applied, let's check the routing table of R3 for R6's Lo0:

```
R3#Show ip route 19.69.0.6

Routing entry for 19.69.0.6/32
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Known via "ospf 1", distance 110, metric 21, type intra area
  Advertised by bgp 6531
  Last update from 19.69.13.1 on FastEthernet0/0.13, 2d16h ago
  Routing Descriptor Blocks:
  * 19.69.35.5, from 19.69.0.6, 2d16h ago, via FastEthernet0/0.35
      Route metric is 21, traffic share count is 1
    19.69.13.1, from 19.69.0.6, 2d16h ago, via FastEthernet0/0.13
      Route metric is 21, traffic share count is 1
```

R3 used the OSPF cost and gave it to BGP and BGP used it as MED, so all we need to do is configure the appripriate interfaces such that they will have a cost of 101; but which interfaces are the ones that we need to configure? The interfaces to R5 and R1, these are F0/0.35 and F0/0.13 respectively. Let's see the existing cost:

```
R3#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               19.69.0.3/32       1     LOOP  0/0
Fa0/0.35     1     0               19.69.35.3/24      10    BDR   1/1
Fa0/0.13     1     0               19.69.13.3/24      10    DR    1/1

R3(config)#Int f0/0.13
R3(config-subif)#IP ospf cost 100

R3(config)#int f0/0.35
R3(config-subif)#IP ospf cost 100
```

Let's verify the configuration:

On R3:

```
R3#Show ip route 19.69.0.6

Routing entry for 19.69.0.6/32
  Known via "ospf 1", distance 110, metric 111, type intra area
  Advertised by bgp 6531
  Last update from 19.69.35.5 on FastEthernet0/0.35, 00:00:56 ago
  Routing Descriptor Blocks:
    19.69.35.5, from 19.69.0.6, 00:00:56 ago, via FastEthernet0/0.35
      Route metric is 111, traffic share count is 1
  * 19.69.13.1, from 19.69.0.6, 00:00:56 ago, via FastEthernet0/0.13
      Route metric is 111, traffic share count is 1
```

This looks good. So why was the cost 21?

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

For R3 to get to R6's Lo0 interface, R3 has to go through R1 and the cost of the F0/0.13 sub-interface to R1 is 10, and R1 has to go to R6 using its F0/0.16 interface which is another FastEthernet sub-interface (F0/0.16) with a cost of 10, plus the cost of R6's Lo0 which is 1, so 10+10+1=21.

To verify this value in BGP:

```
R3#Show ip bgp 19.69.0.6

BGP routing table entry for 19.69.0.6/32, version 17
Paths: (4 available, best #4, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1    2
  Local
    19.69.0.6 (metric 111) from 19.69.0.6 (19.69.0.6)
      Origin IGP, metric 0, localpref 100, valid, internal
  Local
    19.69.56.6 (metric 110) from 19.69.35.5 (19.69.0.5)
      Origin IGP, metric 11, localpref 100, valid, internal
  Local
    19.69.16.6 (metric 110) from 19.69.13.1 (19.69.0.1)
      Origin IGP, metric 11, localpref 100, valid, internal
  Local
    19.69.13.1 from 0.0.0.0 (19.69.0.3)
      Origin IGP, metric 111, localpref 100, weight 32768, valid, sourced,
local, best
```

To test the configuration:

On R6:

Let's clear the access-list counters on both routers:

On R3 and R5:

```
RX#Clear access-list counters
```

On R3:

```
R3#Show access-list 101

Extended IP access list 101
    10 permit icmp host 3.14.0.21 host 19.69.0.6
    20 permit ip any any (6 matches)
```

On R5:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show access-list 100

Extended IP access list 100
    10 permit icmp host 3.14.0.21 host 19.69.0.6
    20 permit ip any any
```

Now, let's ping from R6's Lo0

```
R6#Ping 3.14.0.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.0.21, timeout is 2 seconds:
Packet sent with a source address of 19.69.0.6
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Let's verify the access-lists to see the incoming traffic:

On R3:

```
R3#Show access-list 101

Extended IP access list 101
    10 permit icmp host 3.14.0.21 host 19.69.0.6
    20 permit ip any any (18 matches)
```

On R5:

```
R5#Show access-list 100

Extended IP access list 100
    10 permit icmp host 3.14.0.21 host 19.69.0.6 (5 matches)
    20 permit ip any any (4 matches)
```

Ticket 6

R6 should load balance traffic destined to R8's Lo0 via R1 and R5, but this is not working.

Let's verify the ticket:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Show ip route 3.14.0.22

Routing entry for 3.14.0.22/32
  Known via "bgp 6531", distance 200, metric 0
  Tag 23, type internal
  Last update from 3.14.231.23 1d20h ago
  Routing Descriptor Blocks:
  * 3.14.231.23, from 19.69.16.1, 1d20h ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 23
```

Sure enough it's using R1 only. Let's check R6's BGP table for this prefix:

```
R6#Show ip bgp 3.14.0.22

BGP routing table entry for 3.14.0.22/32, version 10
Paths: (3 available, best #2, table default)
  Advertised to update-groups:
     2
  21 22
    3.14.215.21 (inaccessible) from 19.69.56.5 (19.69.0.5)
      Origin IGP, metric 0, localpref 100, valid, internal
  21 22
    3.14.213.21 (metric 20) from 19.69.0.3 (19.69.0.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
  23 22
    3.14.231.23 (inaccessible) from 19.69.16.1 (19.69.0.1)
      Origin IGP, metric 0, localpref 100, valid, internal
```

Does R6 have reachability to 3.14.231.23? Let's verify:

```
R6#Ping 3.14.231.23

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.14.231.23, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

We can fix this problem by configuring R1 and R5 to redistribute their external links, let's configure this and verify:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#Route-map tst
R1(config-route-map)#Match inter FastEthernet0/0.231

R1(config)#Router ospf 1
R1(config-router)#redistr connected route-map tst sub
```

On R5:

```
R5(config)#Route-map tst
R5(config-route-map)#Match inter FastEthernet0/0.215

R5(config-route-map)#router ospf 1
R5(config-router)#redistribute connected route-map tst sub
```

To verify the configuration:

On R6:

```
R6#Show ip bgp 3.14.0.22

BGP routing table entry for 3.14.0.22/32, version 12
Paths: (3 available, best #3, table default)
  Advertised to update-groups:
     2
  21 22
    3.14.215.21 (metric 20) from 19.69.56.5 (19.69.0.5)
      Origin IGP, metric 0, localpref 100, valid, internal
  21 22
    3.14.213.21 (metric 20) from 19.69.0.3 (19.69.0.3)
      Origin IGP, metric 0, localpref 100, valid, internal
  23 22
    3.14.231.23 (metric 20) from 19.69.16.1 (19.69.0.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R6#Show ip route 3.14.0.22

Routing entry for 3.14.0.22/32
  Known via "bgp 6531", distance 200, metric 0
  Tag 23, type internal
  Last update from 3.14.231.23 00:05:32 ago
  Routing Descriptor Blocks:
  * 3.14.231.23, from 19.69.16.1, 00:05:32 ago
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 23
      MPLS label: none
```

From BGP's perspective, R6 has three paths to R8's Lo0 interface. Maybe load balancing was NOT enabled, let's verify:

```
R6#Show ip protocols | B bgp

Routing Protocol is "bgp 6531"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  IGP synchronization is disabled
  Automatic route summarization is disabled
  Unicast Aggregate Generation:
    19.69.0.0/16
  Neighbor(s):
    Address          FiltIn FiltOut DistIn DistOut Weight RouteMap
    19.69.0.3
    19.69.16.1
    19.69.56.5
    19.69.67.7
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    19.69.16.1           200      00:06:42
    19.69.0.3            200      01:14:55
    19.69.56.5           200      03:12:43
  Distance: external 20 internal 200 local 200
```

We can see that BGP is configured to allow a single path, let's correct this problem:

On R6:

```
R6(config)#router bgp 6531
R6(config-router)#maximum-paths ibgp 2
```

To verify the configuration:

On R6:

```
R6#Show ip bgp 3.14.0.22
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
BGP routing table entry for 3.14.0.22/32, version 14
Paths: (3 available, best #3, table Default-IP-Routing-Table)
Multipath: iBGP
  Advertised to update-groups:
        1
  21 22
    3.14.213.21 (metric 20) from 19.69.0.3 (19.69.0.3)
      Origin IGP, metric 0, localpref 100, valid, internal
  21 22
    3.14.215.21 (metric 20) from 19.69.56.5 (19.69.0.5)
      Origin IGP, metric 0, localpref 100, valid, internal
  23 22
    3.14.231.23 (metric 20) from 19.69.16.1 (19.69.0.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

R6 is still using a single path, but we can se that the "Multipath:iBGP" is enabled.
If the output of the above command is examined closely, you should see that the AS-Path list for this prefix is NOT the same, and according to Cisco's documentation, in order to be a candidate for multipath, paths to the same destination need to have the following characteristics equal to the best path characteristics:

1. Weight
2. Local Preference
3. AS-Path length
4. Origin
5. MED
6. Neighboring AS or sub-AS before the addition of the iBGP Multipath feature
7. AS-Path after the addition of the iBGP Multipath feature

Since the AS-Path is NOT the same, this will NOT happen unless we use a hidden command to ignore or relax this rule. Let's configure this hidden command and verify:

On R6:

```
R6(config)#router bgp 6531
R6(config-router)#bgp bestpath as-path multipath-relax
```

The above hidden command allows Multipath even if the AS-Path is NOT the same.

To verify the configuration:

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Show ip route 3.14.0.22

Routing entry for 3.14.0.22/32
  Known via "bgp 6531", distance 200, metric 0
  Tag 23, type internal
  Last update from 3.14.231.23 00:00:06 ago
  Routing Descriptor Blocks:
    3.14.231.23, from 19.69.16.1, 00:00:06 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 23
      MPLS label: none
  * 3.14.215.21, from 19.69.56.5, 00:00:06 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 23
      MPLS label: none
```

Ticket 7

SW1 is a low end switch with very limited memory. R1 and R6 are configured to summarize all the prefixes in AS6531 and advertise a singe summary to SW1, but SW1 sees all the specific prefixes. The solution should ONLY affect SW1.

Let's verify but looking at the routing table of SW1:

On SW1:

```
SW1#Show ip route bgp

     19.0.0.0/8 is variably subnetted, 8 subnets, 3 masks
B       19.69.0.6/32 [20/0] via 19.69.67.6, 00:06:34
B       19.69.0.5/32 [20/11] via 19.69.67.6, 00:06:34
B       19.69.0.3/32 [20/11] via 19.69.17.1, 2d17h
B       19.69.0.0/16 [20/0] via 19.69.17.1, 2d17h
B       19.69.0.1/32 [20/0] via 19.69.17.1, 2d17h
     3.0.0.0/32 is subnetted, 3 subnets
B       3.14.0.21 [20/0] via 19.69.17.1, 1d21h
B       3.14.0.23 [20/0] via 19.69.17.1, 2d17h
B       3.14.0.22 [20/0] via 19.69.17.1, 2d17h
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

We can see the both summary (19.69.0.0/16) and all the more specific routes which are all the /32 routes. Let's verfiy how R1 and R6 are configured:

On R1:

```
R1#Show run | S bgp

router bgp 6531
 bgp log-neighbor-changes
 network 19.69.0.1 mask 255.255.255.255
 network 19.69.0.3 mask 255.255.255.255
 network 19.69.0.5 mask 255.255.255.255
 network 19.69.0.6 mask 255.255.255.255
 aggregate-address 19.69.0.0 255.255.0.0
 neighbor 3.14.231.23 remote-as 23
 neighbor 3.14.231.23 filter-list 1 out
 neighbor 19.69.0.5 remote-as 6531
 neighbor 19.69.0.5 update-source Loopback0
 neighbor 19.69.13.3 remote-as 6531
 neighbor 19.69.16.6 remote-as 6531
 neighbor 19.69.17.7 remote-as 7
```

On R6:

```
R6#Show run | S bgp

router bgp 6531
 bgp log-neighbor-changes
 bgp bestpath as-path multipath-relax
 network 19.69.0.1 mask 255.255.255.255
 network 19.69.0.3 mask 255.255.255.255
 network 19.69.0.5 mask 255.255.255.255
 network 19.69.0.6 mask 255.255.255.255
 aggregate-address 19.69.0.0 255.255.0.0
 neighbor 19.69.0.3 remote-as 6531
 neighbor 19.69.0.3 update-source Loopback0
 neighbor 19.69.16.1 remote-as 6531
 neighbor 19.69.56.5 remote-as 6531
 neighbor 19.69.67.7 remote-as 7
 maximum-paths ibgp 2
```

We can see that the "summary-only" keyword is missing, but if the "Summary-only" is configured, then R9 will also be affected. There are few ways to configure this, one way to configure this is to use the "Unsuppress-map" which is on a per neighbor basis, another way to accomplish this task is to simply configure filtering. Let's configure Filtering and verify:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1 and R6:

```
R1(config)#ip prefix-list p1SW1 deny 19.69.0.0/16 ge 17 le 32
R1(config)#ip prefix-list p1SW1 permit 0.0.0.0/0 le 32
```

On R1:

```
R1(config)#Router bgp 6531
R1(config-router)#Neighbor 19.69.17.7 prefix-list p1SW1 out
```

On R6:

```
R6(config)#Router bgp 6531
R6(config-router)#Neighbor 19.69.67.7 prefix-list p1SW1 out
```

On R1 and R6:

```
R6#Clear ip bgp 7 out
```

To verify the configuration:

On SW1:

```
SW1#Show ip route bgp

      19.0.0.0/8 is variably subnetted, 4 subnets, 3 masks
B        19.69.0.0/16 [20/0] via 19.69.17.1, 2d17h
      3.0.0.0/32 is subnetted, 3 subnets
B        3.14.0.21 [20/0] via 19.69.17.1, 1d21h
B        3.14.0.23 [20/0] via 19.69.17.1, 2d17h
B        3.14.0.22 [20/0] via 19.69.17.1, 2d17h
```

Ticket 8

Erase the startup configuration and reload the devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

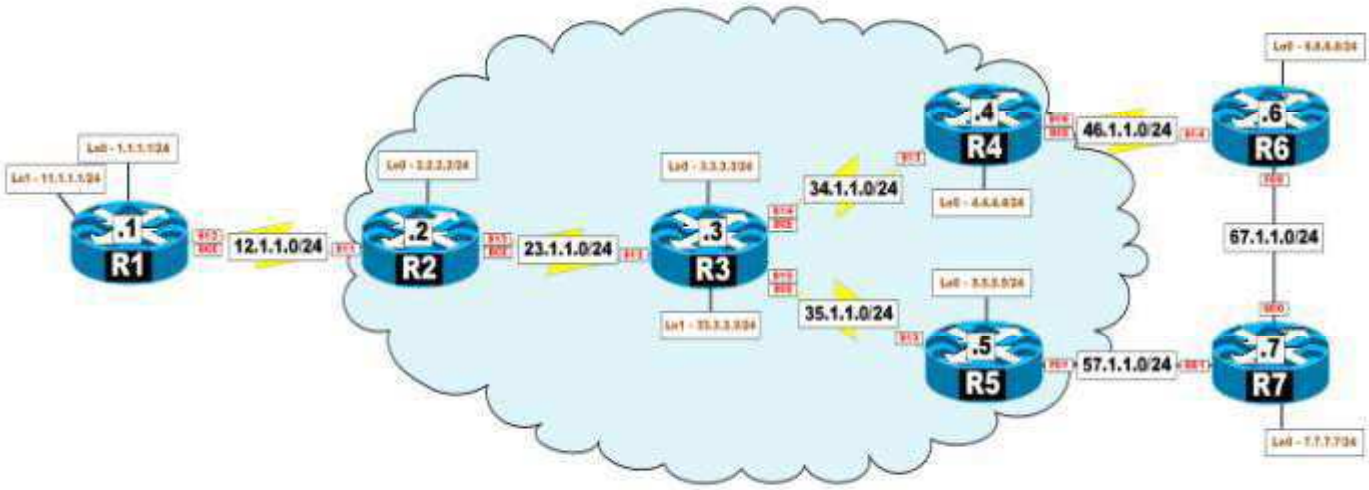
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

MPLS

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Troubleshooting MPLS Scenario - I



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "MPLS" →"MPLS-TS-Lab-1".

Ticket 1

R1, R6 and R7 are the customer edge (CE) routers and are configured in the same VPN; R2, R4 and R5 are the Provider edge (PE) routers, the PEs and the CEs are running OSPF, R6 and R7 can NOT ping 1.1.1.1, or 11.1.1.1 prefixes advertised by R1. DO NOT remove any command/s to fix this ticket.

Let's verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Show ip route ospf | B Gate
Gateway of last resort is not set

      7.0.0.0/32 is subnetted, 1 subnets
O        7.7.7.7 [110/2] via 67.1.1.7, 01:26:53, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/65] via 46.1.1.4, 00:00:45, Serial1/4
      57.0.0.0/24 is subnetted, 1 subnets
O        57.1.1.0 [110/2] via 67.1.1.7, 01:08:06, FastEthernet0/0
```

Let's check R7:

On R7:

```
R7#Show ip route ospf | B Gate
Gateway of last resort is not set

      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/2] via 67.1.1.6, 01:07:43, GigabitEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/66] via 67.1.1.6, 00:00:12, GigabitEthernet0/0
      46.0.0.0/24 is subnetted, 1 subnets
O        46.1.1.0 [110/65] via 67.1.1.6, 00:00:22, GigabitEthernet0/0
```

Do we have a neighbor adjacency with the PE routers?

```
R7#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.6           1   FULL/DR         00:00:32    67.1.1.6        GigabitEthernet0/0
0.0.0.5           1   FULL/DR         00:00:39    57.1.1.5        GigabitEthernet0/1
```

On R6:

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.4           0   FULL/  -        00:00:33    46.1.1.4        Serial1/4
0.0.0.1           1   FULL/BDR        00:00:37    67.1.1.7        FastEthernet0/0
```

The neighbor adjacency is in FULL state between the CE and the PE routers. Let's see if R1 is advertising 1.1.1.1 and 11.1.1.1 prefixes.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip int br | exc unass

Interface                  IP-Address      OK? Method Status                Protocol
Serial1/2                  12.1.1.1        YES manual up                    up
Loopback0                  1.1.1.1         YES manual up                    up
Loopback1                  11.1.1.1        YES manual up                    up

R1#Show run | s router ospf

router ospf 1
 router-id 0.0.0.1
 redistribute connected subnets route-map tst
 network 12.1.1.1 0.0.0.0 area 0
```

R1 is not running OSPF on its lo0 interface, let's fix this problem before we look at Lo1 interface.

```
R1(config)#router ospf 1
R1(config-router)#Network 1.1.1.1 0.0.0.0 a 0
```

To verify the configuration:

On R6:

```
R6#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/129] via 46.1.1.4, 00:00:22, Serial1/4
      7.0.0.0/32 is subnetted, 1 subnets
O        7.7.7.7 [110/2] via 67.1.1.7, 01:39:35, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/65] via 46.1.1.4, 00:07:54, Serial1/4
      57.0.0.0/24 is subnetted, 1 subnets
O        57.1.1.0 [110/2] via 67.1.1.7, 01:20:48, FastEthernet0/0
```

On R7:

```
R7#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/130] via 67.1.1.6, 00:00:49, GigabitEthernet0/0
      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/2] via 67.1.1.6, 01:21:14, GigabitEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O IA     12.1.1.0 [110/66] via 67.1.1.6, 00:08:21, GigabitEthernet0/0
      46.0.0.0/24 is subnetted, 1 subnets
O        46.1.1.0 [110/65] via 67.1.1.6, 00:13:53, GigabitEthernet0/0
```

Let's see if we can ping this IP address:

On R6:

```
R6#Ping 1.1.1.1 source lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 6.6.6.6
.....
Success rate is 0 percent (0/5)
```

On R7:

```
R7#Ping 1.1.1.1 source lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 7.7.7.7
.....
Success rate is 0 percent (0/5)
```

The prefix is in the VPN which means that the control plane is working correctly, but the problem is in the data plane. We can resolve this issue later, let's find out the problem with 11.1.1.1 prefix.

On R1:

```
R1#Show run | s router ospf
router ospf 1
 router-id 0.0.0.1
 redistribute connected subnets route-map tst
 network 1.1.1.1 0.0.0.0 area 0
 network 12.1.1.1 0.0.0.0 area 0
```

It looks like the prefix was redistributed, let's verify the route-map:

```
R1#Show run | s route-map
 redistribute connected subnets route-map tst
route-map tst permit 10
 match interface Loopback1
 set tag 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The Lo1 interface is referenced and tagged in the route-map and the route-map is referenced in the "redistribute connected" command; let's verify the PE (R2) and see if 1.1.1.1 and 11.1.1.1 are in the VPN:

On R2:

```
R2#Show ip route vrf aa | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/65] via 12.1.1.1, 00:01:49, Serial1/1
      6.0.0.0/32 is subnetted, 1 subnets
B        6.6.6.6 [200/65] via 4.4.4.4, 00:04:56
      7.0.0.0/32 is subnetted, 1 subnets
B        7.7.7.7 [200/66] via 4.4.4.4, 00:03:48
      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, Serial1/1
L        12.1.1.2/32 is directly connected, Serial1/1
      46.0.0.0/24 is subnetted, 1 subnets
B        46.1.1.0 [200/0] via 4.4.4.4, 00:05:06
      57.0.0.0/24 is subnetted, 1 subnets
B        57.1.1.0 [200/66] via 4.4.4.4, 00:03:48
      67.0.0.0/24 is subnetted, 1 subnets
B        67.1.1.0 [200/65] via 4.4.4.4, 00:04:56
```

1.1.1.1/32 prefix is in the VRF but not 11.1.1.1, let's see if the route-target is exported/imported and the redistribution has been performed:

```
R2#Show ip vrf detail

VRF aa (VRF Id = 1); default RD 1:10; default VPNID <not set>
  Interfaces:
    Se1/1
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix

R2#Show run | S address-family ipv4 vrf aa

 address-family ipv4 vrf aa
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  redistribute ospf 1

R2#Show run | s router ospf 1

router ospf 1 vrf aa
 domain-tag 100
 redistribute bgp 100 subnets
 network 12.1.1.2 0.0.0.0 area 0
```

Yes, the route-target is imported and exported, and the redistribution is performed correctly from BGP into OSPF and OSPF into BGP. Let's see if the prefixes are in the MP-BGP:

```
R2#Show ip bgp vpnv4 vrf aa | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf aa)
*> 1.1.1.1/32       12.1.1.1                 65         32768 ?
*>i6.6.6.6/32       4.4.4.4                  65     100     0 ?
*>i7.7.7.7/32       4.4.4.4                  66     100     0 ?
*> 12.1.1.0/24      0.0.0.0                   0         32768 ?
*>i46.1.1.0/24      4.4.4.4                   0     100     0 ?
*>i57.1.1.0/24      4.4.4.4                  66     100     0 ?
*>i67.1.1.0/24      4.4.4.4                  65     100     0 ?
```

The 1.1.1.1/32 prefix is there but not the 11.1.1.1, since there are multiple problems, let's fix one at a time, let's focus on 1.1.1.1/32 prefix; this prefix is in the routing table of R6 and R7 but they don't have reachability to it, let's start from R6 and verify the configuration hop by hop all the way to R2:

On R6:

```
R6#Traceroute 1.1.1.1 numeric

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 46.1.1.4 16 msec 16 msec 16 msec
  2  *  *  *
  3  *  *  *
  4  *  *  *
  5  *  *  *
```

The traceroute stopped at PE (R4), let's check the LSP from the PE (R4) all the way back to R2:

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show mpls forwarding-table vrf aa 1.1.1.1 255.255.255.255

Local      Outgoing    Prefix           Bytes Label   Outgoing    Next Hop

Label      Label       or Tunnel Id     Switched      interface
None       203         1.1.1.1/32[V]    0             drop
```

The outgoing interface is drop, does R4 have an LDP session with R3, the "P" router?

```
R4#Show mpls ldp neighbor
R4#
```

No, it does not have an LDP session with R3. Let's see if R4 is running LDP on the interface facing R3:

```
R4#Show mpls interface

Interface              IP            Tunnel   BGP Static Operational
Serial1/3              Yes (ldp)     No       No  No     Yes
```

It is running LDP, let's see why the LDP session has not been established:

```
R4#Show mpls ldp discovery

 Local LDP Identifier:
    4.4.4.4:0
    Discovery Sources:
    Interfaces:
        Serial1/3 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0; no route
```

Let's see if the local router has a route to 3.3.3.3:

```
R4#Show ip route 3.3.3.3

% Network not in table
```

Let's see how R3 is configured:

On R3:

```
R3#Show run | Inc mpls ldp router-id

mpls ldp router-id Loopback0 force
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

OK, the LDP-ID is based on the Lo0 interface of R3, let's see the configuration of this loopback interface:

```
R3#Show run int lo0 | B inter

interface Loopback0
 ip address 3.3.3.3 255.255.255.0
end
```

Let's see if OSPF is running on the loopback0 interface:

```
R3#Show run | s router ospf

router ospf 2
 network 23.1.1.3 0.0.0.0 area 0
 network 33.3.3.3 0.0.0.0 area 0
 network 34.1.1.3 0.0.0.0 area 0
```

OK, we can see two problems:

1. OSPF is not running on the lo0 interface of R3, therefore, R4 does not have reachability to 3.3.3.3 prefix
2. The loopback is configured with a /24 subnet mask.

Let's fix these problems:

```
R3(config)#router ospf 2
R3(config-router)#netw 3.3.3.3 0.0.0.0 a 0

R3(config)#int lo0
R3(config-if)#ip addr 3.3.3.3 255.255.255.255
```

We should see the following two console messages on R3:

```
%LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (1) is UP
%LDP-5-NBRCHG: LDP Neighbor 2.2.2.2:0 (2) is UP
```

To verify the configuration:

On R4:

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 3.3.3.3:0 (1) is UP

R4#Show ip route 3.3.3.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 3.3.3.3/32
  Known via "ospf 2", distance 110, metric 65, type intra area
  Last update from 34.1.1.3 on Serial1/3, 00:00:42 ago
  Routing Descriptor Blocks:
  * 34.1.1.3, from 1.1.1.1, 00:00:42 ago, via Serial1/3
      Route metric is 65, traffic share count is 1

R4#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 3.3.3.3.646 - 4.4.4.4.27763
        State: Oper; Msgs sent/rcvd: 8/12; Downstream
        Up time: 00:00:58
        LDP discovery sources:
          Serial1/3, Src IP addr: 34.1.1.3
        Addresses bound to peer LDP Ident:
          34.1.1.3        35.1.1.3        23.1.1.3        3.3.3.3

R4#Show mpls ldp disc

 Local LDP Identifier:
    4.4.4.4:0
    Discovery Sources:
    Interfaces:
        Serial1/3 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0
```

Let's go to R3 and verify the LDP session:

On R3:

```
R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.61990
        State: Oper; Msgs sent/rcvd: 11/9; Downstream
        Up time: 00:01:00
        LDP discovery sources:
          Serial1/2, Src IP addr: 23.1.1.2
        Addresses bound to peer LDP Ident:
          2.2.2.2         23.1.1.2
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 4.4.4.4.54984 - 3.3.3.3.646
        State: Oper; Msgs sent/rcvd: 11/9; Downstream
        Up time: 00:00:56
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
        LDP discovery sources:
          Serial1/4, Src IP addr: 34.1.1.4
        Addresses bound to peer LDP Ident:
          4.4.4.4         34.1.1.4
```

R3 has not established an LDP session with R5.

```
R3#Show mpls ldp discovery detail

 Local LDP Identifier:
    3.3.3.3:0
    Discovery Sources:
    Interfaces:
        Serial1/2 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 3.3.3.3
            LDP Id: 2.2.2.2:0
              Src IP addr: 23.1.1.2; Transport IP addr: 2.2.2.2
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 2.2.2.2/32
              Password: not required, none, in use
            Clients: IPv4
        Serial1/4 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 3.3.3.3
            LDP Id: 4.4.4.4:0
              Src IP addr: 34.1.1.4; Transport IP addr: 4.4.4.4
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 4.4.4.4/32
              Password: not required, none, in use
            Clients: IPv4
        Serial1/5 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 3.3.3.3
            LDP Id: 5.5.5.5:0; no route to transport addr
              Src IP addr: 35.1.1.5; Transport IP addr: 5.5.5.5
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Password: not required, none, in use
            Clients: IPv4
```

Let's verify the routing table of R5:

On R5:

```
R5#Show ip route 3.3.3.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
% Network not in table
```

Let's see how OSPF was configured:

```
R5#Sh run | s router ospf

router ospf 1 vrf aa
 router-id 0.0.0.5
 redistribute bgp 100 subnets
 network 57.1.1.5 0.0.0.0 area 0
router ospf 2
 network 5.5.5.5 0.0.0.0 area 0
 network 35.1.1.5 0.0.0.0 area 0
```

Let's veify the OSPF configuration on R3:

On R3:

```
R3#Show run | s router ospf

router ospf 2
 network 3.3.3.3 0.0.0.0 area 0
 network 23.1.1.3 0.0.0.0 area 0
 network 33.3.3.3 0.0.0.0 area 0
 network 34.1.1.3 0.0.0.0 area 0
```

We can see the problem, R3 is NOT running OSPF on the interface facing R5, let's run OSPF area 0 on the S1/5 interface of R3:

```
R3(config)#router ospf 2
R3(config-router)#Netw 35.1.1.3 0.0.0.0 area 0
```

You should see the following console message on R3 stating that OSPF and LDP neighbor adjacency has been established with R5:

```
%OSPF-5-ADJCHG: Process 2, Nbr 5.5.5.5 on Serial1/5 from LOADING to FULL,
Loading Done

%LDP-5-NBRCHG: LDP Neighbor 5.5.5.5:0 (3) is UP
```

To verify the configuration:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show mpls ldp neighbor 5.5.5.5

    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 5.5.5.5.17003 - 3.3.3.3.646
        State: Oper; Msgs sent/rcvd: 11/12; Downstream
        Up time: 00:01:47
        LDP discovery sources:
          Serial1/5, Src IP addr: 35.1.1.5
        Addresses bound to peer LDP Ident:
          5.5.5.5         35.1.1.5

R3#Show mpls ldp dis
 Local LDP Identifier:
    3.3.3.3:0
    Discovery Sources:
    Interfaces:
        Serial1/2 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0
        Serial1/4 (ldp): xmit/recv
            LDP Id: 4.4.4.4:0
        Serial1/5 (ldp): xmit/recv
            LDP Id: 5.5.5.5:0
```

This looks pretty good. Let's see if R5's Lo0 is configured with a /32 subnet mask:

On R5:

```
R5#Show run int lo0 | B inter

interface Loopback0
 ip address 5.5.5.5 255.255.255.0
end
```

No, it's not, let's fix this problem:

```
R5(config)#int lo0
R5(config-if)#ip addr 5.5.5.5 255.255.255.255
```

To verify the configuration:

On R5:

```
R5#Show mpls ldp neigh
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 3.3.3.3.646 - 5.5.5.5.47447
        State: Oper; Msgs sent/rcvd: 12/11; Downstream
        Up time: 00:01:03
        LDP discovery sources:
          Serial1/3, Src IP addr: 35.1.1.3
        Addresses bound to peer LDP Ident:
          23.1.1.3        34.1.1.3        35.1.1.3        3.3.3.3
```

Let's verify the configuration of R2's Lo0 interface before we test the data plane from R6 and R7 to the lo0 interface of R1:

On R2:

```
R2#Show run int lo0 | B inter

interface Loopback0
 ip address 2.2.2.2 255.255.255.255
end
```

So far they all look good, let's test the reachability from R6 and R7 to 1.1.1.1/32 prefix advertised by R1:

On R6:

```
R6#Traceroute 1.1.1.1 numeric

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 67.1.1.7 0 msec 0 msec 0 msec
  2 57.1.1.5 0 msec 4 msec 0 msec
  3 35.1.1.3 [MPLS: Labels 300/203 Exp 0] 128 msec 128 msec 128 msec
  4 12.1.1.2 [MPLS: Label 203 Exp 0] 68 msec 64 msec 64 msec
  5 12.1.1.1 52 msec *  48 msec
```

1. This is great we can see that R6 goes through R7 to reach the prefix because it has a better cost.
2. R7 goes to R5.
3. R5 attaches the VPN label of 203 and then it performs a lookup in the LFIB to see which label it needs to impose to reach 2.2.2.2 so it finds label 300 which was given to it by R3.
4. R3 is the PHP, therefore, it removes label 300 and sends the labeled packet to R2.
5. R2 receives a labeled packet of 203, it knows that it was the LSR that imposed that label, so it removes the label and sends the packet with no labels to R1.

Now, let's see the problem with 11.1.1.1 prefix, let's go to the source and work our way back to R6 and R7:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Show ip int br | exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              12.1.1.1        YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
Loopback1              11.1.1.1        YES manual up                    up

R1#Show run int lo1 | B interface
interface Loopback1
 ip address 11.1.1.1 255.255.255.0
end
```

Let's see if OSPF is running on this interface:

```
R1#Show ip ospf inter brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/24         1     LOOP  0/0
Se1/2        1     0               12.1.1.1/24        64    P2P   1/1
```

No, OSPF is not running on this interface, let's see if this interface was redistributed:

```
R1#Show run | s router ospf
router ospf 1
 router-id 0.0.0.1
 redistribute connected subnets route-map tst
 network 1.1.1.1 0.0.0.0 area 0
 network 12.1.1.1 0.0.0.0 area 0
```

Let's check the route-map one more time:

```
R1#Show run | s route-map
 redistribute connected subnets route-map tst
route-map tst permit 10
 match interface Loopback1
 set tag 100
```

The route is redistributed into OSPF with a route tag of 100. Let's see if R2 can see this prefix in the VRF:

On R2:

```
R2#Show ip route vrf aa 11.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing Table: aa
% Network not in table
```

Let's see if it's in the MP-BGP:

```
R2#Show ip bgp vpnv4 vrf aa 11.1.1.1
% Network not in table
```

No, it's not. Was it redistributed?

```
R2#Show ip vrf det

VRF aa (VRF Id = 1); default RD 1:10; default VPNID <not set>
  Interfaces:
    Se1/1
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix

R2#Show run | s address-family ipv4 vrf aa

 address-family ipv4 vrf aa
  redistribute ospf 1
```

Yes it was. So what is the problem?
Let's see the OSPF configuration for VRF "aa":

```
R2#Sh run | s router ospf 1 vrf aa

router ospf 1 vrf aa
 domain-tag 100
 redistribute bgp 100 subnets
 network 12.1.1.2 0.0.0.0 area 0
```

We can see the problem, the "domain-tag" value matches the BGP's AS number, since the ticket states that we can NOT remove any command/s to fix this problem, let's change the domain tag on R2 under "router OSPF 1 vrf aa" to another value, let's change it to "1234".

```
R2(config)#router ospf 1 vrf aa
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#domain-tag 1234
```

To verify the configuration:

On R2:

```
R2#Show ip route vrf aa | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/65] via 12.1.1.1, 00:00:48, Serial1/1
      6.0.0.0/32 is subnetted, 1 subnets
B        6.6.6.6 [200/3] via 5.5.5.5, 00:08:23
      7.0.0.0/32 is subnetted, 1 subnets
B        7.7.7.7 [200/2] via 5.5.5.5, 00:08:23
      11.0.0.0/24 is subnetted, 1 subnets
O E2     11.1.1.0 [110/20] via 12.1.1.1, 00:00:48, Serial1/1
      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, Serial1/1
L        12.1.1.2/32 is directly connected, Serial1/1
      46.0.0.0/24 is subnetted, 1 subnets
B        46.1.1.0 [200/0] via 4.4.4.4, 00:25:37
      57.0.0.0/24 is subnetted, 1 subnets
B        57.1.1.0 [200/0] via 5.5.5.5, 00:08:23
      67.0.0.0/24 is subnetted, 1 subnets
B        67.1.1.0 [200/2] via 5.5.5.5, 00:08:23
```

Now, let's see if it's in the MP-BGP:

```
R2#Show ip bgp vpnv4 vrf aa 11.1.1.1

% Network not in table
```

What else is missing? Let's see how OSPF was redistributed:

```
R2#Show run | s address-family ipv4 vrf aa

 address-family ipv4 vrf aa
  redistribute ospf 1
```

We can see the problem, when OSPF is redistributed into BGP, the route types that are to be redistributed must be configured using the "match" keyword or else ONLY the internal OSPF routes will be redistributed, let's configure this:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#router bgp 100
R2(config-router)#address-family ipv4 vrf aa
R2(config-router-af)#redistr ospf 1 match internal external 1 external 2
```

To verify the configuration:

On R2:

```
R2#Show ip bgp vpnv4 vrf aa 11.1.1.1

BGP routing table entry for 1:10:11.1.1.0/24, version 23
Paths: (1 available, best #1, table aa)
  Advertised to update-groups:
     1
  Local
    12.1.1.1 from 0.0.0.0 (2.2.2.2)
      Origin incomplete, metric 20, localpref 100, weight 32768, valid,
sourced, best
      Extended Community: RT:1:100 OSPF DOMAIN ID:0x0005:0x000000010200
        OSPF RT:0.0.0.0:5:1 OSPF ROUTER ID:12.1.1.2:0
      mpls labels in/out 207/nolabel
```

To verify the configuration:

On R6:

```
R6#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/67] via 67.1.1.7, 00:10:12, FastEthernet0/0
      7.0.0.0/32 is subnetted, 1 subnets
O        7.7.7.7 [110/2] via 67.1.1.7, 01:20:57, FastEthernet0/0
      11.0.0.0/24 is subnetted, 1 subnets
O E2     11.1.1.0 [110/20] via 67.1.1.7, 00:07:04, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/3] via 67.1.1.7, 00:10:12, FastEthernet0/0
      57.0.0.0/24 is subnetted, 1 subnets
O        57.1.1.0 [110/2] via 67.1.1.7, 01:20:57, FastEthernet0/0

R6#Traceroute 11.1.1.1 Numeric

Type escape sequence to abort.
Tracing the route to 11.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VRF info: (vrf in name/id, vrf out name/id)
  1 67.1.1.7 0 msec 0 msec 4 msec
  2 57.1.1.5 0 msec 4 msec 0 msec
  3 35.1.1.3 [MPLS: Labels 300/210 Exp 0] 128 msec 128 msec 128 msec
  4 12.1.1.2 [MPLS: Label 210 Exp 0] 64 msec 64 msec 64 msec
  5 12.1.1.1 48 msec *  48 msec
```

Perfect.

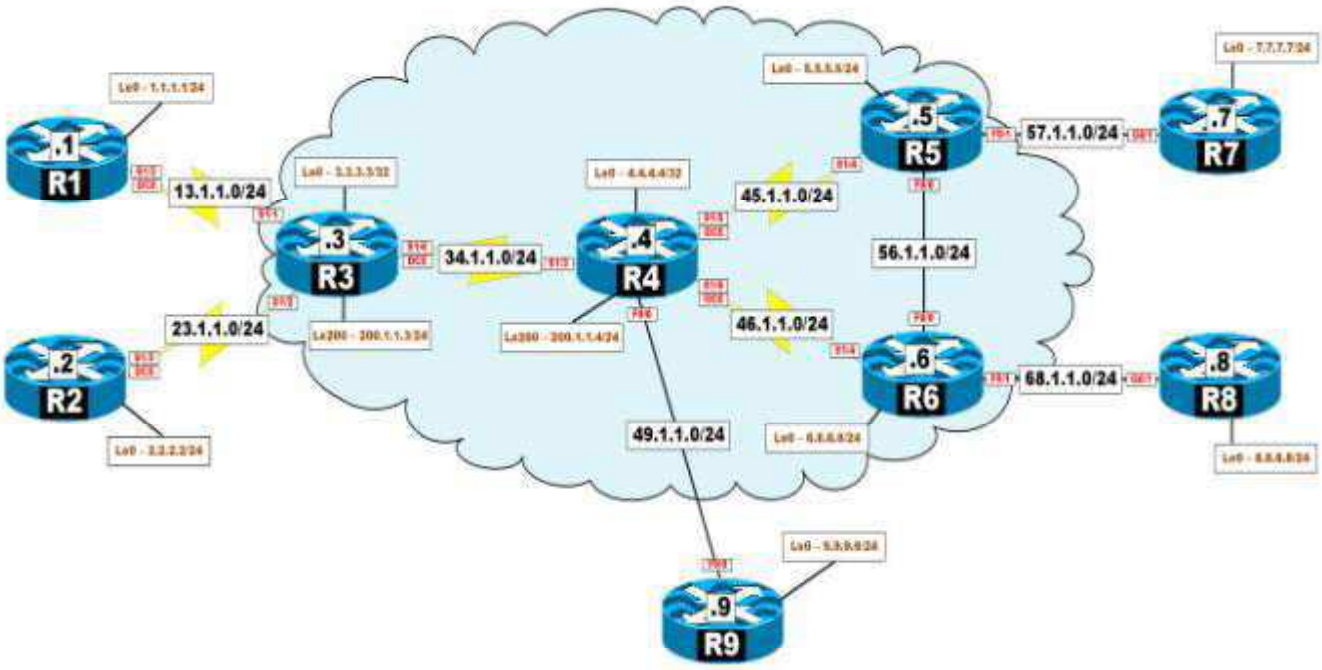
Ticket 2

Erase the startup configuration and reload the devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Troubleshooting MPLS Scenario - 2



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "MPLS" → "MPLS-TS-Lab-2".

Ticket 1

The customer routers (R1 and R9) are configured to run OSPF area 0 on their directly connected interfaces, but they see each other's prefixes as inter-area routes . DO NOT remove any command/s to fix this ticket.

Let's verify the problem:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      9.0.0.0/24 is subnetted, 1 subnets
O IA     9.9.9.0 [110/66] via 13.1.1.3, 00:02:49, Serial1/3
      49.0.0.0/24 is subnetted, 1 subnets
O IA     49.1.1.0 [110/65] via 13.1.1.3, 00:04:46, Serial1/3
```

On R9:

```
R9#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
O IA     1.1.1.0 [110/66] via 49.1.1.4, 00:00:46, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
O IA     13.1.1.0 [110/2] via 49.1.1.4, 00:00:46, FastEthernet0/0
      200.1.1.0/32 is subnetted, 1 subnets
O        200.1.1.4 [110/2] via 49.1.1.4, 00:02:35, FastEthernet0/0
```

Yes, we can see that they see each other's prefixes as "IA", let's see if the CE routers (R1 and R9) have configured OSPF area 0 on their interfaces:

On R1:

```
R1#Show ip ospf inter br

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/24         1     P2P   0/0
Se1/3        1     0               13.1.1.1/24        64    P2P   1/1
```

On R9:

```
R9#Show ip ospf int br

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               9.9.9.9/24         1     P2P   0/0
Fa0/0        1     0               49.1.1.9/24        1     BDR   1/1
```

We can see that they are running their interfaces in area 0. What can be the problem?
In MPLS environment the cloud is known as the MPLS VPN Backbone area or the "Super Backbone Area" and the "PE" routers are known as the ABRs; therefore, the prefixes from R1 traverse through the "Super Backbone" into "Backbone" area and this is the reason that they are seen as Inter Area routes.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

There are two ways to configure the "PE" routers so the routes show up as Intra Area routes, one way is the configuration of "Sham-link", and the second way is simply configuring a GRE tunnel. When a GRE tunnel is configured, we must remember the following two important points:

1. The tunnel interface MUST be advertised in the area that the customer routes are configured in.
2. The tunnel interface MUST be configured in the customers VRF.

Let's see which solution is implemented, in this topology the PEs are R3 and R4, let's see if R3 is configured for "sham-link" configuration:

On R3:

```
R3#Show run | Inc sham-link

 area 0 sham-link 200.1.1.3 200.1.1.4
```

Sure enough sham-link is configured; there are five steps in configuring sham-links:

1. Configure a loopback interface on the PE routers, in our case R3 and R4.
2. Assign an IP address with a 32 bit mask.
3. The loopback interfaces must be configured in the customer VRF.
4. Advertise the IP address of the loopback interfaces in the address-family for the customer VRFs under BGP.
5. Configure the sham-link.

```
R3#Show ip int br | exc unass

Interface       IP-Address      OK? Method Status                Protocol
Serial1/1       13.1.1.3        YES manual up                    up
Serial1/2       23.1.1.3        YES manual up                    up
Serial1/4       34.1.1.3        YES manual up                    up
Loopback0       3.3.3.3         YES manual up                    up
Loopback200     200.1.1.4       YES manual up                    up
```

We can verify steps 1 to 3 by displaying the configuration of the lo200 interface. Let's verify:

On R3:

```
R3#Show run int lo200 | B interface

interface Loopback200
 ip vrf forwarding 11
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 ip address 200.1.1.4 255.255.255.0
end
```

We can see that this interface has an incorrect mask, remember the subnet-mask of this loopback interface should be a /32. Let's fix this and verify:

```
R3(config)#int lo200
R3(config-if)#ip address 200.1.1.4 255.255.255.255
```

Let's do the same on R4:

On R4:

```
R4#Show ip int br | exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        49.1.1.4        YES manual up                    up
Serial1/3              34.1.1.4        YES manual up                    up
Serial1/5              45.1.1.4        YES manual up                    up
Serial1/6              46.1.1.4        YES manual up                    up
Loopback0              4.4.4.4         YES manual up                    up
Loopback200            200.1.1.4       YES manual up                    up

R4#Show run int lo200 | B interface

interface Loopback200
 ip vrf forwarding 99
 ip address 200.1.1.4 255.255.255.0
end
```

We can se the same problem here as well:

```
R4(config)#int lo200
R4(config-if)#ip addr 200.1.1.4 255.255.255.255
```

But we also have a duplicate IP address, based on the lab topology, the lo200 interface of R3 should be 200.1.1.3. Let's go back to R3 and change it:

On R3:

```
R3(config)#int lo200
R3(config-if)#ip address 200.1.1.3 255.255.255.255
```

WOW, you should see the following console message:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 200.1.1.4 from
200.1.1.4 on interface OSPF_SL0
```

Let's check the router-id of these two routers:

On R3:

We need to find out which OSPF process is configured for the VRF

```
R3#Show run | s router ospf

router ospf 1 vrf 11
 area 0 sham-link 200.1.1.3 200.1.1.4
 redistribute bgp 100 subnets
 network 13.1.1.3 0.0.0.0 area 0
 network 200.1.1.3 0.0.0.0 area 0
router ospf 2
 network 3.3.3.3 0.0.0.0 area 0
 network 34.1.1.3 0.0.0.0 area 0
```

Now that we know the process ID, let's verify the RID:

```
R3#Show ip ospf 1 | Inc ID

 Routing Process "ospf 1" with ID 200.1.1.4
   Domain ID type 0x0005, value 0.0.0.1
```

On R4:

```
R4#Show run | s router ospf

router ospf 1 vrf 99
 area 0 sham-link 200.1.1.4 200.1.1.3
 redistribute bgp 100 subnets
 network 49.1.1.4 0.0.0.0 area 0
router ospf 2
 network 4.4.4.4 0.0.0.0 area 0
 network 34.1.1.4 0.0.0.0 area 0
 network 45.1.1.4 0.0.0.0 area 0
 network 46.1.1.4 0.0.0.0 area 0

R4#Show ip ospf 1 | Inc ID

 Routing Process "ospf 1" with ID 200.1.1.4
   Domain ID type 0x0005, value 0.0.0.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

If the router-id is not configured statically, OSPF takes numerically highest IP address of any loopback interface, let's see if R3 has other loopback interfaces configured:

On R3:

```
R3#Show ip int br | Inc Loo

Loopback0              3.3.3.3         YES manual up                    up
Loopback200            200.1.1.3       YES manual up                    up
```

As we can see Lo200 does have the highest IP address, since we have already changed the IP address of the Lo200 interface, let's clear the OSPF process for the changes to take effect:

```
R3#Clear ip ospf 1 process
Reset OSPF process 1? [no]: y
```

To verify the configuration:

On R3:

```
R3#Show ip ospf 1 | Inc ID

 Routing Process "ospf 1" with ID 200.1.1.3
   Domain ID type 0x0005, value 0.0.0.1
```

Let's see if the sham-link is up:

```
R3#Show ip ospf sham-links

Sham Link OSPF_SL0 to address 200.1.1.4 is down
Area 0 source address 200.1.1.3
  Run as demand circuit
  DoNotAge LSA allowed. Cost of using 1 State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
```

OK, we have verifies the first three steps, let's verify step number four. In this step we need to ensure that the IP address of the Lo200 interfaces of R3 and R4 are advertised in the address-family for the customer VRFs under BGP.

On R3:

```
R3#Show run | S router bgp

router bgp 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 4.4.4.4 remote-as 100
 neighbor 4.4.4.4 update-source Loopback0
 neighbor 5.5.5.5 remote-as 100
 neighbor 5.5.5.5 update-source Loopback0
 neighbor 6.6.6.6 remote-as 100
 neighbor 6.6.6.6 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community extended
  neighbor 5.5.5.5 activate
  neighbor 5.5.5.5 send-community extended
  neighbor 6.6.6.6 activate
  neighbor 6.6.6.6 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf 11
  redistribute ospf 1 match internal external 1 external 2
 exit-address-family
 !
 address-family ipv4 vrf 22
  redistribute eigrp 100
 exit-address-family
```

It is not advertised under BGP, where is it advertised? Let's verify:

```
R3#Show run | Inc 200.1.1.3

 ip address 200.1.1.3 255.255.255.255
 area 0 sham-link 200.1.1.3 200.1.1.4
 network 200.1.1.3 0.0.0.0 area 0
```

We can see the problem, let's remove it from OSPF and advertise it in BGP under the customer's VRF:

```
R3(config)#router ospf 1 vrf 11
R3(config-router)#No network 200.1.1.3 0.0.0.0 area 0
```

Let's configure Lo200 interface under the address-family of customer's VRF in BGP:

```
R3(config)#router bgp 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-router)#address-family ipv4 vrf 11
R3(config-router-af)#Network 200.1.1.3 mask 255.255.255.255
```

Let's check R4:

On R4:

```
R4#Show run | Inc 200.1.1.4

ip address 200.1.1.4 255.255.255.255
 area 0 sham-link 200.1.1.4 200.1.1.3
 network 200.1.1.4 0.0.0.0 area 0
```

The lo200 interface is advertised in OSPF and it should be advertised in BGP, let's fix this problem:

```
R4(config)#router ospf 1 vrf 99
R4(config-router)#No network 200.1.1.4 0.0.0.0 a 0

R4(config)#router bgp 100
R4(config-router)#address-family ipv4 vrf 99
R4(config-router-af)#Network 200.1.1.4 mask 255.255.255.255
```

We should se the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 200.1.1.3 on OSPF_SL0 from LOADING to
FULL, Loading Done
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      9.0.0.0/24 is subnetted, 1 subnets
O        9.9.9.0 [110/67] via 13.1.1.3, 00:01:05, Serial1/3
      49.0.0.0/24 is subnetted, 1 subnets
O        49.1.1.0 [110/66] via 13.1.1.3, 00:01:05, Serial1/3
      200.1.1.0/32 is subnetted, 2 subnets
O E2     200.1.1.3 [110/1] via 13.1.1.3, 00:03:29, Serial1/3
O E2     200.1.1.4 [110/1] via 13.1.1.3, 00:01:20, Serial1/3
```

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
O        1.1.1.0 [110/67] via 49.1.1.4, 00:01:39, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
O        13.1.1.0 [110/66] via 49.1.1.4, 00:01:39, FastEthernet0/0
      200.1.1.0/32 is subnetted, 2 subnets
O E2     200.1.1.3 [110/1] via 49.1.1.4, 00:04:04, FastEthernet0/0
O E2     200.1.1.4 [110/1] via 49.1.1.4, 00:01:54, FastEthernet0/0

R9#Traceroute 1.1.1.1 source Lo0 numeric

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 49.1.1.4 4 msec 0 msec 4 msec
  2 13.1.1.3 [MPLS: Label 28 Exp 0] 32 msec 32 msec 32 msec
  3 13.1.1.1 32 msec *  32 msec

R9#Ping 1.1.1.1 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 9.9.9.9
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

The problem is fixed.

Ticket 2

The customer routers (R2, R7 and R8) are saying that they can not see each other's routers. DO NOT remove any commands to fix this problem.

Let's verify the problem:

On R7:

```
R7#Show ip protocol | Inc Routing Protocol
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing Protocol is "eigrp 100"
```

They are running EIGRP in AS 100, let's verify the routing table:

```
R7#Show ip route eigrp | B Gate
Gateway of last resort is not set
R7#
```

Let's verify R8's routing table:

On R8:

```
R8#Show ip route eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 68.1.1.6, 00:57:01, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 68.1.1.6, 00:57:01, FastEthernet0/1
```

R8 can see some routes but not all. Let's verify R2's routing table:

On R2:

```
R2#Show ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 100"

R2#Show ip route eigrp | B Gate
Gateway of last resort is not set

      8.0.0.0/24 is subnetted, 1 subnets
D        8.8.8.0 [90/20642560] via 23.1.1.3, 00:57:45, Serial1/3
      68.0.0.0/24 is subnetted, 1 subnets
D        68.1.1.0 [90/20514560] via 23.1.1.3, 00:58:29, Serial1/3
```

R2 can see the link between the PE (R6) and R8, and network 8.8.8.0/24. Does R2 have reachability to these routes? Let's verify:

```
R2#Ping 68.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 68.1.1.8, timeout is 2 seconds:
.....
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

R2#Ping 8.8.8.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's start by verifying connectivity from R2 to R8:

On R2:

```
R2#Show ip int br | Exc unass

Interface        IP-Address      OK? Method Status                Protocol
Serial1/3        23.1.1.2        YES manual up                    up
Loopback0        2.2.2.2         YES manual up                    up

R2#Show run int Lo0 | B interface

interface Loopback0
 ip address 2.2.2.2 255.255.255.0
end
```

Let's see if R2 is advertising its loopback 0 interface:

```
R2#Show run | s router eigrp 100

router eigrp 100
 network 2.2.2.2 0.0.0.0
 network 23.1.1.2 0.0.0.0
```

Yes it is running EIGRP on the link. Let's see if the PE to which R2 is directly connected to can see the prefix:

On R3:

```
R3#Sh ip vrf

  Name                             Default RD          Interfaces
  11                               1:10                Se1/1
                                                       Lo200
  22                               1:20                Se1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's see if R3 can see the route in its VRF:

```
R3#Show ip route vrf 22 eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2297856] via 23.1.1.2, 01:38:56, Serial1/2
```

The next step is to see if the prefix was redistributed into MB-BGP:

```
R3#Show ip bgp vpnv4 vrf 22 | B Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf 22)
*> 2.2.2.0/24       23.1.1.2           2297856         32768 ?
*>i8.8.8.0/24       6.6.6.6             156160    100      0 ?
*> 23.1.1.0/24      0.0.0.0                  0         32768 ?
*>i68.1.1.0/24      6.6.6.6                  0    100      0 ?
```

let's see which interface the local router uses as its LDP ID so we can verify the proper operation of LDP:

```
R3#Show run | Inc mpls ldp router-id

mpls ldp router-id Loopback0 force
```

Let's verify the configuration of this loopback interface:

```
R3#Show run int lo0 | B inter

interface Loopback0
 ip address 3.3.3.3 255.255.255.255
```

Next is to verify the LDP session:

```
R3#Show mpls ldp neighbor 4.4.4.4

    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 4.4.4.4.25671 - 3.3.3.3.646
        State: Oper; Msgs sent/rcvd: 128/128; Downstream
        Up time: 01:42:29
        LDP discovery sources:
          Serial1/4, Src IP addr: 34.1.1.4
        Addresses bound to peer LDP Ident:
          4.4.4.4         34.1.1.4        45.1.1.4        46.1.1.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's move to R4:

On R4:

```
R4#Show run | Inc mpls ldp router

mpls ldp router-id Loopback0 force

R4#Show run int lo0 | B interface

interface Loopback0
 ip address 4.4.4.4 255.255.255.255
end

R4#Show mpls ldp neigh 5.5.5.5

    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 5.5.5.5.13755 - 4.4.4.4.646
        State: Oper; Msgs sent/rcvd: 130/132; Downstream
        Up time: 01:44:00
        LDP discovery sources:
          Serial1/5, Src IP addr: 45.1.1.5
        Addresses bound to peer LDP Ident:
          5.5.5.5         56.1.1.5        45.1.1.5

R4#Show mpls ldp neigh 6.6.6.6

    Peer LDP Ident: 6.6.6.6:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 6.6.6.6.34662 - 4.4.4.4.646
        State: Oper; Msgs sent/rcvd: 129/131; Downstream
        Up time: 01:44:25
        LDP discovery sources:
          Serial1/6, Src IP addr: 46.1.1.6
        Addresses bound to peer LDP Ident:
          6.6.6.6         56.1.1.6        46.1.1.6
```

We should verify R5:

On R5:

```
R5#Show run | Inc mpls ldp router

mpls ldp router-id Loopback0 force

R5#Show run int lo0 | B inter
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
interface Loopback0
 ip address 5.5.5.5 255.255.255.0
end
```

We can see the problem, R5's Lo0 interface has a 24 bit mask and it should have a 32 bit mask, let's fix this problem:

On R5:

```
R5(config)#int lo0
R5(config-if)#ip addr 5.5.5.5 255.255.255.255
```

We should check R6:

On R6:

```
R6#Show run | Inc mpls ldp router

mpls ldp router-id Loopback0 force

R6#Show run int Lo0 | B inter

interface Loopback0
 ip address 6.6.6.6 255.255.255.0
end
```

We can see a similar problem on R6, R6's Lo0 interface has a 24 bit mask and it should have a 32 bit mask, let's fix this problem:

```
R6(config)#int lo0
R6(config-if)#ip addr 6.6.6.6 255.255.255.255
```

Let's see if R6 can see any prefixes form R8:

```
R6#Show ip vrf

  Name                              Default RD          Interfaces
  88                                1:80                Fa0/1
```

We can see that the VRF was configured and applied it R6's F0/1 interface.

```
R6#Show ip bgp vpnv4 vrf 88 | B Net

   Network          Next Hop            Metric LocPrf Weight Path
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Route Distinguisher: 1:80 (default for vrf 88)
*>i2.2.2.0/24        3.3.3.3                2297856    100        0 ?
*> 8.8.8.0/24        68.1.1.8                156160              32768 ?
*>i23.1.1.0/24       3.3.3.3                       0    100        0 ?
*> 68.1.1.0/24       0.0.0.0                       0              32768 ?
```

Let's check and see if the route is in the VRF:

```
R6#Show ip route vrf 88 | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
B        2.2.2.0 [200/2297856] via 3.3.3.3, 02:58:29
      8.0.0.0/24 is subnetted, 1 subnets
D        8.8.8.0 [90/156160] via 68.1.1.8, 03:19:45, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
B        23.1.1.0 [200/0] via 3.3.3.3, 02:58:29
      68.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        68.1.1.0/24 is directly connected, FastEthernet0/1
L        68.1.1.6/32 is directly connected, FastEthernet0/1
```

We can see Network 2.2.2.0/24 advertised by R2 and network 8.8.8.0/24 advertised by R8. Let's check and see if R8 can see network 2.2.2.0/24 in its routing table:

On R8:

```
R8#Show ip route eigrp | b Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 68.1.1.6, 03:00:18, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 68.1.1.6, 03:00:18, FastEthernet0/1
```

Yes we can see network 2.2.2.0/24 in R8's routing table, let's verify reachability to network 2.2.2.0/24:

```
R8#Ping 2.2.2.2 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 8.8.8.8
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/89/92 ms
```

Let's see if R5 can see the 7.7.7.0/24 prefix in its MP-BGP table:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show ip bgp vpnv4 vrf 77 | B Net

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:70 (default for vrf 77)
*>i2.2.2.0/24       3.3.3.3            2297856    100      0 ?
*>i8.8.8.0/24       6.6.6.6             156160    100      0 ?
*>i23.1.1.0/24      3.3.3.3                  0    100      0 ?
*>i68.1.1.0/24      6.6.6.6                  0    100      0 ?
```

Let's check to see if R5 can see network 7.7.7.0/24 in its VRF:

```
R5#Show ip route vrf 77 | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
B        2.2.2.0 [200/2297856] via 3.3.3.3, 03:23:20
      8.0.0.0/24 is subnetted, 1 subnets
B        8.8.8.0 [200/156160] via 6.6.6.6, 03:44:36
      23.0.0.0/24 is subnetted, 1 subnets
B        23.1.1.0 [200/0] via 3.3.3.3, 03:23:20
      57.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        57.1.1.0/24 is directly connected, FastEthernet0/1
L        57.1.1.5/32 is directly connected, FastEthernet0/1
      68.0.0.0/24 is subnetted, 1 subnets
B        68.1.1.0 [200/0] via 6.6.6.6, 03:45:12
```

Let's verify the EIGRP and BGP configuration of R5:

```
R5#Show run | s router eigrp 100

router eigrp 100
 !
 address-family ipv4 vrf 77
  redistribute bgp 100 metric 1 1 1 1 1
  network 57.1.1.5 0.0.0.0
 exit-address-family
```

We can see the problem, R5 is not configured wit hthe autonomous-system number, we should remember that the VRF does not inherit the EIGRP AS number and it must be configured statically:

```
R5(config)#router eigrp 100
R5(config-router)#address-family ipv4 vrf 77
R5(config-router-af)#autonomous-system 100
```

Remember that the "Autonomous-system" command even though a required command, in IOS 15.1(4)M7

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

behaves like a hidden command and it must be statically configured, because the TAB key does not work.

Once the above command is eneterd you should see the following console message stating that a neighbor adjacency is established between R5 and R7:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 57.1.1.7 (FastEthernet0/1) is
up: new adjacency
```

To verify the configuration:

On R5:

```
R5#Show ip route vrf 77 | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
B        2.2.2.0 [200/2297856] via 3.3.3.3, 03:31:37
      7.0.0.0/24 is subnetted, 1 subnets
D        7.7.7.0 [90/156160] via 57.1.1.7, 00:04:17, FastEthernet0/1
      8.0.0.0/24 is subnetted, 1 subnets
B        8.8.8.0 [200/156160] via 6.6.6.6, 03:52:53
      23.0.0.0/24 is subnetted, 1 subnets
B        23.1.1.0 [200/0] via 3.3.3.3, 03:31:37
      57.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        57.1.1.0/24 is directly connected, FastEthernet0/1
L        57.1.1.5/32 is directly connected, FastEthernet0/1
      68.0.0.0/24 is subnetted, 1 subnets
B        68.1.1.0 [200/0] via 6.6.6.6, 03:53:29
```

On R7:

```
R7#Show ip route eigrp | b Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 57.1.1.5, 00:05:35, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 57.1.1.5, 00:05:35, FastEthernet0/1
```

To verify reachability:

On R7:

```
R7#Ping 2.2.2.2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/90/92 ms
```

So far the routes are being exchanged between R7 and R2, and R8 and R2, but for whatever reason R7 and R8 are not exchanging routes, let's see why:

On R5:

```
R5#Show ip bgp vpnv4 vrf 77 7.7.7.0/24

BGP routing table entry for 1:70:7.7.7.0/24, version 23
Paths: (1 available, best #1, table 77)
  Advertised to update-groups:
     1
  Local
    57.1.1.7 from 0.0.0.0 (5.5.5.5)
      Origin incomplete, metric 156160, localpref 100, weight 32768, valid,
sourced, best
      Extended Community: SoO:1:555 RT:1:200 Cost:pre-bestpath:128:156160
        0x8800:32768:0 0x8801:100:130560 0x8802:65281:25600 0x8803:65281:1500
        0x8806:0:117901063
      mpls labels in/out 16/nolabel
```

We can see the extended community SoO 1:555 that is applied to this prefix, let's verify prefix 8.8.8.0/24:

```
R5#Show ip bgp vpnv4 vrf 77 8.8.8.0/24

BGP routing table entry for 1:70:8.8.8.0/24, version 13
Paths: (1 available, best #1, table 77)
  Not advertised to any peer
  Local, imported path from 1:80:8.8.8.0/24
    6.6.6.6 (metric 2) from 6.6.6.6 (6.6.6.6)
      Origin incomplete, metric 156160, localpref 100, valid, internal, best
      Extended Community: SoO:1:555 RT:1:200 Cost:pre-bestpath:128:156160
        0x8800:32768:0 0x8801:100:130560 0x8802:65281:25600 0x8803:65281:1500
        0x8806:0:134744072
      mpls labels in/out nolabel/23
```

OK we can see the problem, these routes have identical SoO value; let's see if they are applied to the F0/1 interface?

On R5:

```
R5#Show run int f0/1 | B interface
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
interface FastEthernet0/1
 ip vrf forwarding 77
 ip vrf sitemap tst
 ip address 57.1.1.5 255.255.255.0
 duplex auto
 speed auto
end

R5#Show route-map tst

route-map tst, permit, sequence 10
  Match clauses:
  Set clauses:
    extended community SoO:1:555
  Policy routing matches: 0 packets, 0 bytes
```

On R6:

```
R6#Show run int f0/1 | B interface

interface FastEthernet0/1
 ip vrf forwarding 88
 ip vrf sitemap tst
 ip address 68.1.1.6 255.255.255.0
 duplex auto
 speed auto
end

R6#Show route-map tst

route-map tst, permit, sequence 10
  Match clauses:
  Set clauses:
    extended community SoO:1:555
  Policy routing matches: 0 packets, 0 bytes
```

Let's remove the "sitemap" applied to the F0/1 interface of R5:

On R5:

```
R5(config)#route-map tst permit 10
R5(config-route-map)#No set extcommunity soo 1:555
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R7:

```
R7#Show ip route eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 57.1.1.5, 00:01:37, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 57.1.1.5, 00:01:37, FastEthernet0/1
      68.0.0.0/24 is subnetted, 1 subnets
D        68.1.1.0 [90/30720] via 57.1.1.5, 00:01:37, FastEthernet0/1
```

Why can't we see network 8.8.8.0/24? Let's check R8 and see if R8 is seeing routes from R7:

On R8:

```
R8#Show ip route eigrp | b Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 68.1.1.6, 03:10:23, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 68.1.1.6, 03:10:23, FastEthernet0/1
      57.0.0.0/24 is subnetted, 1 subnets
D        57.1.1.0 [90/30720] via 68.1.1.6, 00:03:06, FastEthernet0/1
```

Let's check the EIGRP configuration:

```
R8#Show run | s router eigrp

router eigrp 100
 network 8.8.8.8 0.0.0.0
 network 68.1.1.8 0.0.0.0
 network 78.1.1.8 0.0.0.0

R8#Show eigrp protocol

EIGRP-IPv4 Protocol for AS(100)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  NSF-aware route hold timer is 240
  Router-ID: 200.1.1.78
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    Maximum hopcount 100
    Maximum metric variance 1
```

On R7:

```
R7#Show run | S router eigrp

router eigrp 100
 network 7.7.7.7 0.0.0.0
 network 57.1.1.7 0.0.0.0
 network 78.1.1.7 0.0.0.0

R7#Show eigrp protocol

EIGRP-IPv4 Protocol for AS(100)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  NSF-aware route hold timer is 240
  Router-ID: 200.1.1.78
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 4
    Maximum hopcount 100
    Maximum metric variance 1
```

OK, we can see the problem, these two routers have identical router IDs, in the earlier versions of IOS (12.4 or ealier) router IDs ONLY effected the external routes, but in 15 code it effects both internal and external routes. Let's fix this problem:

On R7:

```
R7(config)#router eigrp 100
R7(config-router)#eigrp router-id 7.7.7.7
```

To verify the configuration:

On R7:

```
R7#Show ip route eigrp | b Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 57.1.1.5, 00:04:25, FastEthernet0/1
      8.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D        8.8.8.0 [90/158720] via 57.1.1.5, 00:04:25, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 57.1.1.5, 00:04:25, FastEthernet0/1
      68.0.0.0/24 is subnetted, 1 subnets
D        68.1.1.0 [90/30720] via 57.1.1.5, 00:04:25, FastEthernet0/1

R7#Ping 8.8.8.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R8:

```
R8#Show ip route eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/2300416] via 68.1.1.6, 00:24:10, FastEthernet0/1
      7.0.0.0/24 is subnetted, 1 subnets
D        7.7.7.0 [90/158720] via 68.1.1.6, 00:01:27, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/2172416] via 68.1.1.6, 00:24:10, FastEthernet0/1
      57.0.0.0/24 is subnetted, 1 subnets
D        57.1.1.0 [90/30720] via 68.1.1.6, 00:01:14, FastEthernet0/1

R8#Ping 7.7.7.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Ticket 3

Erase the startup configuration on all devices and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

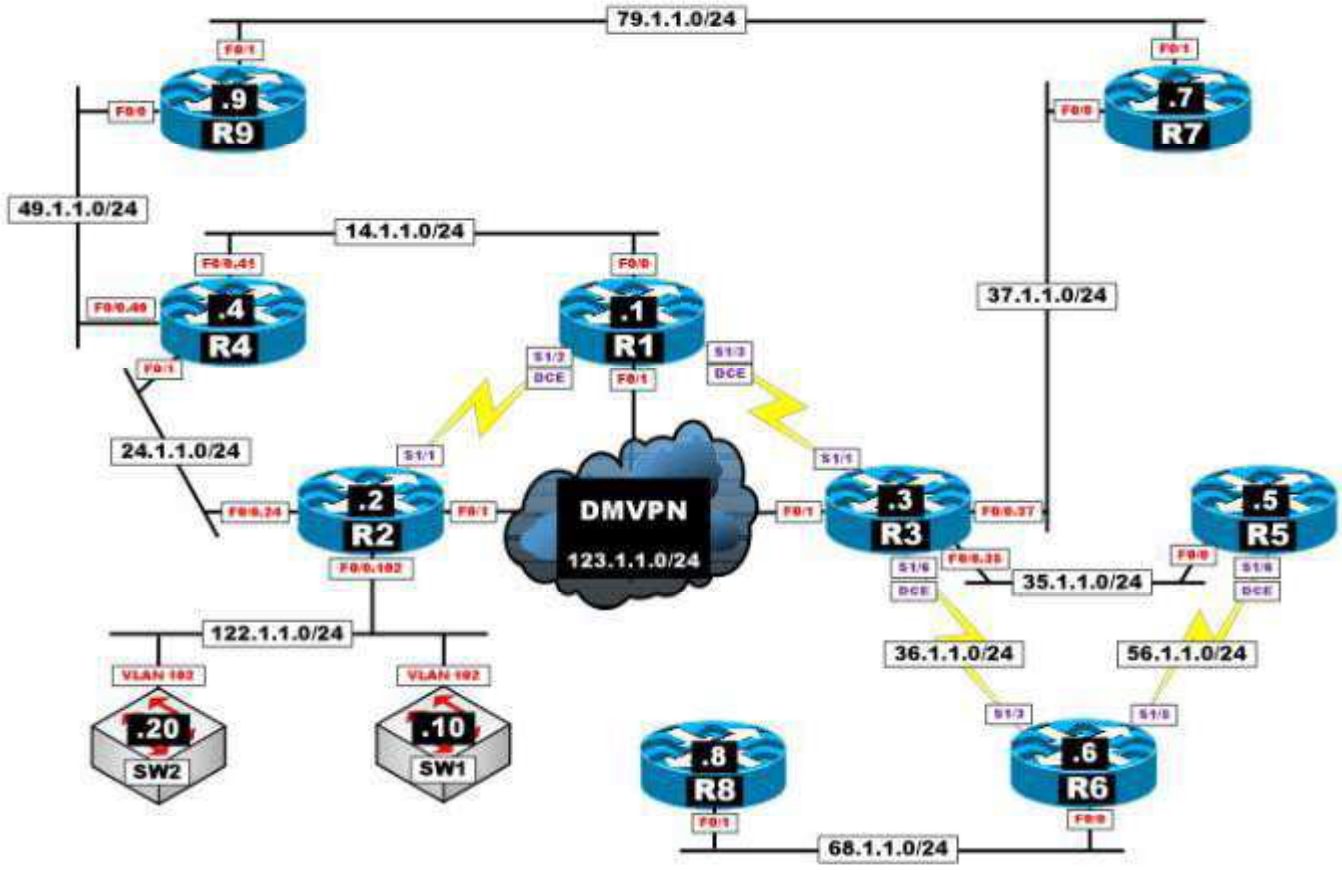
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

TS-Mock-Lab-1

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

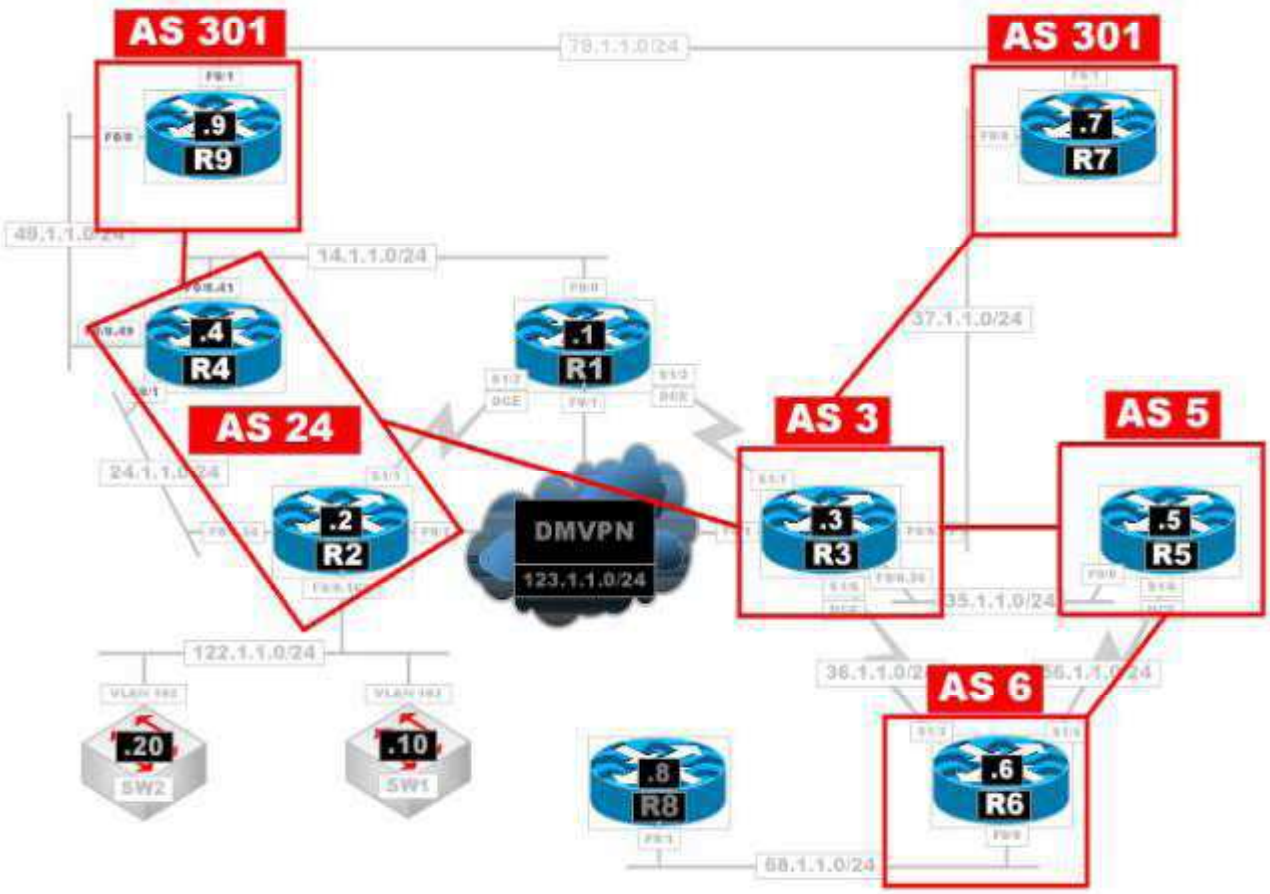
IPv4 Topology



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

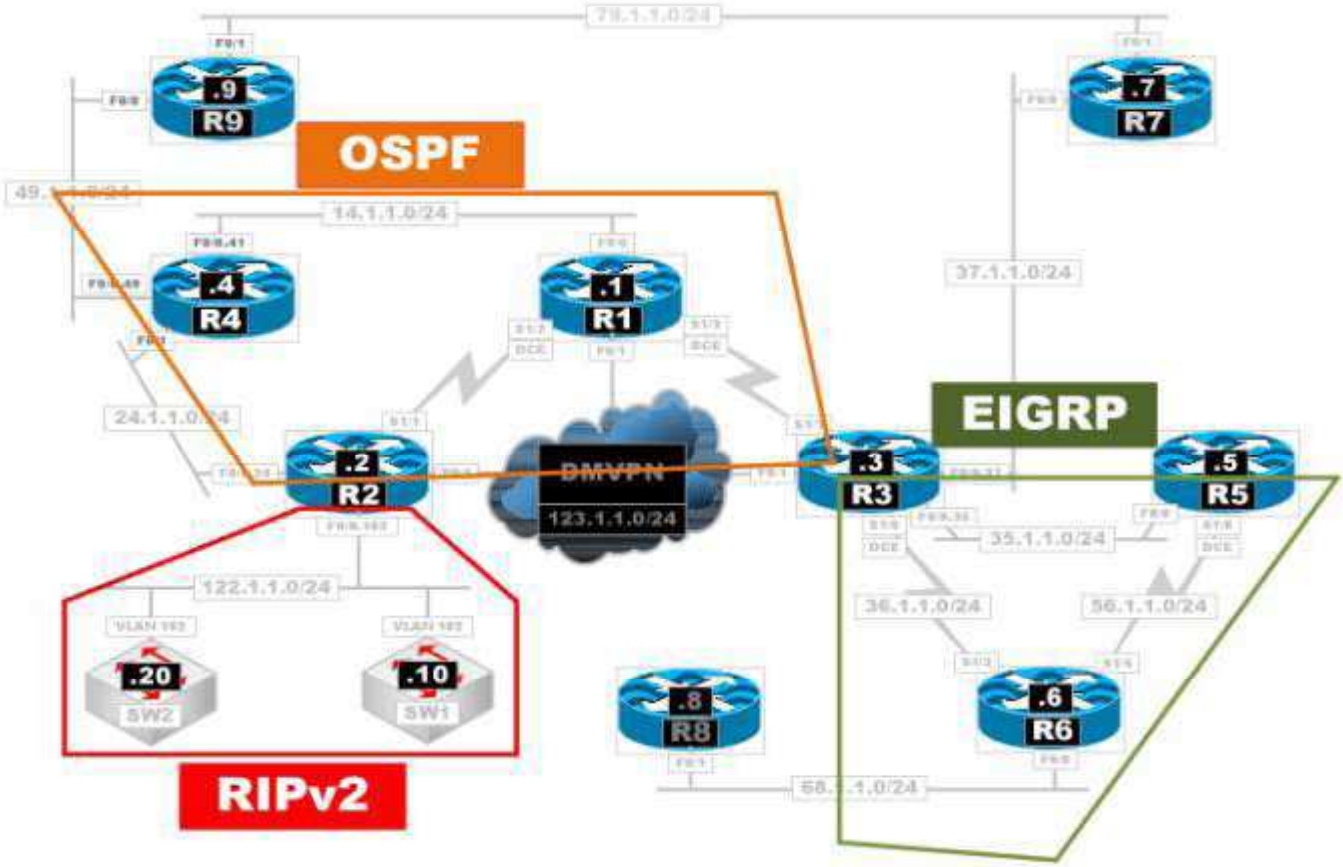
BGP Topology



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

IGP Topology



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "TS-Mock-Labs" →"TS-Mock-Lab-1".

Lab Rules:

- No static routes, no default routes, no Policy Based Routing is allowed, unless otherwise specified.
- The default password is "CCIE"

Ticket 1

The EtherChannel link between SW1 and SW2 is not working. Do not change the EtherChannel protocol or configure SW2 to fix this ticket.

To verify the problem:

On SW1:

```
SW1#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
12     Po12(SD)         PAgP      Fa0/19(I)   Fa0/20(I)
```

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------
12     Po12(SD)         PAgP      Fa0/19(I)   Fa0/20(I)
```

Looking at the output of the above show command, both switches are using PAgP, and the Etherchannel is a layer 2 Etherchannel BUT it is down (The letter "D" in the SD indicates that it is down). Let's check the configuration of these two switches:

On SW1:

```
SW1#Show run int F0/19 | B interface

interface FastEthernet0/19
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 12 mode auto
end

SW1#Show run int F0/20 | B interface

interface FastEthernet0/20
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 12 mode auto
end
```

On SW2:

```
SW2#Show run int F0/19 | B interface

interface FastEthernet0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 12 mode auto
end

SW2#Show run int F0/20 | B interface

interface FastEthernet0/20
 switchport trunk encapsulation dot1q
 switchport mode trunk
 channel-group 12 mode auto
end
```

The output of the above configuration reveals that both switches are configured in "Auto" mode, in "Auto" mode none of the switches initiate the negotiation process. To fix this problem, only one side needs to be changed:

On SW1:

```
SW1(config)#int range f0/19-20
SW1(config-if-range)#Channel-group 12 mode desirable
```

You should see the following console messages:

```
Line protocol on Interface FastEthernet0/19, changed state to down
Line protocol on Interface FastEthernet0/20, changed state to down
Line protocol on Interface FastEthernet0/19, changed state to up
Line protocol on Interface FastEthernet0/20, changed state to up

Interface Port-channel12, changed state to up
Line protocol on Interface Port-channel12, changed state to up
```

To verify the configuration:

On SW1:

```
SW1#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
12     Po12(SU)         PAgP      Fa0/19(P)   Fa0/20(P)
```

On SW2:

```
SW2#Show etherchannel summary

Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
12     Po12(SU)         PAgP      Fa0/19(P)   Fa0/20(P)
```

Ticket 2

R1 cannot establish an OSPF session with R3. DO NOT configure OSPF, or R1 to fix this problem.

On R1:

Once you go to R1, you should see the following console message stating that R1 and R3 have identical OSPF router-ids. let's check the router-id of these routers:

```
%OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 100.1.0.1 from 13.1.1.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
on interface Serial1/3

R1#Show ip ospf | Inc ID

 Routing Process "ospf 1" with ID 100.1.0.1
```

On R3:

Identical Router-ids

```
R3#Show ip ospf | Inc ID

 Routing Process "ospf 1" with ID 100.1.0.1
```

Let's check the configuration of OSPF and see if the router-id was configured statically:

On R1:

```
R1#Sh run | S router ospf

router ospf 1
 network 1.1.1.1 0.0.0.0 area 0
 network 13.1.1.1 0.0.0.0 area 0
 network 14.1.1.1 0.0.0.0 area 0
 network 123.1.1.1 0.0.0.0 area 0
 neighbor 123.1.1.3
 neighbor 123.1.1.2
```

On R3:

```
R3#Sh run | S router ospf

router ospf 1
 network 3.3.0.3 0.0.0.0 area 0
 network 13.1.1.3 0.0.0.0 area 0
 network 123.1.1.3 0.0.0.0 area 0
```

If the router-id is NOT statically configured, OSPF will take the numerically highest IP address of any loopback interface, let's check the loopback interfaces of these routers:

On R1:

```
R1#Sh ip int br | Inc Loop

Loopback0                  1.1.0.1         YES manual up                    up
Loopback12                 12.12.12.1      YES NVRAM  up                    up
Loopback100                100.1.0.1       YES manual up                    up
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3#Show ip int br | Inc Loop

Loopback0                  3.3.0.3         YES manual up                    up
Loopback1                  3.3.1.3         YES NVRAM  up                    up
Loopback100                100.1.0.1       YES manual up                    up
```

Now we see the reason for this problem. The numerically highest IP address of the loopback interfaces of these two routers is identical. We can fix this problem in two different ways, statically configure the router-ids, or fix the IP address of one of the routers.

Fixing the IP address is better than just configuring the router-id:

On R3:

```
R3(config)#int lo100
R3(config-if)#ip addr 100.1.0.3 255.255.255.0
```

We must reset the OSPF process to inmplement this change:

On R3:

```
R3#Clear ip ospf process

Reset ALL OSPF processes? [no]: Yes
```

To verify the configuration:

On R3:

```
R3#Show ip ospf | Inc ID

 Routing Process "ospf 1" with ID 100.1.0.3
```

The outputs of the following show commands with the exception of R1 and R3 may be different, this could be a flapping neighbor; you should NOT fix any problem/s unless stated.

```
R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
100.1.0.1         1   FULL/DR         00:00:35    123.1.1.1       Tunnel123
100.1.0.1         0   FULL/  -        00:00:36    13.1.1.1        Serial1/1
```

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/DROTHER    00:00:39    123.1.1.2       Tunnel123
100.1.0.3         0   FULL/DROTHER    00:00:33    123.1.1.3       Tunnel123
4.4.4.4           1   FULL/BDR        00:00:31    14.1.1.4        FastEthernet0/0
100.1.0.3         0   FULL/  -        00:00:33    13.1.1.3        Serial1/3
```

Ticket 3

For some reason R1 is not exchanging routes with R2 through its Tunnel 123 interface.
DO NOT configure R1 to fix this problem.

Let's see if the two routers are running OSPF on the correct interface:

On R1:

```
R1#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/24         1     LOOP  0/0
Tu123        1     0               123.1.1.1/24       64    DR    2/2
Fa0/0        1     0               14.1.1.1/24        1     DR    1/1
Se1/3        1     0               13.1.1.1/24        781   P2P   1/1
```

On R2:

```
R2#Show ip ospf int brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Tu123        1     0               123.1.1.2/24       1000  P2P   1/1
Lo0          1     24              2.2.2.2/32         1     LOOP  0/0
Tu1          1     24              Unnumbered Lo0     1000  P2P   0/0
Tu0          1     24              Unnumbered Lo0     1000  P2P   0/0
Fa0/0.24     1     24              24.1.1.2/24        1     BDR   1/1
```

On R1:

```
R1#Show ip ospf neighbor
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           0   FULL/DROTHER    00:00:36    123.1.1.2       Tunnel123
100.1.0.3         0   FULL/DROTHER    00:00:32    123.1.1.3       Tunnel123
4.4.4.4           1   FULL/BDR        00:00:39    14.1.1.4        FastEthernet0/0
100.1.0.3         0   FULL/  -        00:00:34    13.1.1.3        Serial1/3
```

Note the output of the above show command clearly states that the two routers have established a neighbor adjacency. Let's verify this on R2:

On R2:

```
R2#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
100.1.0.1         0   FULL/  -        00:00:36    123.1.1.1       Tunnel123
4.4.4.4           1   FULL/DR         00:00:36    24.1.1.4        FastEthernet0/0.24
```

We can see that R2 is in FULL state with R1 over the tunnel 123 interface, let's verify the routing table of R1:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/3] via 14.1.1.4, 00:19:46, FastEthernet0/0
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.0.3 [110/65] via 123.1.1.3, 00:11:52, Tunnel123
      4.0.0.0/32 is subnetted, 1 subnets
O IA     4.4.4.4 [110/2] via 14.1.1.4, 00:19:46, FastEthernet0/0
      24.0.0.0/24 is subnetted, 1 subnets
O IA     24.1.1.0 [110/2] via 14.1.1.4, 00:19:46, FastEthernet0/0
```

Let's shutdown the F0/0 interface on R1 and see if 2.2.2.2/32 prefix can be learned through the DMVPN cloud:

```
R1(config)#int F0/0
R1(config-if)#Shut
```

Let's check the routing table:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.0.3 [110/65] via 123.1.1.3, 00:13:44, Tunnel123
```

The local router is NOT learning any prefixes from R2. Let's see if we have unicast reachability to R2:

```
R1#Ping 123.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Let's check the parameters of the tunnel 123 interface of these two routers, we know that if everything except the OSPF Network type matches, the two routers will transition into "FULL" state but they won't exchange routes, let's verify this information:

On R2:

```
R2#Show ip ospf inter tunne 123 | Inc Network|Hello

  Internet Address 123.1.1.2/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 2.2.2.2, Network Type POINT_TO_POINT, Cost: 1000
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:09
```

On R1:

```
R1#Show ip ospf inter tunne 123 | Inc Network|Hello

  Internet Address 123.1.1.1/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 100.1.0.1, Network Type BROADCAST, Cost: 64
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
```

We can see the problem, R1's tunnel 123 interface is configured as "OSPF broadcast" network type whereas, R2's tunnel 123 interface is configured as "OSPF Point-to-point". Since the ticket states that we should NOT configure R1 to fix this ticket, let's change R2's OSPF's network type to Broadcast:

On R2:

```
R2(config)#Int tunnel 123
R2(config-if)#IP ospf network broadcast
```

We should see the following console messages stating that the local router has established an OSPF adjacency with R1 through the tunnel interface, let's see if the routers are exchanging prefixes:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/65] via 123.1.1.2, 00:01:51, Tunnel123
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.0.3 [110/65] via 123.1.1.3, 00:23:00, Tunnel123
      4.0.0.0/32 is subnetted, 1 subnets
O IA     4.4.4.4 [110/66] via 123.1.1.2, 00:01:51, Tunnel123
      24.0.0.0/24 is subnetted, 1 subnets
O IA     24.1.1.0 [110/65] via 123.1.1.2, 00:01:51, Tunnel123
```

Sure enough, they are exchanging routes through the tunnel 123 interface, let's "No Shut" the F0/0 interface of R1:

On R1:

```
R1(config)#int F0/0
R1(config-if)#No shut
```

Ticket 4

R5 is using its Serial link to R6 to reach R3's Lo0 interface, while it has a direct Ethernet link to R3.

Let's check R5's routing entry for R3's Loopback0, but before we do that let's find out the IP address of R3's Lo0 interface:

On R3:

```
R3#Show ip int brief Lo0

Interface              IP-Address      OK? Method Status                Protocol
Loopback0              3.3.0.3         YES manual up                    up
```

On R5:

```
R5#Show ip route 3.3.0.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 3.3.0.0/24
  Known via "eigrp 16", distance 90, metric 21152000, type internal
  Redistributing via eigrp 16
  Last update from 56.1.1.6 on Serial1/6, 00:00:37 ago
  Routing Descriptor Blocks:
  * 56.1.1.6, from 56.1.1.6, 00:00:37 ago, via Serial1/6
      Route metric is 21152000, traffic share count is 1
      Total delay is 45000 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

Let's see why R5 is NOT taking the direct Ethernet connection to reach the 3.3.0.3 IP address. Let's check if R5 sees R3 as a neighbor:

On R5:

```
R5#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(16)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   56.1.1.6                Se1/6            10 00:13:41   41   246  0  29
```

R3 is NOT a neighbor. Let's check and see if R5 is running EIGRP on its F0/0 interface:

```
R5#Show ip eigrp interfaces

EIGRP-IPv4 Interfaces for AS(16)
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0              0        0/0        0        0/1           0           0
Se1/6              1        0/0       33        5/190       310          0
Lo0                0        0/0        0        0/1           0           0
```

Even though R5 is running EIGRP on its F0/0 interface, we do not see any peers through that interface. Let's test reachability by pinging R3's F0/0.35 sub-interface from R5:

On R5:

```
R5#Ping 35.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 35.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check R5's ARP table:

On R5:

```
R5#Show ip arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  35.1.1.3               0   Incomplete      ARPA
Internet  35.1.1.5               -   0007.0e1f.bf60  ARPA   FastEthernet0/0
```

It seems like R5 can't resolve the MAC address of R3's F0/0.35 sub-interface. Let's check SW1's configuration and see if these routers are both in the same VLAN:

On SW1:

```
SW1#Show run int f0/3 | B interface

interface FastEthernet0/3
 switchport trunk encapsulation dot1q
 switchport mode trunk
end

SW1#Show run int f0/5 | B interface

interface FastEthernet0/5
 switchport access vlan 35
 switchport mode dynamic desirable
end
```

This looks fine, before we check the Spanning-tree status for these interfaces let's verify R3's configuration:

On R3:

```
R3#Show run | S FastEthernet0/0

interface FastEthernet0/0
 no ip address
 duplex auto
 speed auto

interface FastEthernet0/0.31
 encapsulation dot1Q 31
 ip address 37.1.1.3 255.255.255.0
interface FastEthernet0/0.35
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
encapsulation dot1Q 35
ip address 35.1.1.3 255.255.255.0
ipv6 enable
ipv6 rip RIPNG enable
```

R3's F0/0 interface is configured with two sub-interfaces, and the F0/0.35 sub-interface is configured correctly in the VLAN, let's verify STP:

On SW1:

```
SW1#Show spanning-tree interface f0/3

Vlan                Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
VLAN0001            Desg FWD 19        128.5    P2p
VLAN0014            Desg FWD 19        128.5    P2p
VLAN0024            Desg FWD 19        128.5    P2p
VLAN0031            Desg FWD 19        128.5    P2p
VLAN0035            Desg FWD 19        128.5    P2p
VLAN0043            Desg FWD 19        128.5    P2p
VLAN0062            Desg FWD 19        128.5    P2p
VLAN0102            Desg FWD 19        128.5    P2p
VLAN0301            Desg FWD 19        128.5    P2p

SW1#Show spanning-tree interface f0/5

no spanning tree info available for FastEthernet0/5
```

Let's check the status of F0/5:

```
SW1#Show interface F0/5 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/5                        monitoring   35         a-full  a-100 10/100BaseTX
```

Why is F0/5 interface on SW1 in monitoring status? Let's verify:

```
SW1#Show monitor session all

Session 1
---------
Type                   : Local Session
Source Ports           :
    Both               : Fa0/16
Destination Ports      : Fa0/5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    Encapsulation      : Native
          Ingress      : Disabled
```

Well, we can see that the F0/5 interface of SW1 is used as the destination of the monitor session. Let's remove this configuration so R3 and R5 can form an OSPF neighbor adjacency:

On SW1:

```
SW1(config)#No monitor session 1
```

Let's check the spanning-tree's status:

```
SW1#Sh spanning-tree int f0/5

Vlan                Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------
VLAN0035            Desg LRN 19        128.7    P2p

SW1#Show spanning-tree interface F0/5

Vlan                   Role Sts Cost      Prio.Nbr Type
---------------------- ---- --- --------- -------- -----
VLAN0035               Desg FWD 19        128.7    P2p
```

Let's go to R5 and verify the configuration:

On R5:

Once you go to R5, you should see the following console message within the 10 seconds:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 16: Neighbor 35.1.1.3 (FastEthernet0/0) is
up: new adjacency

%BGP-5-ADJCHANGE: neighbor 35.1.1.3 Up
```

It looks like we fixed couple of problems.

To verify the configuration:

On R5:

```
R5#Show ip route 3.3.0.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 3.3.0.0/24
  Known via "eigrp 16", distance 90, metric 156160, type internal
  Redistributing via eigrp 16
  Last update from 35.1.1.3 on FastEthernet0/0, 00:03:23 ago
  Routing Descriptor Blocks:
  * 35.1.1.3, from 35.1.1.3, 00:03:23 ago, via FastEthernet0/0
      Route metric is 156160, traffic share count is 1
      Total delay is 5100 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Ticket 5

R4 can't reach the lo0 interfaces of SW1 or SW2. You should fix this problem while making sure that R4 sees both routes in its routing table. DO NOT use the "IP summary-address" or the "Summary-address" commands to accomplish this task.

Before verifying the problem, let's check the lo0 interface of SW1 and SW2:

On SW1:

```
SW1#Show run int lo0 | B interface

interface Loopback0
 ip address 10.1.1.10 255.255.255.255
end
```

On SW2:

```
SW2#Show run int lo0 | B interface

interface Loopback0
 ip address 10.1.1.20 255.255.255.255
end
```

Let's see the routing entry on R4:

On R4:

```
R4#Show ip route 10.1.1.10
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
% Network not in table

R4#Show ip route 10.1.1.20

% Network not in table
```

Since R2 is the next device in the path toward SW1 and SW2, let's check the routing table of R2:

On R2:

```
R2#Show ip route 10.1.1.10

Routing entry for 10.0.0.0/8
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 122.1.1.10 on FastEthernet0/0.102, 00:00:03 ago
  Routing Descriptor Blocks:
  * 122.1.1.20, from 122.1.1.20, 00:00:17 ago, via FastEthernet0/0.102
      Route metric is 1, traffic share count is 1
    122.1.1.10, from 122.1.1.10, 00:00:03 ago, via FastEthernet0/0.102
      Route metric is 1, traffic share count is 1

R2#Show ip route 10.1.1.20

Routing entry for 10.0.0.0/8
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 122.1.1.10 on FastEthernet0/0.102, 00:00:12 ago
  Routing Descriptor Blocks:
  * 122.1.1.20, from 122.1.1.20, 00:00:25 ago, via FastEthernet0/0.102
      Route metric is 1, traffic share count is 1
    122.1.1.10, from 122.1.1.10, 00:00:12 ago, via FastEthernet0/0.102
      Route metric is 1, traffic share count is 1
```

We have two routes in R2's routing table, but R4 does NOT have them. We can see that these routes are learned via RIP, BUT based on the IGP topology, we see that RIP is NOT active on the link between R2 and R4. Therefore, a mutual redistribution needs to be configured between RIP and OSPF on R2:

On R2:

```
R2(config)#router ospf 1
R2(config-router)#redistribute rip subnets

R2(config)#router rip
R2(config-router)#redistribute ospf 1 metric 2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Now, let's check the routing table of R4 for the two prefixes:

On R4:

```
R4#Show ip route 10.1.1.10

Routing entry for 10.0.0.0/8
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 24.1.1.2 on FastEthernet0/1, 00:00:24 ago
  Routing Descriptor Blocks:
  * 24.1.1.2, from 2.2.2.2, 00:00:24 ago, via FastEthernet0/1
      Route metric is 20, traffic share count is 1

R4#Show ip route 10.1.1.20

Routing entry for 10.0.0.0/8
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 24.1.1.2 on FastEthernet0/1, 00:00:30 ago
  Routing Descriptor Blocks:
  * 24.1.1.2, from 2.2.2.2, 00:00:30 ago, via FastEthernet0/1
      Route metric is 20, traffic share count is 1
```

Let's verify the configuration by using "Traceroute" on R4:

On R4

```
R4#Traceroute 10.1.1.10

Type escape sequence to abort.
Tracing the route to 10.1.1.10
VRF info: (vrf in name/id, vrf out name/id)
  1 24.1.1.2 0 msec 0 msec 0 msec
  2 122.1.1.10 4 msec *  0 msec

R4#Traceroute 10.1.1.20

Type escape sequence to abort.
Tracing the route to 10.1.1.20
VRF info: (vrf in name/id, vrf out name/id)
  1 24.1.1.2 4 msec 0 msec 4 msec
  2 122.1.1.20 0 msec *  0 msec

R4#Show ip route ospf | Inc E2

       E1 - OSPF external type 1, E2 - OSPF external type 2
O E2  10.0.0.0/8 [110/20] via 24.1.1.2, 00:03:20, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O E2     122.1.1.0 [110/20] via 24.1.1.2, 00:03:20, FastEthernet0/1
O E2  192.1.2.0/24 [110/20] via 24.1.1.2, 00:03:20, FastEthernet0/1
```

We can see network 10.0.0.0/8 in the routing table of R4 as an OSPF route, but we need to see both routes in our routing table. Let's check RIP's configuration on SW1 and SW2.

On SW1:

```
SW1#Show run | B router rip

router rip
 version 2
 network 10.0.0.0
 network 122.0.0.0
(The rest of the output is omitted for brevity)
```

On SW2:

```
SW2#Show run | B router rip

router rip
 version 2
 network 10.0.0.0
 network 122.0.0.0
(The rest of the output is omitted for brevity)
```

Let's disable auto summarization, this could be the problem:

On SW1 and SW2:

```
SWx(config)#router rip
SWx(config-router)#No auto-summary
```

To verify the configuration:

On R2:

```
R2#Show ip route 10.1.1.10

Routing entry for 10.1.1.10/32
  Known via "rip", distance 120, metric 1
  Redistributing via ospf 1, rip
  Advertised by ospf 1 subnets
  Last update from 122.1.1.10 on FastEthernet0/0.102, 00:00:03 ago
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Routing Descriptor Blocks:
  * 122.1.1.10, from 122.1.1.10, 00:00:03 ago, via FastEthernet0/0.102
      Route metric is 1, traffic share count is 1

R2#Show ip route 10.1.1.20

Routing entry for 10.1.1.20/32
  Known via "rip", distance 120, metric 1
  Redistributing via ospf 1, rip
  Advertised by ospf 1 subnets
  Last update from 122.1.1.20 on FastEthernet0/0.102, 00:00:26 ago
  Routing Descriptor Blocks:
  * 122.1.1.20, from 122.1.1.20, 00:00:26 ago, via FastEthernet0/0.102
      Route metric is 1, traffic share count is 1
```

Let's verify the routing table of R4:

On R4:

```
R4#Show ip route ospf | inc E2

        E1 - OSPF external type 1, E2 - OSPF external type 2
O E2     10.0.0.0/8 [110/20] via 24.1.1.2, 00:10:21, FastEthernet0/1
O E2     10.1.1.10/32 [110/20] via 24.1.1.2, 00:00:11, FastEthernet0/1
O E2     10.1.1.20/32 [110/20] via 24.1.1.2, 00:02:28, FastEthernet0/1
O E2     122.1.1.0 [110/20] via 24.1.1.2, 00:10:21, FastEthernet0/1
O E2  192.1.2.0/24 [110/20] via 24.1.1.2, 00:10:21, FastEthernet0/1
```

Ticket 6

R6 should form an EIGRP neighbor adjacency with R8 but something prevents it from working.

Let's check reachability by Pinging R8 from R6:

On R6:

```
R6#Ping 68.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 68.1.1.8, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Let's check and verify if R6 is running EIGRP on the interface facing R8 router:

```
R6#Show ip eigrp interfaces

EIGRP-IPv4 Interfaces for AS(16)
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0              0        0/0       0        0/1           0           0
Se1/3              1        0/0      25        5/190       298          0
Se1/5              1        0/0      24        5/190       314          0
Lo0                0        0/0       0        0/1           0           0
Lo63               0        0/0       0        0/1           0           0
```

We can see that R6 is running EIGRP on its F0/0 interface. Let's check R8:

On R8:

```
R8#Show ip eigrp inter

EIGRP-IPv4 Interfaces for AS(16)
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Lo0                0        0/0       0        0/1           0           0
Fa0/1              0        0/0       0        0/1           0           0
```

Let's verify the adjacency on R6:

On R6:

```
R6#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(16)

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   56.1.1.5                Se1/5            11 01:10:01   22   1140  0  21
0   36.1.1.3                Se1/3            12 02:07:29   20   1140  0  22
```

Let's enable debugging of EIGRP for neighbor 68.1.1.8 on R6:

```
R6#Debug ip eigrp neighbor 68.1.1.8
Neighbor target enabled on AS 16 for 68.1.1.8
```

Since we do not see any output, we should enable conditional debugging on the F0/0 interface of R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Undebug All
All possible debugging has been turned off

R6#Debug condition interface F0/0
Condition 1 set

R6#Debug ip packet det

IP packet debugging is on (detailed)

IP: s=68.1.1.6 (local), d=224.0.0.10 (FastEthernet0/0), len 100, sending
broad/multicast, proto=88

IP: s=68.1.1.6 (local), d=224.0.0.10 (FastEthernet0/0), len 100, sending
full packet, proto=88

IP: s=68.1.1.8 (FastEthernet0/0), d=224.0.0.10, len 100, rcvd 2, proto=88

IP: s=68.1.1.8 (FastEthernet0/0), d=224.0.0.10, len 100, stop process pak
for forus packet, proto=88
```

We can see that the two routers are sending EIGRP Hellos. Let's disable the debug and check the configuration of R6's F0/0 interface:

```
R6#Undebug All
All possible debugging has been turned off

R6#Show run int F0/0 | B interface

interface FastEthernet0/0
 ip address 4.9.62.6 255.255.255.0
 ip authentication mode eigrp 16 md5
 ip authentication key-chain eigrp 16 TST
 duplex auto
 speed auto
```

Let's check the configuration of the key-chain called "TST":

```
R6#Show key chain TST
```

We can see the space.

```
Key-chain TST:
    key 1 -- text "CCIE "
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
```

Let's check the F0/1 configuration of R8:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R8:

```
R8#Show run int F0/1 | B interface

interface FastEthernet0/1
 ip address 4.9.62.22 255.255.255.0
 ip authentication mode eigrp 16 md5
 ip authentication key-chain eigrp 16 TST
 duplex auto
 speed auto
```

Next, we should check the configuration of the "TST" key chain on R8:

```
R8#Sh key chain TST

Key-chain TST:
    key 1 -- text "CCIE"
        accept lifetime (always valid) - (always valid) [valid now]
        send lifetime (always valid) - (always valid) [valid now]
```

If you can NOT see the space after the "CCIE" configured on R6, the space can also be verified using the following Show command:

On R8:

```
R8#Show run | Inc CCIE$

   key-string CCIE
```

On R6:

```
R6#Show run | Inc CCIE$
R6#
```

The above Show command will ONLY work if the "Service password-encryption" command is NOT configured. To fix the problem, the space is removed:

On R6:

```
R6(config)#key chain TST
R6(config-keychain)#key 1
R6(config-keychain-key)#key-string CCIE
```

Once the above key chain is configured with the correct password, the following console messages are displayed stating the new adjacency:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 16: Neighbor 68.1.1.8 (FastEthernet0/0) is up: new
adjacency
```

Ticket 7

After correcting all the OSPF neighboring issues, users connected to R3 complain that they can NOT communicate with the Lo0 interface of R2. Fix the problem without changing any OSPF settings.

Let's verify the problem:

On R3:

```
R3#Ping 2.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see if R3 has 2.2.2.2 prefix in its routing table:

On R3:

```
R3#Show ip route 2.2.2.2

Routing entry for 2.2.2.2/32
  Known via "ospf 1", distance 110, metric 65, type inter area
  Last update from 13.1.1.1 on Serial1/1, 00:14:38 ago
  Routing Descriptor Blocks:
  * 123.1.1.2, from 2.2.2.2, 00:14:43 ago, via Tunnel123
      Route metric is 65, traffic share count is 1
    13.1.1.1, from 4.4.4.4, 00:14:38 ago, via Serial1/1
      Route metric is 65, traffic share count is 1
```

There are two ways to connect to 2.2.2.2, but for some strange reason the local router can not reach that prefix. Let's debug and analyze the output:

On R3:

Before writing any access-list, a "**Show access-list**" must be done to ensure that the access-list does NOT

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

exist.

```
R3#Sh access-list
R3#

R3(config)#access-list 100 permit icmp any any

R3#Debug ip packet detail 100
IP packet debugging is on (detailed) for access list 100

R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:

IP: s=123.1.1.3 (local), d=2.2.2.2 (Tunnel123), len 100, sending
     ICMP type=8, code=0
IP: s=123.1.1.3 (local), d=2.2.2.2 (Tunnel123), len 100, encapsulation
failed
     ICMP type=8, code=0

IP: s=123.1.1.3 (local), d=2.2.2.2 (Serial1/1), len 100, sending
     ICMP type=8, code=0
IP: s=123.1.1.3 (local), d=2.2.2.2 (Serial1/1), len 100, sending full
packet
     ICMP type=8, code=0
```

Let's disable the debug:

```
R3#Undebug All
All possible debugging has been turned off
```

NOTE: Some of the packets were routed through S1/1 interface, and some through the tunnel 123 interface, but the packets that were going to the tunnel 123 failed with encapsulation failure.

Since there are two paths to that prefix and the Ping failed through the tunnel 123 interface, we should check the NHRP mapping on R3 and see if it is correctly configured, but let's check our routing table before we check the tunnel configuration:

```
R3#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/63] via 13.1.1.1, 00:28:08, Serial1/1
      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/65] via 123.1.1.2, 00:24:16, Tunnel123
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                [110/65] via 13.1.1.1, 00:24:11, Serial1/1
      4.0.0.0/32 is subnetted, 1 subnets
O IA     4.4.4.4 [110/64] via 13.1.1.1, 00:24:11, Serial1/1
      10.0.0.0/32 is subnetted, 2 subnets
O E2     10.1.1.10 [110/20] via 123.1.1.2, 00:14:34, Tunnel123
O E2     10.1.1.20 [110/20] via 123.1.1.2, 00:18:14, Tunnel123
      14.0.0.0/24 is subnetted, 1 subnets
O        14.1.1.0 [110/63] via 13.1.1.1, 00:24:16, Serial1/1
      24.0.0.0/24 is subnetted, 1 subnets
O IA     24.1.1.0 [110/64] via 13.1.1.1, 00:24:11, Serial1/1
      122.0.0.0/24 is subnetted, 1 subnets
O E2     122.1.1.0 [110/20] via 123.1.1.2, 00:21:28, Tunnel123
      192.1.2.0/24 is variably subnetted, 2 subnets, 2 masks
O E2     192.1.2.0/24 [110/20] via 123.1.1.2, 00:21:28, Tunnel123
```

So the next-hop to reach 2.2.2.2 prefix through the tunnel 123 interface is 123.1.1.2, so let's see if the local router has an NHRP map for that next-hop:

```
R3#Show ip nhrp detail

R3#sh ip nhrp
123.1.1.1/32 via 123.1.1.1
   Tunnel123 created 00:17:52, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
123.1.1.2/32
   Tunnel123 created 00:02:17, expire 00:00:47
   Type: incomplete, Flags: negative
   Cache hits: 7
```

The NHRP mapping for 123.1.1.2 is missing, let's add that mapping and verify the result, remember we also need to map Multicast so OSPF can for an adjacency:

On R3:

```
R3(config)#Int tunnel 123
R3(config-if)#ip nhrp map 123.1.1.2 192.1.2.2
R3(config-if)#ip nhrp map multicast 192.1.2.2
```

To verify the configuration:

On R3:

```
R3#Show ip nhrp detail
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
123.1.1.1/32 via 123.1.1.1
   Tunnel123 created 00:40:26, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
123.1.1.2/32 via 123.1.1.2
   Tunnel123 created 00:01:56, never expire
   Type: static, Flags: used
   NBMA address: 192.1.2.2
```

To test the configuration:

On R3:

```
R3#Ping 123.1.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

We can't ping the next-hop IP address of 123.1.1.2, let's see if R2 has a mapping back to R3 for return traffic:

On R2:

```
R2#Show ip nhrp detail

123.1.1.1/32 via 123.1.1.1
   Tunnel123 created 00:42:44, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
123.1.1.3/32
   Tunnel123 created 00:02:37, expire 00:00:27
   Type: incomplete, Flags: negative
   Cache hits: 7
```

The mapping is incomplete, let's statically add an NHRP mapping for 123.1.1.3:

```
R2(config)#Int tunnel 123
R2(config-if)#ip nhrp map 123.1.1.3 192.1.3.3
R2(config-if)#ip nhrp map multicast 192.1.3.3
```

To verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Show ip nhrp detail

123.1.1.1/32 via 123.1.1.1
   Tunnel123 created 00:44:24, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
123.1.1.3/32 via 123.1.1.3
   Tunnel123 created 00:00:38, never expire
   Type: static, Flags:
   NBMA address: 192.1.3.3
```

To test the configuration:

On R3:

```
R3#Ping 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/8/20 ms
```

Ticket 8

R3 and R4 are configured to establish an EBGP session, but for some unknown reason they cannot. Please refer to the BGP diagram.

Let's check the configuration of R3 and R4:

On R3:

```
R3#Sh ip bgp summ | B Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4    24       0       0        1    0    0 never    Idle
35.1.1.5        4     5      41      42        7    0    0 00:33:19        1
37.1.1.7        4   301      50      50        7    0    0 00:42:15        4
```

Let's check the configuration of R3 and R4:

```
R3#Show run | S router bgp
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
bgp log-neighbor-changes
 network 3.3.1.0 mask 255.255.255.0
 neighbor 4.4.4.4 remote-as 24
 neighbor 4.4.4.4 update-source Loopback0
 neighbor 35.1.1.5 remote-as 5
 neighbor 37.1.1.7 remote-as 301
```

On R4:

```
R4#Show run | S router bgp

router bgp 24
 bgp log-neighbor-changes
 neighbor 3.3.0.3 remote-as 3
 neighbor 3.3.0.3 update-source Loopback0
 neighbor 24.1.1.2 remote-as 24
 neighbor 49.1.1.9 remote-as 301
```

We can see that both routers are using their Lo0 interface to form a peer session. When 2 BGP speakers establish a peer session using their loopback interfaces we should see if they have reachability to each other's Loopback interfaces, let's check if they have reachability:

On R3:

```
R3#Ping 4.4.4.4 source loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 3.3.0.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Well it looks pretty good, another issue here is the TTL of the initial SYN packet, since the two routers are establishing an EBGP peer session, the TTL is set to 1, we should increase it since the routers are NOT directly connected:

On R3:

```
R3(config)#Router bgp 3
R3(config-router)#Neighbor 4.4.4.4 ebgp-multihop
```

On R4:

```
R4(config)#Router bgp 24
R4(config-router)#Neighbor 3.3.0.3 ebgp-multihop
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R3:

You should see the following console message on R3:

```
%BGP-5-ADJCHANGE: neighbor 4.4.4.4 Up
```

```
R3#Show ip bgp summary | B Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
4.4.4.4	4		24	6	6	10	0	0 00:00:49	4
35.1.1.5	4		5	21	22	11	0	0 00:14:16	1
37.1.1.7	4		301	81	80	10	0	0 01:10:12	4

On R4:

```
R4#Show ip bgp summ | B Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.3.0.3	4	3	7	7	11	0	0	00:00:56	6
24.1.1.2	4	24	81	84	10	0	0	01:11:05	0
49.1.1.9	4	301	77	80	10	0	0	01:06:52	4

Ticket 9

R3 should use its direct link to get to R6's Lo63 interface, and R6 should do the same to reach R3's Lo1, but for some strange reason R3 prefers the indirect route via R5 and vice versa. You should fix this problem without using the "distance" or an "aggregate" command, or shutdown interfaces or static routes, you should provide redundancy.

Let's find out the IP address of R6's Lo63 interface:

On R6:

```
R6#Show run int lo63 | B interface

interface Loopback63
 ip address 6.6.63.6 255.255.255.255
end
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check R3's routing table entry for R6's Lo63:

On R3:

```
R3#Show ip route 6.6.63.6

Routing entry for 6.6.63.6/32
  Known via "bgp 3", distance 20, metric 0
  Tag 5, type external
  Last update from 35.1.1.5 00:01:27 ago
  Routing Descriptor Blocks:
  * 35.1.1.5, from 35.1.1.5, 00:01:27 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 5
      MPLS label: none
```

The output of the above show command reveals the following:

- The route was learned via BGP in AS 3.
- The distance 20 and the "Type external" indicate that the prefix was advertised by an EBGP and NOT an IBGP peer.
- Last update from 35.1.1.5, indicates that the local router received the update for this prefix via 35.1.1.5 or R5.
- The next hop IP address of 35.1.1.5 means that the local router has to traverse through R5 to get to the Lo63 interface of R6.

Let's shutdown R5 as a BGP Neighbor and see if the alternate path will be injected into the BGP table:

On R3:

```
R3(config)#Router bgp 3
R3(config-router)#Neighbor 35.1.1.5 Shutdown

R3#Show ip route 6.6.63.6

Routing entry for 6.6.63.6/32
  Known via "eigrp 16", distance 90, metric 20640000, type internal
  Redistributing via eigrp 16
  Last update from 36.1.1.6 on Serial1/6, 00:00:26 ago
  Routing Descriptor Blocks:
  * 36.1.1.6, from 36.1.1.6, 00:00:26 ago, via Serial1/6
      Route metric is 20640000, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      Loading 1/255, Hops 1
```

Therefore, in order to get the path directly through R6, R5 should be shutdown as a BGP neighbor.

One should ask why when we have an active EBGP session between R3 and R5, the preferred path is the EBGP and NOT EIGRP path. However, since you are a CCIE candidate, you should NOT be the one asking this type of question (just kidding!). Well this is because the administrative distance of EBGP is 20 whereas, the administrative distance of EIGRP is 90.

There are several ways to fix this problem; the most elegant way is to use the "**BGP Backdoor**" feature, which causes the IOS to set the administrative distance of an EBGP prefix to 200:

On R3:

```
R3(config)#Router bgp 3
R3(config-router)#No Neighbor 35.1.1.5 Shutdown
R3(config-router)#Network 6.6.63.6 mask 255.255.255.255 backdoor
```

To verify the configuration:

On R3:

```
R3#Show ip route 6.6.63.6

Routing entry for 6.6.63.6/32
  Known via "eigrp 16", distance 90, metric 20640000, type internal
  Redistributing via eigrp 16
  Last update from 36.1.1.6 on Serial1/6, 00:00:37 ago
  Routing Descriptor Blocks:
  * 36.1.1.6, from 36.1.1.6, 00:00:37 ago, via Serial1/6
      Route metric is 20640000, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Now let's see R6's path to R3's Lo1 interface:

On R6:

```
R6#Show ip route 3.3.1.3

Routing entry for 3.3.1.0/24
  Known via "bgp 6", distance 20, metric 0
  Tag 5, type external
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Last update from 56.1.1.5 00:04:26 ago
  Routing Descriptor Blocks:
  * 56.1.1.5, from 56.1.1.5, 00:04:26 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 5
      MPLS label: none
```

We can do the same for R6:

On R6:

```
R6(config)#Router bgp 6
R6(config-router)#Network 3.3.1.0 mask 255.255.255.0 backdoor
```

To verify the configuration:

On R6:

```
R6#Show ip route 3.3.1.3

Routing entry for 3.3.1.0/24
  Known via "eigrp 16", distance 90, metric 20640000, type internal
  Redistributing via eigrp 16
  Last update from 36.1.1.3 on Serial1/3, 00:00:09 ago
  Routing Descriptor Blocks:
  * 36.1.1.3, from 36.1.1.3, 00:00:09 ago, via Serial1/3
      Route metric is 20640000, traffic share count is 1
      Total delay is 25000 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Ticket 10

The administrator of R3 complains that he cannot Ping R6's IPv6 Lo0 (FC00::6).

Looking at the diagram we can see that R3, R5 and R6 are participating in a domain called "RIPng". Let's have a look at R3's routing table:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ipv6 route

IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external

C   2002:3::/64 [0/0]
     via Loopback0, directly connected
L   2002:3::3/128 [0/0]
     via Loopback0, receive
R   FEC0::/64 [120/2]
     via FE80::207:EFF:FE1F:BF60, FastEthernet0/0.35
L   FF00::/8 [0/0]
     via Null0, receive
```

We can see a RIPng route but it is for R5's Lo0 and NOT R6's Lo0. Let's check R5's routing table:

On R5:

```
R5#Show ipv6 route

IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
R   FC00::/64 [120/2]
     via FE80::21C:58FF:FEF6:F7C8, Serial1/6
C   FEC0::/64 [0/0]
     via Loopback0, directly connected
L   FEC0::5/128 [0/0]
     via Loopback0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

We can see R6's Lo0 in R5's routing table. So far we know the following:

1. **R5 is exchanging RIPng routes with R3 and R6**
2. **R5 has R6's Lo0 in its routing table, BUT R3 does NOT.**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Maybe the reason is route filtering, let's check R5 for route filtering:

```
R5#Show run | s ipv6 router rip

ipv6 router rip RiPNG
ipv6 router rip RIPNG
```

We can see that filtering was NOT configured, But we also see that R5 is running two different instances of RIPng, "RIPNG" and "RiPNG", let's check the interfaces and see which instance is running on which interface:

```
R5#Show run | Inc interface|rip

interface Loopback0
 ipv6 rip RiPNG enable
interface FastEthernet0/0
 ipv6 rip RiPNG enable
interface FastEthernet0/1
interface Serial1/0
interface Serial1/1
interface Serial1/2
interface Serial1/3
interface Serial1/4
interface Serial1/5
interface Serial1/6
 ipv6 rip RIPNG enable
interface Serial1/7
ipv6 router rip RiPNG
ipv6 router rip RIPNG
```

We can see that the F0/0 interface that is facing R3 is running "RiPNG", whereas, the interface that is facing R6 is running "RIPNG", there are many ways to fix this problem, but one easy way is to fix this problem is to remove the "RIPNG" process on the interface facing R6 and add the "RiPNG" process instead.

```
R5(config)#Int S1/6
R5(config-subif)#ipv6 rip RiPNG enable
R5(config-subif)#No ipv6 rip RIPNG enable
```

Let's check the routing table of R3, you may have to clear the processes using the "Clear ipv6 RIPNG or RiPNG":

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3#Show ipv6 route

IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
C   2002:3::/64 [0/0]
     via Loopback0, directly connected
L   2002:3::3/128 [0/0]
     via Loopback0, receive
R   FC00::/64 [120/3]
     via FE80::207:EFF:FE1F:BF60, FastEthernet0/0.35
R   FEC0::/64 [120/2]
     via FE80::207:EFF:FE1F:BF60, FastEthernet0/0.35
L   FF00::/8 [0/0]
     via Null0, receive
```

Now that we can see R6's Lo0 interface in R3's routing table, let's check reachability using Ping:

To test the configuration:

On R3:

```
R3#Ping Fc00::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00::6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check and see if R6 has a route back to R3:

On R6:

```
R6#Show ipv6 route

IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP

       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
C    FC00::/64 [0/0]
      via Loopback0, directly connected
L    FC00::6/128 [0/0]
      via Loopback0, receive
R    FEC0::/64 [120/2]
      via FE80::207:EFF:FE1F:BF60, Serial1/5
L    FF00::/8 [0/0]
      via Null0, receive
```

On R5:

```
R5#Show ipv6 route

IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
R    FC00::/64 [120/2]
      via FE80::21C:58FF:FEF6:F7C8, Serial1/6
C    FEC0::/64 [0/0]
      via Loopback0, directly connected
L    FEC0::5/128 [0/0]
      via Loopback0, receive
L    FF00::/8 [0/0]
      via Null0, receive
```

Loopback 0 interface of R3 is missing here as well. Let's see if R3 is advertising its Lo0:

On R3:

```
R3#Show ipv6 rip database

RIP process "RIPNG", local RIB
 FC00::/64, metric 3, installed
     FastEthernet0/0.35/FE80::207:EFF:FE1F:BF60, expires in 169 secs
 FEC0::/64, metric 2, installed
     FastEthernet0/0.35/FE80::207:EFF:FE1F:BF60, expires in 169 secs
```

R3 does NOT have its Lo0 in its database. Let's see if "RIPNG" is enabled on this interface:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show run int Lo0 | B interface

interface Loopback0
 ip address 9.18.99.3 255.255.255.255
 ip pim sparse-mode
 ipv6 address 2002:3::3/64
```

Let's enable RIPNG on this interface and try again:

```
R3(config)#Int lo0
R3(config-if)#ipv6 rip RIPNG enable
```

To verify the configuration:

On R3:

```
R3#Ping Fc00::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00::6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

Ticket 11

R1 and R2 are connected with VRF-lite configurations. R1's Lo12 is not reachable on R2.

Let's check the VRF configurations on R1 and R2:

On R1:

```
R1#Show ip vrf

  Name                              Default RD          Interfaces
  V2                                1:2                 Lo12
                                                        Se1/2

R1#Show run int S1/2 | B interface

interface Serial1/2
 ip vrf forwarding V2
 ip address 12.1.1.1 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 clock rate 64000
end

R1#Show run int Lo12 | B interface

interface Loopback12
 ip vrf forwarding V2
 ip address 12.12.12.1 255.255.255.0
end
```

We can see that VRF "V2" is configured on this router and there are two interfaces in this VRF: S1/2 and Lo12. Let's check R2's configuration:

On R2:

```
R2#Sh ip vrf

  Name                              Default RD          Interfaces
  V2                                1:1                 Se1/1

R2#Show run int S1/1 | B interface

interface Serial1/1
 ip vrf forwarding V2
 ip address 12.1.1.2 255.255.255.0
end
```

On R2, VRF called "V2" is configured and ONLY interface S1/1 is placed into this VRF. Let's check the routing table for the configured VRF:

```
R2#Sh ip rou vrf V2 | B Gate
Gateway of last resort is not set

      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, Serial1/1
L        12.1.1.2/32 is directly connected, Serial1/1
```

The Lo12's IP address of R1 is NOT in R2's routing table. Let's check R2's RIPv2 configuration:

On R2:

```
R2#Show run | S router rip

router rip
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 version 2
 redistribute ospf 1 metric 2
 network 0.0.0.0
 no auto-summary
 !
 address-family ipv4 vrf V2
  network 12.0.0.0
  no auto-summary
  version 2
 exit-address-family
```

R2 is configured properly, let's check R1's RIPv2 configuration:

On R1:

```
R1#Sh run | s router rip

router rip
 version 2
 network 12.0.0.0
 no auto-summary
 !
 address-family ipv4 vrf V2
  no auto-summary
  version 2
 exit-address-family
```

The address-family and the VRF is configured correctly BUT, the network statement is NOT configured under the VRF, let's correct that:

On R1:

```
R1(config)#router rip
R1(config-router)#address-family ipv4 vrf V2
R1(config-router-af)#Network 12.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip route vrf V2 | B Gate
Gateway of last resort is not set

      12.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C        12.1.1.0/24 is directly connected, Serial1/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
L        12.1.1.2/32 is directly connected, Serial1/1
R        12.12.12.0/24 [120/1] via 12.1.1.1, 00:00:03, Serial1/1
```

To test the configuration:

On R2:

```
R2#Ping vrf V2 12.12.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.12.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
```

Ticket 12

A worm is spreading on the network. The worm is spreading using ICMP vulnerability with DSCP value of AF12 coming from R4 to R1.

R1 was configured with ACL #101 on its F0/0 interface to mitigate the problem by filtering all AF12 traffic. However it is not working.

Let's start with checking the Access-list on R1:

On R1:

```
R1#Show access-list 101

Extended IP access list 101
    10 deny ip any any dscp af12
    20 permit ip any any (12723 matches)

R1#Show run int F0/0 | B interface

interface FastEthernet0/0
 ip address 14.1.1.1 255.255.255.0
 ip access-group 101 in
 duplex auto
 speed auto
end
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The access-list (101) is configured correctly and it is applied to the correct interface.
Let's test the configuration by triggering the Access-list from R4, first we will use a regular Ping and then, we will Ping with modified TOS field:

On R4:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R4#Ping
Protocol [ip]:
Target IP address: 1.1.1.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: Y
Source address or interface:
Type of service [0]: 48
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 9.18.99.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note the "Type of Service" value entered is 48, this translates to AF12 (please refer to the "Advanced CCIE Routing and Switching" Workbook for an explanation), but for now.... AF12 is equal to a DSCP value of 12 (AFxy, you multiply "x" by 8, and then multiply "y" by 2 and add the two values, which should equal to 12), then, multiply 12 by 4 to get the ToS value.
But why did the second Ping work? It should have failed.

R4 is sending the packets marked with AF12, and it seems like the marking is cleared at the destination.

What can clear the marking?
Since both routers are connected to SW1, let's check SW1's configuration:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show run int F0/1 | B interface

interface FastEthernet0/1
 switchport access vlan 14
 switchport mode dynamic desirable
end

SW1#Show run int F0/4 | B interface

interface FastEthernet0/4
 switchport access vlan 14
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 14,43
 switchport mode trunk
end
```

The interfaces are configured properly, let's check and see if the "MLS QOS" is enabled on this switch:

On SW1:

```
SW1#Show mls qos

QoS is enabled
QoS ip packet dscp rewrite is enabled
```

MLS qos is enabled, to see the behavior:

```
SW1#Show mls qos interface F0/4

FastEthernet0/4
trust state: not trusted
trust mode: not trusted
trust enabled flag: ena
COS override: dis
default COS: 0
DSCP Mutation Map: Default DSCP Mutation Map
Trust device: none
qos mode: port-based
```

The port is NOT trusted; therefore, the switch will rewrite the TOS field according to the configured default CoS, which is ZERO. To see the table that the switch consults for this action:

On SW1:

```
SW1#Show mls qos maps cos-dscp
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cos-dscp map:

     cos:   0  1  2  3  4  5  6  7
   --------------------------------
    dscp:   0  8 16 24 32 40 48 56
```

Note the switch consults the above table and rewrites the incoming CoS value of 0 to outgoing DSCP value of 0.

This problem can be resolved in two different ways:

- **MLS QOS can be disabled**
- **Trust the port**

In this case, since there are no other QoS tasks or tickets we have decided to disable mls qos:

On SW1:

```
SW1(config)#No mls qos
```

To test the configuration:

On R4:

```
R4#Ping
Protocol [ip]:
Target IP address: 1.1.1.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: Y
Source address or interface:
Type of service [0]: 48
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 9.18.99.1, timeout is 2 seconds:

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

Now a regular ping should be successful:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Ticket 13

Erase the startup configuration on the Routers and the Switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

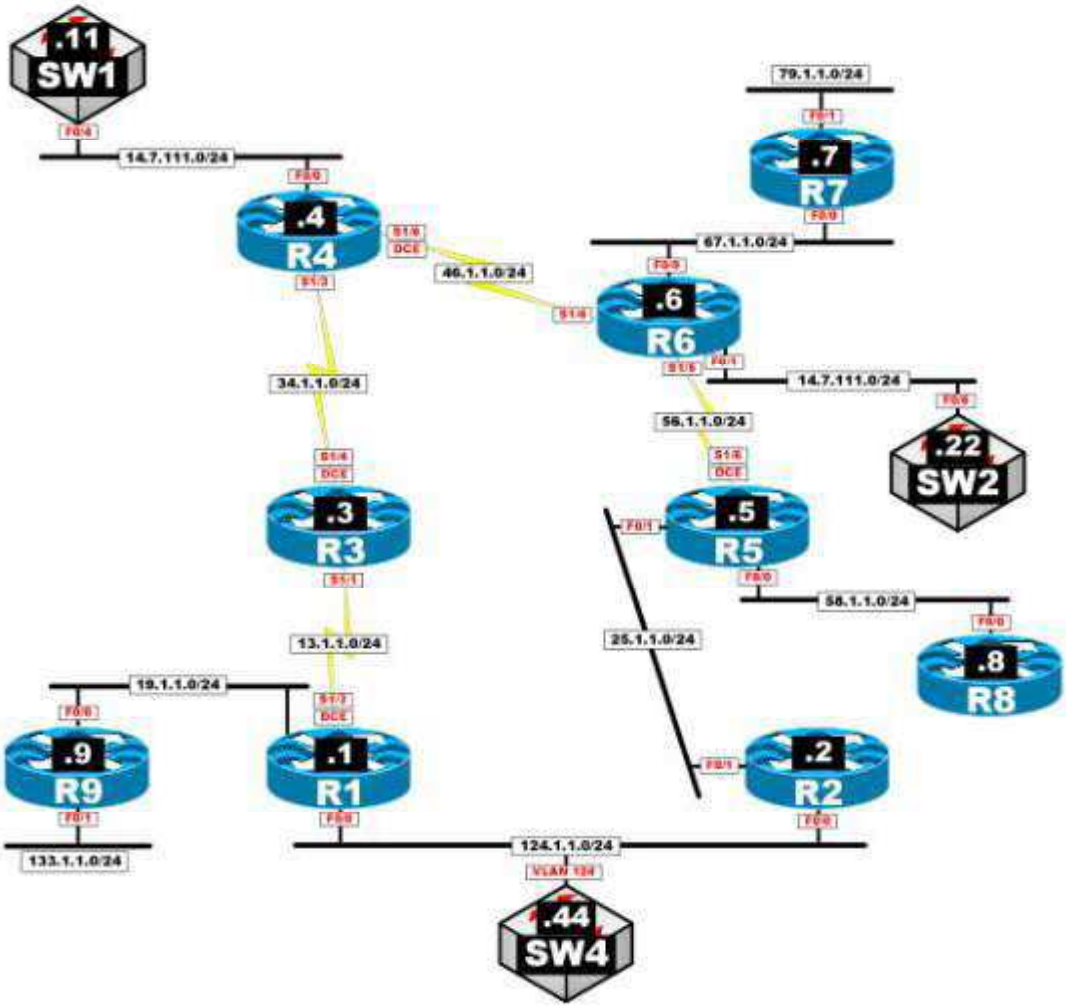
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

TS-Mock-Lab-2

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

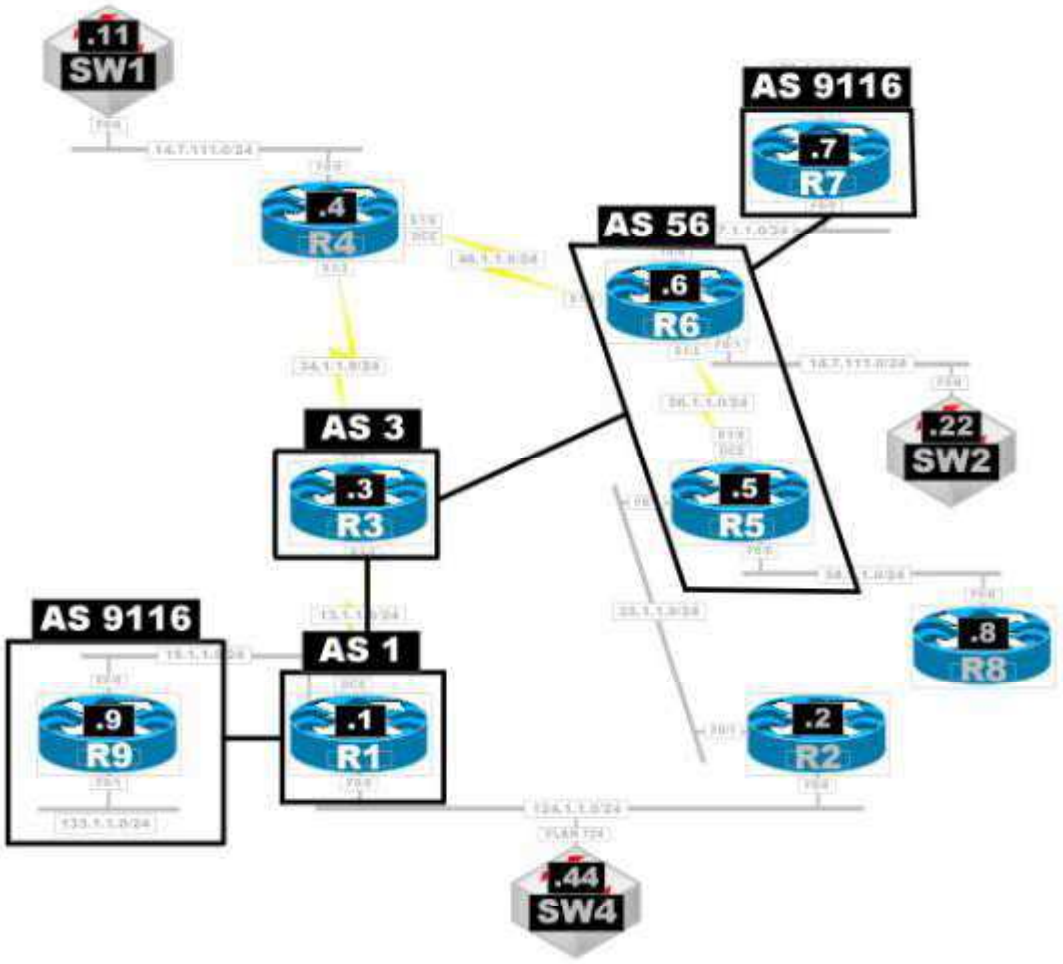
IPv4 Topology



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

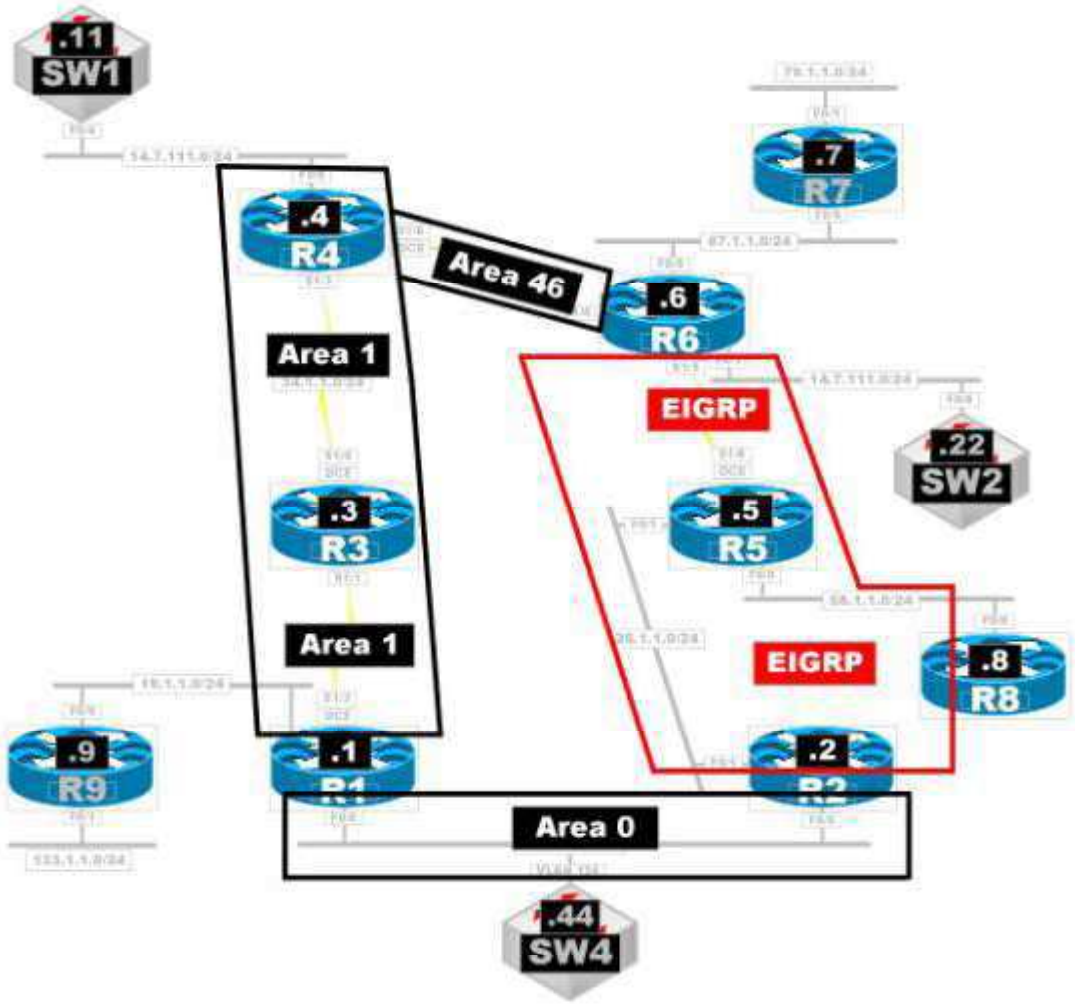
BGP Topology



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

IGP Topology



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "TS-Mock-Labs" → "TS-Mock-Lab-2".

Lab Setup:

- No Static routes, No default routes, NO Policy Based Routing is allowed, unless otherwise specified.
- Use the following IP addressing chart for the BB routers:

Ticket 1

SW4 cannot reach R1's Loopback 0 IP address. You should fix this problem WITHOUT changing the configuration on SW4, R1 or R2.

Before we start the troubleshooting process, the problem should be verified, therefore, let's try to Ping R1's Lo0 from SW4.

To find out R1's Lo0's IP address:

On R1:

```
R1#Show run int Lo0 | B interface

interface Loopback0
 ip address 1.1.1.1 255.255.255.255
 ip ospf 1 area 0
end
```

On SW4:

```
SW4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check SW4's routing table to see if it has a route to that destination:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW4:

```
SW4#Show ip route

Default gateway is 124.1.1.254

Host              Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

It is obvious that this switch is configured as a host and NOT a router, WELL... this should NOT matter but let's see if the default gateway address is reachable from the local switch:

On SW4:

```
SW4#Ping 124.1.1.254

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 124.1.1.254, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Hmm.....Let's check the ARP table for 124.1.1.254:

```
SW4#Show ARP

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  124.1.1.44               -  000c.301f.af00  ARPA   Vlan124
Internet  124.1.1.254              0  Incomplete      ARPA
```

It is obvious that this switch can't resolve the MAC address for its default gateway. To fix this problem, we must find out which router has the 124.1.1.254 IP address. Looking at the diagram, we can see that there are ONLY two routers in VLAN 124, R1 and R2. Let's check R1's F0/0 interface:

On R1:

```
R1#Show run int F0/0 | B interface

interface FastEthernet0/0
 ip address 124.1.1.1 255.255.255.0
 standby 1 ip 124.1.1.254
 standby 1 priority 109
 standby 1 preempt
 standby 1 mac-address 0000.1111.2222
 ip ospf 1 area 0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 duplex auto
 speed auto
end
```

The output of the above show command reveals that SW4's gateway address is an HSRP virtual IP address with a virtual MAC address of 0000.1111.2222.

Let's verify SW4's reachability to R1's F0/0 IP address:

On SW4:

```
SW4#Ping 124.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 124.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1000 ms
```

The result of the above ping command verifies that SW4 has reachability to R1, therefore, HSRP configuration should be verified next:

On R1:

```
R1#Show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Fa0/0       1    109 P Active  local           unknown         124.1.1.254
```

GOOD NEWS, the output of the above show command reveals a potential problem; R1 thinks of itself as the active router but it cannot see the standby router. Let us check R2's configuration, may be R2 HSRP is incorrectly configured?

On R2:

```
R2#Show run int F0/0 | B interface

interface FastEthernet0/0
 ip address 124.1.1.2 255.255.255.0
 standby 1 ip 124.1.1.254
 standby 1 mac-address 0000.1111.2222
 ip ospf 1 area 0
 duplex auto
 speed auto
end
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R2's configuration looks correct; it is similar to R1's configuration. Let's check the status of R2's HSRP:

```
R2#Show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Fa0/0       1    100   Active  local           unknown         124.1.1.254
```

R2's HSRP status is identical to the status of R1's HSRP. R2 thinks of itself as the active router but does NOT see a standby router. The question is can R2 and R1 communicate? Based on the diagram they should be running OSPF and they should be in the same VLAN. Are they OSPF neighbors? Do they have reachability to each other?
Since they are connected via SW1, let's check the VLAN assignment on SW1:

On SW1:

```
SW1#Show vlan brie | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/3, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/25, Fa0/26, Fa0/27, Fa0/28
                                                Fa0/29, Fa0/30, Fa0/31, Fa0/32
                                                Fa0/33, Fa0/34, Fa0/35, Fa0/36
                                                Fa0/37, Fa0/38, Fa0/39, Fa0/40
                                                Fa0/41, Fa0/42, Fa0/43, Fa0/44
                                                Fa0/45, Fa0/46, Fa0/47, Fa0/48
                                                Gi0/1, Gi0/2, Gi0/3, Gi0/4
25   VLAN0025                         active
99   VLAN0099                         active
124  VLAN0124                         active    Fa0/1, Fa0/2
216  VLAN0216                         active    Fa0/6, Fa0/7
225  VLAN0225                         active    Fa0/5, Fa0/8
231  VLAN0231                         active    Fa0/9
```

YES, they are configured in the correct VLAN, Let's see if they can Ping each other's IP address:

On R1:

```
R1#Ping 124.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 124.1.1.2, timeout is 2 seconds:
!!!!!
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

The ping is successful as well, Have they established an OSPF neighbor adjacency?

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
4.4.4.4           0   FULL/  -        -           34.1.1.4        OSPF_VL1
3.3.3.3           0   FULL/  -        -           13.1.1.3        OSPF_VL0
2.2.2.2           1   FULL/BDR        00:00:37    124.1.1.2       FastEthernet0/0
3.3.3.3           0   FULL/  -        00:00:32    13.1.1.3        Serial1/3
```

The output of the above commands verifies that they are neighbors and obviously, they have reachability because Ping was successful. What if something is blocking HSRP on VLAN 124? And maybe that is why SW4 could NOT resolve HSRP's IP address and R1 and R2 cannot exchange HSRP messages?

Well....Since no access-list was seen on the F0/0 interface of R1 and R2, can it be the switch to which they are connected?

Let's check SW1's configuration for the MAC address configured for HSRP:

On SW1:

```
SW1#Show run | Inc 0000.1111.2222

mac-address-table static 0000.1111.2222 vlan 124 drop
```

YES... We can see that SW1 is configured to drop all packets with a MAC address of 0000.1111.2222 as source OR destination.

Let's remove this bogus command from SW1 and see the results:

```
SW1(config)#NO mac-address-table static 0000.1111.2222 VLAN 124 DROP
```

Once this is removed, you should see the following messages on R2 regarding HSRP:

```
%HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Active -> Speak
%HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Speak -> Standby
```

Note: R2 transitioned from Active to standby, which means that they can NOW send and receive HSRP hellos.

On R1:

```
R1#Show standby brief
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                    P indicates configured to preempt.
                    |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Fa0/0       1    109 P Active  local           124.1.1.2       124.1.1.254
```

On R2:

```
R2#Show standby brief
                    P indicates configured to preempt.
                    |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Fa0/0       1    100   Standby 124.1.1.1       local           124.1.1.254
```

The outputs of the above show commands verify that R1 and R2 are exchanging HSRP messages. Let's see if SW4 can ping R1's Lo0 interface:

On SW4

```
SW4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1000 ms

SW4#Show ARP

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  124.1.1.44             -   000c.301f.af00  ARPA   Vlan124
Internet  124.1.1.1              9   0007.0e09.8be0  ARPA   Vlan124
Internet  124.1.1.254            0   0000.1111.2222  ARPA   Vlan124
```

Success!

Important issues:

In HSRP, the hello packets are sent with destination MAC address of the HSRP group MAC address, in this case 0000.1111.2222, and because SW1 was filtering that MAC address, R1 and R2 could NOT exchange hellos.

As far as ARP, remember that the ARP replies are sent with the source MAC address of the HSRP group, and since SW1 was filtering that MAC address, ARP replies NEVER passed from R1/R2 through SW1 all the way to SW4.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The "**Mac address-table static drop**" global configuration command enables MAC address filtering and instructs the switch to DRP traffic with the configured MAC address as source OR destination.

Ticket 2

R3 is configured to deny ingress telnet session. A route-map is configured to implement this policy. When testing this policy you realized that Telnet is still working.

To verify the problem:

On R1:

```
R1#Telnet 13.1.1.3
Trying 14.7.13.3 ... Open

Password required, but none set

[Connection to 14.7.13.3 closed by foreign host]
```

Note Telnet traffic is getting in.

Let's verify the configured route-map on R3:

On R3:

```
R3#Show route-map

route-map TST, permit, sequence 10
  Match clauses:
    ip address (access-lists): 103
  Set clauses:
    interface Null0
  Policy routing matches: 0 packets, 0 bytes
```

The output of the above show command reveals the following:

- The route-map is called TST
- It's matching on packets using access-list 103
- The route-map is configured to drop packets by sending the matched packets to the Null 0 interface.
- BUT there are NO packets matched against the route-map

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's verify the access-list (103):

```
R3#Show access-list 103

Extended IP access list 103
    10 permit tcp any host 13.1.1.3 eq telnet (10 matches)

    20 permit tcp any host 34.1.1.3 eq telnet
    30 permit tcp any host 3.3.3.3 eq telnet
```

NOTE: The packets are matched against the Telnet coming from R3, which means the following:

- The access-list (103) is configured correctly
- The route-map "TST" is configured correctly

So....why is it that R3 does NOT send Telnet packets to NULL 0 interface?
Let's have a look at RFC 1812 requirements for IP routers:

When a router receives an IP packet, it must decide whether the packet is addressed to the router (and should be delivered locally) or the packet is addressed to another system (and should be handled by the forwarder).

Translation:

When a router receives a packet it must determine the destination of the packet, if the packet is destined to one of the router's interface's IP address, then, it must deliver it locally, or else, it should route the packets. This means that Telnet packets to R3 will NOT be routed, hence, will NOT apply the policy based routing using the route-map called "TST".
Since the task requires configuring a route-map to prevent Telnet traffic, the "IP local Policy" command which applies the route-map to locally generated traffic must be configured.

To accomplish this task using a route-map:

- Configure an Access-list to match Telnet traffic
- Configure a route-map to reference the access-list from previous step and send the packet/s to NULL interface
- Apply the route-map on locally originated traffic

On R3:

```
R3(config)#Access-list 133 permit tcp any eq Telnet any

R3(config)#Route-map TST permit 10
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-route-map)#Match ip address 133
R3(config-route-map)#Set interface NULL 0

R3(config)#IP local policy route-map TST
```

To verify the configuration:

On R1:

```
R1#Telnet 13.1.1.3

Trying 14.7.13.3 ...
% Connection timed out; remote host not responding
```

Ticket 3

R8 is configured to filter all incoming EIGRP routes. R5 cannot get EIGRP routes from R8. Fix this problem while ensuring that the routes received from R8 are reachable by R5.

NOTE: R5 is getting the following messages which tells us that its neighbor adjacency to R8 is flapping:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 58.1.1.8 (FastEthernet0/0) is
down: Interface PEER-TERMINATION received

%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 58.1.1.8 (FastEthernet0/0) is
up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 58.1.1.8 (FastEthernet0/0) is
down: Interface PEER-TERMINATION received

%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 58.1.1.8 (FastEthernet0/0) is
up: new adjacency
```

The routing table of R5 will reveal No route/s from R8 when it's checked while the adjacency is UP:

On R5:

```
R5#Show ip route eigrp | B Gate
Gateway of last resort is not set
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      1.0.0.0/32 is subnetted, 1 subnets
D EX     1.1.1.1 [170/2560002816] via 25.1.1.2, 00:51:21, FastEthernet0/1
      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/156160] via 25.1.1.2, 00:51:21, FastEthernet0/1
      3.0.0.0/32 is subnetted, 1 subnets
D EX     3.3.3.3 [170/2560002816] via 25.1.1.2, 00:51:21, FastEthernet0/1
      4.0.0.0/32 is subnetted, 1 subnets
D EX     4.4.4.4 [170/2560002816] via 25.1.1.2, 00:51:21, FastEthernet0/1
      6.0.0.0/32 is subnetted, 1 subnets
D EX     6.6.6.6 [170/2560002816] via 25.1.1.2, 00:51:21, FastEthernet0/1
      13.0.0.0/24 is subnetted, 1 subnets
D EX     13.1.1.0 [170/2560002816] via 25.1.1.2, 00:51:21, FastEthernet0/1
      34.0.0.0/24 is subnetted, 1 subnets
D EX     34.1.1.0 [170/2560002816] via 25.1.1.2, 00:51:21, FastEthernet0/1
      46.0.0.0/24 is subnetted, 1 subnets
D        46.1.1.0 [90/21024000] via 56.1.1.6, 00:52:19, Serial1/6
      124.0.0.0/24 is subnetted, 1 subnets
D        124.1.1.0 [90/30720] via 25.1.1.2, 00:51:21, FastEthernet0/1
```

The ticket states that R8 is filtering all incoming EIGRP routes. But if R8 does NOT know about the networks, how can its routes be reachable? R8 will NOT know how to return the traffic.

One way to allow communication with R8's routes is to use NAT. But R5 is configured with NAT, Let's check and see R5's NAT configuration:

On R5:

```
R5#Show run | Inc interface|nat

interface Loopback0
 ip nat inside
interface FastEthernet0/0
 ip nat outside
interface FastEthernet0/1
 ip nat inside
interface Serial1/0
interface Serial1/1
interface Serial1/2
interface Serial1/3
interface Serial1/4
interface Serial1/5
interface Serial1/6
 ip nat inside
interface Serial1/7
ip nat inside source list 122 interface FastEthernet0/0 overload
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Based on the above output, it is obvious that NAT is used to hide the network behind R5's F0/0 interface, which allows R8 to return traffic without knowing about the networks that are present behind R5.

To cross eliminate NAT as a problem, let's enable debugging NAT on R5:

On R5:

```
R5#Debug ip nat
IP NAT debugging is on

NAT: translation failed (A), dropping packet s=58.1.1.5 d=58.1.1.8
NAT: translation failed (A), dropping packet s=58.1.1.5 d=58.1.1.8
NAT: translation failed (A), dropping packet s=58.1.1.5 d=58.1.1.8
```

Based on the above output, it looks like R5 is trying to perform NAT on EIGRP packets and because of its failure, the packets are dropped. Let's disable the debug and check the access-list (122):

On R5:

```
R5#U all
All possible debugging has been turned off

R5#Show ip access-list 122

Extended IP access list 122
    10 permit ip any any (778 matches)
```

We can see that R5 is configured to translate all packets including EIGRP, it should be configured to NAT everything except EIGRP.

```
R5(config)#No access-list 122

R5(config)#Access-list 122 deny eigrp any any
R5(config)#Access-list 122 permit ip any any
```

Instead of waiting for the neighbor adjacency to come up, let's clear the neighbor table of R5:

```
R5#Clear ip eigrp neighbor
```

You should receive the following messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 56.1.1.6 (Serial1/6) is up:
new adjacency
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 25.1.1.2 (FastEthernet0/1) is
up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 256: Neighbor 58.1.1.8 (FastEthernet0/0) is
up: new adjacency
```

To verify the configuration:

On R5:

```
R5#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
D EX     1.1.1.1 [170/2560002816] via 25.1.1.2, 00:00:41, FastEthernet0/1
      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/156160] via 25.1.1.2, 00:00:41, FastEthernet0/1
      3.0.0.0/32 is subnetted, 1 subnets
D EX     3.3.3.3 [170/2560002816] via 25.1.1.2, 00:00:41, FastEthernet0/1
      4.0.0.0/32 is subnetted, 1 subnets
D EX     4.4.4.4 [170/2560002816] via 25.1.1.2, 00:00:41, FastEthernet0/1
      6.0.0.0/32 is subnetted, 1 subnets
D EX     6.6.6.6 [170/2560002816] via 25.1.1.2, 00:00:41, FastEthernet0/1
      8.0.0.0/32 is subnetted, 1 subnets
D        8.8.8.8 [90/156160] via 58.1.1.8, 00:00:40, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D EX     13.1.1.0 [170/2560002816] via 25.1.1.2, 00:00:41, FastEthernet0/1
      34.0.0.0/24 is subnetted, 1 subnets
D EX     34.1.1.0 [170/2560002816] via 25.1.1.2, 00:00:41, FastEthernet0/1
      46.0.0.0/24 is subnetted, 1 subnets
D        46.1.1.0 [90/21024000] via 56.1.1.6, 00:00:41, Serial1/6
      124.0.0.0/24 is subnetted, 1 subnets
D        124.1.1.0 [90/30720] via 25.1.1.2, 00:00:41, FastEthernet0/1
```

To test reachability:

On R5:

```
R5#Ping 8.8.8.8 source loopback 0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
Packet sent with a source address of 5.5.5.5
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R5

```
R5#Show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 58.1.1.5:0        5.5.5.5:0          8.8.8.8:0          8.8.8.8:0
```

Ticket 4

Routers R2 and R6 are NOT getting routes from R8. You may break one of the rules of this lab.

Let's see the networks advertised by R8:

On R8:

```
R8#Show run | S router eigrp

router eigrp 256
 distribute-list 22 in
 network 8.0.0.0
 network 58.0.0.0

R8#Show ip int brief Exc unass
Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        58.1.1.8        YES manual up                    up
Loopback0              8.8.8.8         YES manual up                    up
```

Let's verify if R5 receives routes from R8:

On R5:

```
R5#Show ip route eigrp | B Gate

      1.0.0.0/32 is subnetted, 1 subnets
D EX     1.1.1.1 [170/2560002816] via 25.1.1.2, 00:11:37, FastEthernet0/1
      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/156160] via 25.1.1.2, 00:11:37, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      3.0.0.0/32 is subnetted, 1 subnets
D EX     3.3.3.3 [170/2560002816] via 25.1.1.2, 00:11:37, FastEthernet0/1
      4.0.0.0/32 is subnetted, 1 subnets
D EX     4.4.4.4 [170/2560002816] via 25.1.1.2, 00:11:37, FastEthernet0/1
      6.0.0.0/32 is subnetted, 1 subnets
D EX     6.6.6.6 [170/2560002816] via 25.1.1.2, 00:11:37, FastEthernet0/1
      8.0.0.0/32 is subnetted, 1 subnets
D        8.8.8.8 [90/156160] via 58.1.1.8, 00:11:36, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D EX     13.1.1.0 [170/2560002816] via 25.1.1.2, 00:11:37, FastEthernet0/1
      34.0.0.0/24 is subnetted, 1 subnets
D EX     34.1.1.0 [170/2560002816] via 25.1.1.2, 00:11:37, FastEthernet0/1
      46.0.0.0/24 is subnetted, 1 subnets
D        46.1.1.0 [90/21024000] via 56.1.1.6, 00:11:37, Serial1/6
      124.0.0.0/24 is subnetted, 1 subnets
D        124.1.1.0 [90/30720] via 25.1.1.2, 00:11:37, FastEthernet0/1
```

The routing table of R5 verifies that it receives 8.8.8.8 /32 prefix from R8. Let's check R6's routing table:

On R6:

```
R6#Show ip route 8.8.8.8

% Network not in table
```

Does R6 receive any prefix/s from R5?

```
R6#Show ip route eigrp | B Gate
Gateway of last resort is not set

      5.0.0.0/32 is subnetted, 1 subnets
D        5.5.5.5 [90/20640000] via 56.1.1.5, 00:13:04, Serial1/5
      25.0.0.0/24 is subnetted, 1 subnets
D        25.1.1.0 [90/20514560] via 56.1.1.5, 00:13:04, Serial1/5
      58.0.0.0/24 is subnetted, 1 subnets
D        58.1.1.0 [90/20514560] via 56.1.1.5, 00:13:04, Serial1/5
```

R6 has all R5's connected routes, but NOT prefix 8.8.8.8 /32. Let's check R5's EIGRP configuration:

On R5:

```
R5#Show run | S router eigrp

router eigrp 256
 network 0.0.0.0
 eigrp stub connected
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Note the "eigrp stub connected" configuration causes R5 to advertise ONLY its connected routes and this is the reason that R6 can ONLY see R5's directly connected prefixes. Since the ticket did not prohibit us from removing any commands, the easiest solution is to remove the configuration.

```
R5(config)#Router eigrp 256
R5(config-router)#No eigrp stub
```

Once the configuration is removed the local router will re-establish its neighbor adjacency with its directly connected neighbors.

To verify the configuration:

On R6:

```
R6#Show ip route eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
D        2.2.2.2 [90/20642560] via 56.1.1.5, 00:00:27, Serial1/5
      5.0.0.0/32 is subnetted, 1 subnets
D        5.5.5.5 [90/20640000] via 56.1.1.5, 00:00:27, Serial1/5
      8.0.0.0/32 is subnetted, 1 subnets
D        8.8.8.8 [90/20642560] via 56.1.1.5, 00:00:27, Serial1/5
      25.0.0.0/24 is subnetted, 1 subnets
D        25.1.1.0 [90/20514560] via 56.1.1.5, 00:00:27, Serial1/5
      58.0.0.0/24 is subnetted, 1 subnets
D        58.1.1.0 [90/20514560] via 56.1.1.5, 00:00:27, Serial1/5
      124.0.0.0/24 is subnetted, 1 subnets
D        124.1.1.0 [90/20517120] via 56.1.1.5, 00:00:27, Serial1/5

R6#Ping 8.8.8.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
```

To verify the configuration on R2:

On R2:

```
R2#Show ip route eigrp | B Gate
Gateway of last resort is not set

      5.0.0.0/32 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D        5.5.5.5 [90/156160] via 25.1.1.5, 00:01:37, FastEthernet0/1
      8.0.0.0/32 is subnetted, 1 subnets
D        8.8.8.8 [90/158720] via 25.1.1.5, 00:01:37, FastEthernet0/1
      46.0.0.0/24 is subnetted, 1 subnets
D        46.1.1.0 [90/21026560] via 25.1.1.5, 00:01:35, FastEthernet0/1
      56.0.0.0/24 is subnetted, 1 subnets
D        56.1.1.0 [90/20514560] via 25.1.1.5, 00:01:37, FastEthernet0/1
      58.0.0.0/24 is subnetted, 1 subnets
D        58.1.1.0 [90/30720] via 25.1.1.5, 00:01:37, FastEthernet0/1

R2#Ping 8.8.8.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 8.8.8.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Ticket 5

When testing this network, you realized that R1 cannot reach R6's Lo0 interface. Fix this problem using minimum number of commands.

Before we can verify the problem, we need to know the IP address of R6's Lo0 interface:

On R6:

```
R6#Show run int lo0 | B interface
interface Loopback0
 ip address 6.6.6.6 255.255.255.255
 ip ospf 1 area 46
end
```

Let's verify the problem:

On R1:

```
R1#Ping 6.6.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check R1's routing table:

```
R1#Show ip route 6.6.6.6

Routing entry for 6.6.6.6/32
  Known via "ospf 1", distance 110, metric 2344, type inter area
  Last update from 13.1.1.3 on Serial1/3, 01:31:06 ago
  Routing Descriptor Blocks:
  * 13.1.1.3, from 4.4.4.4, 01:31:06 ago, via Serial1/3
      Route metric is 2344, traffic share count is 1
```

Explaining the highlighted areas:

- It's a route to 6.6.6.6 /32
- It's known via OSPF, it's an inter-area route, and the ABR is R4
- The route is pointing in the right direction through R3 (13.1.1.3).

To verify the route:

```
R1#Traceroute 6.6.6.6 Numeric

Type escape sequence to abort.
Tracing the route to 6.6.6.6
VRF info: (vrf in name/id, vrf out name/id)
  1 13.1.1.3 12 msec 16 msec 16 msec
  2 13.1.1.1 28 msec 32 msec 28 msec
  3 13.1.1.3 24 msec 28 msec 24 msec
  4 13.1.1.1 40 msec 40 msec 36 msec
  5 13.1.1.3 32 msec 32 msec 36 msec
  6 13.1.1.1 52 msec 48 msec 48 msec
  7 13.1.1.3 48 msec 44 msec 48 msec
  8 13.1.1.1 60 msec 60 msec 60 msec
  9 13.1.1.3 56 msec 56 msec 56 msec
 10 13.1.1.1 72 msec 68 msec 72 msec
 11 13.1.1.3 64 msec 68 msec 64 msec
 12 13.1.1.1 80 msec 80 msec 80 msec
 13 13.1.1.3 76 msec 76 msec 76 msec
 14 13.1.1.1 92 msec 88 msec 92 msec
 15 13.1.1.3 84 msec 88 msec 88 msec
 16 13.1.1.1 100 msec 100 msec 100 msec
 17 13.1.1.3 96 msec 96 msec 96 msec
 18 13.1.1.1 112 msec 108 msec 112 msec
 19 13.1.1.3 108 msec 108 msec 104 msec
 20 13.1.1.1 120 msec 124 msec 120 msec
 21 13.1.1.3 116 msec 116 msec 116 msec
 22 13.1.1.1 132 msec 132 msec 128 msec
 23 13.1.1.3 128 msec 128 msec 128 msec
 24 13.1.1.1 140 msec 140 msec 144 msec
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 25 13.1.1.3 136 msec 136 msec 140 msec
 26 13.1.1.1 152 msec 148 msec 148 msec
 27 13.1.1.3 144 msec 148 msec 148 msec
 28 13.1.1.1 160 msec 164 msec 160 msec
 29 13.1.1.3 156 msec 160 msec 156 msec
 30 13.1.1.1 172 msec 172 msec 168 msec
```

Note the route from R1 to R6's Lo0 is through R3, and it looks like the problem is at R3, because R3 is going through R1 and R1 is going through R3. Let's look at R3's routing table:

On R3:

```
R3#Show ip route 6.6.6.6

Routing entry for 6.6.6.6/32
  Known via "ospf 1", distance 110, metric 3125, type inter area
  Last update from 13.1.1.1 on Serial1/1, 01:33:48 ago
  Routing Descriptor Blocks:
  * 13.1.1.1, from 4.4.4.4, 01:33:48 ago, via Serial1/1
      Route metric is 3125, traffic share count is 1

R3#Show ip route | Inc 6.6.6.6

O IA     6.6.6.6 [110/3125] via 13.1.1.1, 01:34:32, Serial1/1
```

Sure enough, R3 uses R1 as the next-hop.

Let's have a look at R3's OSPF interfaces:

```
R3#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
VL0          1     0               13.1.1.3/24        781   P2P   1/1
Se1/1        1     1               13.1.1.3/24        781   P2P   1/1
Se1/4        1     1               34.1.1.3/24        781   P2P   1/1
Lo0          1     3               3.3.3.3/32         1     LOOP  0/0
```

We can see that R3 has a virtual-link configured and Lo0 of this router is configured in area 3, and that is why this router has a virtual-link configured.

BUT why is R3 pointing back to R1 for a route whose ABR is R4?

We all know that by default all inter-area routes must traverse through the backbone area (Area 0), therefore, it is logical why R3 must go through area 0, but in Cisco IOS 12.3(7) T and better transit capability is enabled by default. Let's check R3's OSPF configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3#Show run | S router ospf

router ospf 1
 no capability transit
 area 1 virtual-link 1.1.1.1
```

WOW… The transit capability is disabled, so R3 must go through area 0 to connect to inter-area routes. Let's enable this feature and verify the path:

```
R3(config)#Router ospf 1
R3(config-router)#capability transit
```

You should immediately see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 6.6.6.6 Up
```

Let's verify the configuration:

On R3:

```
R3#Show ip route 6.6.6.6

Routing entry for 6.6.6.6/32
  Known via "ospf 1", distance 110, metric 1563, type inter area
  Last update from 34.1.1.4 on Serial1/4, 00:02:15 ago
  Routing Descriptor Blocks:
  * 34.1.1.4, from 4.4.4.4, 00:02:15 ago, via Serial1/4
      Route metric is 1563, traffic share count is 1

R3#Traceroute 6.6.6.6 numeric

Type escape sequence to abort.
Tracing the route to 6.6.6.6
VRF info: (vrf in name/id, vrf out name/id)
  1 34.1.1.4 16 msec 16 msec 16 msec
  2 46.1.1.6 28 msec *  28 msec
```

NOTE: R3 is taking the correct route toward R6's Lo0 IP address. The following verifies R1's path to R6's Lo0 IP address.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Traceroute 6.6.6.6 numeric

Type escape sequence to abort.
Tracing the route to 6.6.6.6
VRF info: (vrf in name/id, vrf out name/id)
  1 13.1.1.3 16 msec 16 msec 16 msec
  2 34.1.1.4 28 msec 28 msec 28 msec
  3 46.1.1.6 40 msec *  40 msec
```

To test the connection:

On R1:

```
R1#Ping 6.6.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/87/88 ms
```

Ticket 6

R5 is NOT getting any BGP routes in its routing table. Configure R5 to fix this problem, ensure that R6 sees valid BGP routes; R5 should also have these routes in its routing table. R5 may NOT have reachability to 7.7.7.7 prefix, whereas, R6 should.

Let's verify the problem:

```
R5#Show ip route bgp | B Gate
Gateway of last resort is not set
R5#
```

R5 does NOT have any BGP routes in its routing table.

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i7.7.7.7/32       67.1.1.7                 0    100      0 9116 i
```

The route is in BGP table but did NOT get injected into the routing table, but why? Let's check this route in detail:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show ip bgp 7.7.7.7
BGP routing table entry for 7.7.7.7/32, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  9116
    67.1.1.7 (inaccessible) from 56.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 100, valid, internal
```

We can clearly see the problem; the next hop IP address (67.1.1.7) is inaccessible. This means that it is NOT in R5's routing table. Let's verify that:

```
R5#Show ip route 67.1.1.7

% Network not in table
```

Since the ticket states that R5 should be configured to fix the problem, the following configures R5 to set the next hop IP address to be 56.1.1.6:

```
R5(config)#route-map TST permit 10
R5(config-route-map)#set ip next-hop 56.1.1.6

R5(config)#Router bgp 56
R5(config-router)#Neighbor 56.1.1.6 route-map TST in

R5#Clear ip bgp * in
```

Let's have a look at R5's BGP entry:

On R5:

```
R5#Show ip bgp 7.7.7.7

BGP routing table entry for 7.7.7.7/32, version 2
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  9116
    56.1.1.6 from 56.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R5#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*>i7.7.7.7/32       56.1.1.6                 0    100      0 9116 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

NOW.....Let's check the routing table of R5:

```
R5#Show ip route bgp | B Gate
Gateway of last resort is not set

      7.0.0.0/32 is subnetted, 1 subnets
B        7.7.7.7 [200/0] via 56.1.1.6, 00:01:30
```

Ticket 7

R3 is configured to receive multicast flow for 239.3.3.3 group. When testing this connection, you realized that R6 could not ping 239.3.3.3.

To verify the problem:

On R6:

```
R6#Ping 239.3.3.3

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.3.3.3, timeout is 2 seconds:
.
```

Let's check the mroute table of R6:

On R6:

```
R6#Show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 02:12:59/stopped, RP 0.0.0.0, flags: DCL
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/4, Forward/Sparse, 02:12:59/stopped
```

There are NO entries for 239.3.3.3 group, which means that R6 does not send Multicast traffic toward R3.

Let's verify the multicast configuration of R6:

On R6:

```
R6#Show run | Inc multi|pim|interface

ip multicast-routing
multilink bundle-name authenticated
interface Loopback0
interface Tunnel46
interface FastEthernet0/0
interface FastEthernet0/1
interface Serial1/0
interface Serial1/1
interface Serial1/2
interface Serial1/3
interface Serial1/4
 ip pim sparse-mode
interface Serial1/5
interface Serial1/6
interface Serial1/7
 neighbor 3.3.3.3 ebgp-multihop 255
```

We can see that Multicast routing is enabled and the interface toward R3 is configured to operate in PIM sparse-mode. Therefore, there should be an RP, let's verify R6's RP mapping:

```
R6#Show ip pim rp mapping
PIM Group-to-RP Mappings
R6#
```

Since we did not see a static RP configuration in R6, therefore, It is safe to assume that there should be a BSR or a Mapping-agent.

Let's check the neighbor adjacency to ensure that R6 has a neighbor adjacency with R4:

On R6:

```
R6#Show ip pim neighbor
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
46.1.1.4          Serial1/4                02:16:42/00:01:22 v2    1 / S P G
```

Since R6 has established a neighbor adjacency with R4, we should move closer to R3 and examine R4's RP mapping:

On R4:

```
R4#Show ip pim rp mapping

PIM Group-to-RP Mappings
This system is a candidate RP (v2)
R4#
```

```
R4#Show run | Inc ip pim

 ip pim sparse-mode
 ip pim sparse-mode
 ip pim bsr-candidate Loopback0 0
 ip pim rp-candidate Loopback0
```

The output of the above show commands reveal that R4 is the RP candidate, but it's configured to be both the BSR and the RP. Let's check the BSR status of R4:

On R4:

```
R4#Show ip pim bsr-router

PIMv2 Bootstrap information
This system is a candidate BSR
  Candidate BSR interface Loopback0
    PIMv2 is not configured - BSR messages not originated
  Candidate RP: 4.4.4.4(Loopback0)
    PIMv2 is not configured on Loopback0 - not advertised
```

The Lo0 interface was used as the source of BSR-candidate and RP-candidate BUT PIM is NOT configured on the lo0 interface of this router, this can be verified by the following show command:

```
R4#Show run int Lo0 | B interface

interface Loopback0
 ip address 4.4.4.4 255.255.255.255
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 ip ospf 1 area 1
end
```

To configure the loopback 0 interface:

On R4:

```
R4(config)#int lo0
R4(config-if)#ip pim sparse-mode
```

We should see the following message:

```
%PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 4.4.4.4 on interface
Loopback0
```

To verify the configuration:

On R4

```
R4#Show ip pim rp mapping
```

You may have to wait few seconds for your output to resemble the following.

```
PIM Group-to-RP Mappings
This system is a candidate RP (v2)
This system is the Bootstrap Router (v2)

Group(s) 224.0.0.0/4
  RP 4.4.4.4 (?), v2
    Info source: 4.4.4.4 (?), via bootstrap, priority 0, holdtime 150
         Uptime: 00:00:32, expires: 00:01:53
```

R4 is the elected BSR and the RP for all Multicast groups.

Let's have a look at R6's RP mappings:

On R6:

```
R6#Show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 4.4.4.4 (?), v2
    Info source: 4.4.4.4 (?), via bootstrap, priority 0, holdtime 150
         Uptime: 00:01:25, expires: 00:02:01
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To test the configuration:

On R6

```
R6#Ping 239.3.3.3

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.3.3.3, timeout is 2 seconds:

Reply to request 0 from 34.1.1.3, 92 ms
```

On R4

```
R4#Show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,

       Y - Joined MDT-data group, y - Sending to MDT-data group
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.3.3.3), 00:02:12/00:03:15, RP 4.4.4.4, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/3, Forward/Sparse, 00:02:12/00:03:15

(46.1.1.6, 239.3.3.3), 00:00:25/00:02:34, flags: T
  Incoming interface: Serial1/6, RPF nbr 46.1.1.6
  Outgoing interface list:
    Serial1/3, Forward/Sparse, 00:00:25/00:03:15

(*, 224.0.1.40), 02:24:29/00:00:47, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/3, Forward/Sparse, 02:24:29/00:00:47
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 8

SW1 and SW2 cannot ping each other over the 14.7.111.0 /24 segment.

Let's verify the problem:

On SW1:

```
SW1#Ping 14.7.111.22

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 14.7.111.22, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the configuration of the interfaces on these switches:

```
SW1#Show run int f0/4 | B interface

interface FastEthernet0/4
 no switchport
 ip address 14.7.111.11 255.255.255.0
end
```

On SW2:

```
SW2#Show run int f0/6 | B interface

interface FastEthernet0/6
 no switchport
 ip address 14.7.111.22 255.255.255.0
end
```

Well, the configuration of the switches looks fine. But these interfaces on SW1 and SW2 are connected to R4 and R6 respectively, and R4 and R6 are connected via a Serial link. Let's check the configuration of R4's F0/0 and R6's F0/1 interface:

On R4:

```
R4#Show run int F0/0 | B interface

interface FastEthernet0/0
 no ip address
 duplex auto
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 speed auto
 xconnect 6.6.6.6 46 encapsulation mpls
end
```

On R6:

```
R6#Show run int F0/1 | B interface

interface FastEthernet0/1
 no ip address
 duplex auto
 speed auto
 xconnect 4.4.4.4 64 encapsulation mpls
end
```

It is obvious that Any Transport Over MPLS (AToM) is used to provide reachability. Let's check LDP neighbor adjacency between the routers:

```
R6#Show mpls ldp neighbor

    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 6.6.6.6:0
        TCP connection: 4.4.4.4.646 - 6.6.6.6.59077
        State: Oper; Msgs sent/rcvd: 94/93; Downstream
        Up time: 01:06:56
        LDP discovery sources:
          Serial1/4, Src IP addr: 46.1.1.4
          Targeted Hello 6.6.6.6 -> 4.4.4.4, active, passive
        Addresses bound to peer LDP Ident:
          34.1.1.4        4.4.4.4         46.1.1.4
```

We can see that R4 is the ONLY neighbor to the local router. This means that both R4 and R6 are configured correctly as far as LDP and MPLS.

To see the interfaces that are running MPLS:

```
R6#Show MPLS interfaces

Interface              IP            Tunnel   BGP Static Operational
Serial1/4              Yes (ldp)     No       No  No     Yes
```

Since LDP and MPLS is configured correctly let's look at AToM's configuration closely:

On R4:

```
R4#Show mpls ldp neighbor
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    Peer LDP Ident: 6.6.6.6:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 6.6.6.6.59077 - 4.4.4.4.646
        State: Oper; Msgs sent/rcvd: 93/94; Downstream
        Up time: 01:06:47
        LDP discovery sources:
          Serial1/6, Src IP addr: 46.1.1.6
          Targeted Hello 4.4.4.4 -> 6.6.6.6, active, passive
        Addresses bound to peer LDP Ident:
          6.6.6.6         67.1.1.6        46.1.1.6        56.1.1.6
```

AToM uses targeted LDP session between the routers to form an AToM VC. For this to work, the two routers had to go through a discovery process using UDP 646 and once they discovered each other, a TCP session will be established using port 646 on the server side and a high port on the client side, this is displayed in the output of the above show command. In the lower part of the output, the details of the targeted LDP session is displayed, in this case the two routers are directly connected and therefore, the output displays the direction of the Hellos going from 4.4.4.4 → 6.6.6.6.

Since the information shows no errors, let's check the status of the VC on R4 and then on R6:

On R4:

```
R4#Show MPLS L2transport vc

Local intf     Local circuit                 Dest address    VC ID      Status
-------------  ----------------------------- --------------- ---------- ------
Fa0/0          Ethernet                      6.6.6.6         46         DOWN
```

On R6:

```
R6#Show MPLS L2transport vc

Local intf     Local circuit                 Dest address    VC ID      Status
-------------  ----------------------------- --------------- ---------- --------
Fa0/1          Ethernet                      4.4.4.4         64         DOWN
```

The vc is down on both routers; VCs are negotiated over the targeted LDP session. Let's check the "Xconnect" command once again:

On R4:

```
R4#Show run | Inc xconnect

 xconnect 6.6.6.6 46 encapsulation mpls
```

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Show run | Inc xconnect
```

Note the VC-IDs DO NOT match

```
 xconnect 4.4.4.4 64 encapsulation mpls
```

In AToM the VC-IDs MUST match or else the VC cannot be negotiated properly, let's change R4's VC-ID to match R6's.

On R4:

```
R4(config)#int F0/0
R4(config-if)#No xconnect 6.6.6.6 46 encapsulation MPLS
R4(config-if)#xconnect 6.6.6.6 64 encapsulation MPLS
```

To verify the configuration:

On R4:

```
R4#Show MPLS L2transport vc

Local intf     Local circuit                 Dest address    VC ID      Status
-------------  ----------------------------  --------------  ----------  ---------
Fa0/0          Ethernet                      6.6.6.6         64          UP
```

On R6:

```
R6#Show MPLS L2transport vc

Local intf     Local circuit                 Dest address    VC ID      Status
-------------  ----------------------------  --------------  ----------  ---------
Fa0/1          Ethernet                      4.4.4.4         64          UP
```

To test the configuration:

On SW1:

```
SW1#Ping 14.7.111.22

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 14.7.111.22, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 67/67/68 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 9

You were called to fix a configuration problem that implements the following policy: SW1 should be configured to transmit OSPF packets immediately to R1 and R2's F0/0 interfaces, this policy should be implemented no matter what other traffic is accumulated on R1 and R2's interface buffers. SW1 should also be configured to limit OSPF bandwidth to 33.3 Mbits/sec.

Let's check SW1's interface configuration facing R1 and R2:

On SW1:

```
SW1#Show run int F0/1 | B interface

interface FastEthernet0/1
 switchport access vlan 124
 srr-queue bandwidth shape  0  0  0  4
 priority-queue out
 mls qos trust dscp
end

SW1#Show run int F0/2 | B interface

interface FastEthernet0/2
 switchport access vlan 124
 srr-queue bandwidth shape  0  0  0  4
 priority-queue out
 mls qos trust dscp
end
```

Let's start checking one item at a time, the first thing to check is if "MLS qos" is enabled on this switch:

On SW1:

```
SW1#Show mls qos

QoS is enabled
QoS ip packet dscp rewrite is enabled
```

Mls Qos is enabled. The next step is to see if priority queueing is enabled and if the priority queue is shaped to 1/3 of the interface's bandwidth:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show run int F0/1 | B interface

interface FastEthernet0/1
 switchport access vlan 124
 srr-queue bandwidth shape  0  0  0  4
 priority-queue out
 mls qos trust dscp
end
```

Even though Priority Queueing is enabled, there are two errors here and they are as follows:

1. Queue number, Q4 is NOT the priority queue; Q1 is the priority queue.
2. The shaping value is incorrect; it should have been shaped to 1/3rd and NOT 1/4th.

Let's correct the errors and verify the configuration:

```
SW1(config)#int range f0/1-2
SW1(config-if-range)#srr-queue bandwidth shape 3 0 0 0
```

Before the configuration is verified, we must map OSPF to Q1. By default the Precedence value of the routing protocols are set to 6, which is the same as DSCP CS6 with a decimal value of 48.

Let's view the mapping of DSCP values to the queues:

```
SW1#Show MLS qos maps dscp-output-q

   Dscp-outputq-threshold map:

     d1 :d2    0     1     2     3     4     5     6     7     8     9
     ------------------------------------------------------------

      0 :    02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01
      1 :    02-01 02-01 02-01 02-01 02-01 02-01 03-01 03-01 03-01 03-01
      2 :    03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01
      3 :    03-01 03-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      4 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 04-01 04-01
      5 :    04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      6 :    04-01 04-01 04-01 04-01
```

Note DSCP value of 48 maps to Queue 4 Threshold 1, this should be changed to Q1T1, as follows:

On SW1:

```
SW1(config)#mls qos srr-queue output dscp-map queue 1 48

SW1#Show MLS qos maps dscp-output-q
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Dscp-outputq-threshold map:
  d1 :d2    0     1     2     3     4     5     6     7     8     9
 ------------------------------------------------------------------
  0 :     02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01
  1 :     02-01 02-01 02-01 02-01 02-01 02-01 03-01 03-01 03-01 03-01
  2 :     03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01
  3 :     03-01 03-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
  4 :     01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 04-01
  5 :     04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
  6 :     04-01 04-01 04-01 04-01
```

This looks MUCH better. The last thing to do is to check the trust state of the interfaces running OSPF. This MUST be done because if the switch is NOT configured to trust the DSCP values, then they are dropped.

On SW1:

```
SW1#Show run int F0/2 | Inc mls

 mls qos trust dscp

SW1#Show run int F0/1 | Inc mls

 mls qos trust dscp
```

Success!

Ticket 10

R2 is configured to be accessed via SSH, but the network administrator can ONLY access R2 via Telnet. The username and the password is "CCIE".

Let's verify the configuration by using SSH to R2 from R1:

On R1:

```
R1#SSH -l CCIE 2.2.2.2

% Connection refused by remote host
```

Let's check to see if SSH is enabled on R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Show ip ssh

SSH Disabled - version 1.99
%Please create RSA keys to enable SSH (and of atleast 768 bits for SSH v2).
Authentication timeout: 120 secs; Authentication retries: 3
Minimum expected Diffie Hellman key size : 1024 bits
IOS Keys in SECSH format(ssh-rsa, base64 encoded): NONE
```

NOTE: The output of the above command actually tells us what needs to be done. Let's create the keys:

```
R2(config)#IP domain name CCIE
R2(config)#Crypto key generate rsa general-keys modulus 1024
```

You should see the following messages:

```
The name for the keys will be: R2.CCIE

The name for the keys will be: R2.CCIE

% The key modulus size is 1024 bits
% Generating 1024 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 2 seconds)

%SSH-5-ENABLED: SSH 1.99 has been enabled
```

On R2:

```
R2#Show run | S vty

line vty 0 4
 login
 transport input telnet
```

NOTE: SSH is NOT enabled as a valid method for accessing the VTY lines. Let's fix this:

On R2:

```
R2(config)#line VTY 0 4
R2(config-line)#Transport input ssh
```

Since SSH is enabled, let's SSH from R1 to R2 again:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#SSH -l CCIE 2.2.2.2

Password: ← ENTER CCIE

 [Connection to 2.2.2.2 closed by foreign host]
```

Even though the correct password "CCIE" was entered, connection was refused due to failed authentication. Let's check to see if R2 is configured with a username and password for the user CCIE:

On R2:

```
R2#Show run | Inc username
R2#
```

The above output reveals that aaa is NOT enabled, username and pass pair is NOT configured, and the "line vty 0 4" is NOT configured with "login local", let's start configuring:

On R2:

```
R2(config)#username CCIE password CCIE

R2#Sh run | S vty

line vty 0 4
 login
 transport input ssh
```

The "Login Local" command is also missing:

```
R2(config)#line vty 0 2
R2(config-line)#login local
```

To verify the configuration:

On R1:

```
R1#SSH -l CCIE 2.2.2.2

Password: ← ENTER CCIE

R2>
```

→ You can type "quit" to get back to R1's console.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Perfect, it is working.

Ticket 11

R2 is configured to detect unreachable Telnet clients and once detected it should disconnect them. But your client is stating that it is NOT working.

Remember, Telnet does NOT have a built-in keepalive operation and it totally relies on TCP keepalives:

On R2:

```
R2#Show run | Inc tcp

service tcp-keepalives-out
R2#
```

NOTE: R2 is ONLY activated for TCP keepalives for outgoing TCP sessions and NOT for incoming. Let's configure R2 so it will also do it for incoming sessions:

```
R2(config)#service tcp-keepalives-in
```

Ticket 12

The BGP network was configured to use MED such that R3 prefers R6 as a gateway to reach the prefixes in AS 9116, but for what ever reason R3 is using R1 as the preferred gateway. You should fix this problem such that R3 prefers R6 as its primary gateway, you should use MED to accomplish this task.

Let's verify the problem:

On R3:

```
R3#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*  7.7.7.7/32       6.6.6.6                 10             0 56 9116 i
*>                  13.1.1.1                               0 1 9116 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Note R3 is going through 13.1.1.1 (R1) and NOT R6. Let's look at this entry closer:

```
R3#Show ip bgp 7.7.7.7

BGP routing table entry for 7.7.7.7/32, version 2
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     1
  56 9116
    6.6.6.6 (metric 1563) from 6.6.6.6 (6.6.6.6)

      Origin IGP, metric 10, localpref 100, valid, external
  1 9116
    13.1.1.1 from 13.1.1.1 (1.1.1.1)
      Origin IGP, localpref 100, valid, external, best
```

We can see that R3 takes R1 (13.1.1.1) because it has a lower metric value (No value means zero) assign. When a router compares the MED attribute coming from two routers, it will always take the lower value, because the lower value has more preference. But it's VERY important to know the path selection process and order, the follows identifies the steps:

1. Prefer the path with the highest **WEIGHT**.
   Note: WEIGHT is a Cisco-specific parameter. It is local to the router on which it is configured.
   Both path have the same weight.

2. Prefer the path with the highest **LOCAL_PREF**.
   Note: A path without LOCAL_PREF is considered to have had the value set with the "bgp default local-preference" command, or to have a value of 100 by default.
   Both paths have the same local preference.

3. Prefer the path that was locally originated via a network or aggregate BGP subcommand or through redistribution from an IGP.
   Local paths that are sourced by the **Network** or **Redistribute** commands are preferred over local aggregates that are sourced by the **aggregate-address** command.
   None of the paths are locally originated.

4. Prefer the path with the shortest **AS_PATH**.
   Note: Be aware of these items:

   - This step is skipped if you have configured the **bgp Bestpath as-path ignore** command.
   - An **AS_SET** counts as 1, no matter how many ASes are in the set.
   - The **AS_CONFED_SEQUENCE** and **AS_CONFED_SET** are not included in the AS_PATH length.
     Both paths have the same as path length.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

5. Prefer the path with the lowest Origin type.
   Note: IGP is lower than Exterior Gateway Protocol (EGP), and EGP is lower than INCOMPLETE.
   Both paths have the same origin.

6. Prefer the path with the lowest Multi-Exit Discriminator (MED).
   Note: Be aware of these items:
   - This comparison only occurs if the first (the neighboring) AS is the same in the two paths. Any confederation sub-ASs are ignored.

     In other words, MEDs are compared only if the first AS in the AS_SEQUENCE is the same for multiple paths. Any preceding AS_CONFED_SEQUENCE is ignored.
   - If bgp always-compare-med is enabled, MEDs are compared for all paths.
     You must disable this option over the entire AS. Otherwise, routing loops can occur.
   - If bgp Bestpath med-confed is enabled, MEDs are compared for all paths that consist only of AS_CONFED_SEQUENCE.
     These paths originated within the local confederation.
   - THE MED of paths that are received from a neighbor with a MED of 4,294,967,295 is changed before insertion into the BGP table. The MED changes to to 4,294,967,294.
   - Paths received with no MED are assigned a MED of 0, unless you have enabled bgp Bestpath med missing-as-worst .
     If you have enabled "bgp Bestpath med missing-as-worst", the paths are assigned a MED of 4,294,967,294.
   - The bgp deterministic med command can also influence this step.
     Refer to How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection for a demonstration.

     Both paths have different MED values. Lets stop here and check MED configurations.

Therefore, in order to influence R3's preference such that it takes R6, we need to configure the following:

On R3:

```
R3(config)#router bgp 3
R3(config-router)#bgp Bestpath med missing-as-worst
```

To affect BGP with this policy:

```
R3#Clear ip bgp *
```

To verify the configuration:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip bgp 7.7.7.7

BGP routing table entry for 7.7.7.7/32, version 2
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     2
  56 9116
    6.6.6.6 (metric 1563) from 6.6.6.6 (6.6.6.6)
      Origin IGP, metric 10, localpref 100, valid, external, best
  1 9116

    13.1.1.1 from 13.1.1.1 (1.1.1.1)
      Origin IGP, metric 4294967295, localpref 100, valid, external
```

Ticket 13

Your client is complaining that timezone cannot be configured on R2, he wants to configure the following Timezone:

```
R2(config)#Clock Timezone UTC +12
```

Let's verify the configuration by configuring the requested timezone on R2:

On R2:

```
R2(config)#Clock timezone UTC +12
exit timezone UTC +12
     ^
% Invalid input detected at '^' marker.
```

WOW…What could be wrong here? The commands are entered correctly, but what is EXIT? Let's try one keyword at a time:

On R2:

```
R2(config)#Clock ?
  <cr>
R2(config)#exit
```

OK…..It looks like an alias was configured, let's check for aliases on R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2

```
R2#Show run | Inc alias

alias configure clock exit
```

HAHAHAHA almost got you..... Let's remove the bogus alias:

On R2

```
R2(config)#No alias configure clock exit
```

To verify the configuration:

On R2

```
R2(config)#Clock timezone UTC +12
```

We should see the following console messages:

```
%SYS-6-CLOCKUPDATE: System clock has been updated from 01:59:43 UTC Sat
Mar 2 2002 to 01:59:43 UTC Sat Mar 2 2002, configured from console by
console.
```

Ticket 14

Erase the startup configuration on the Routers and the Switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

www.MicronicsTraining.com

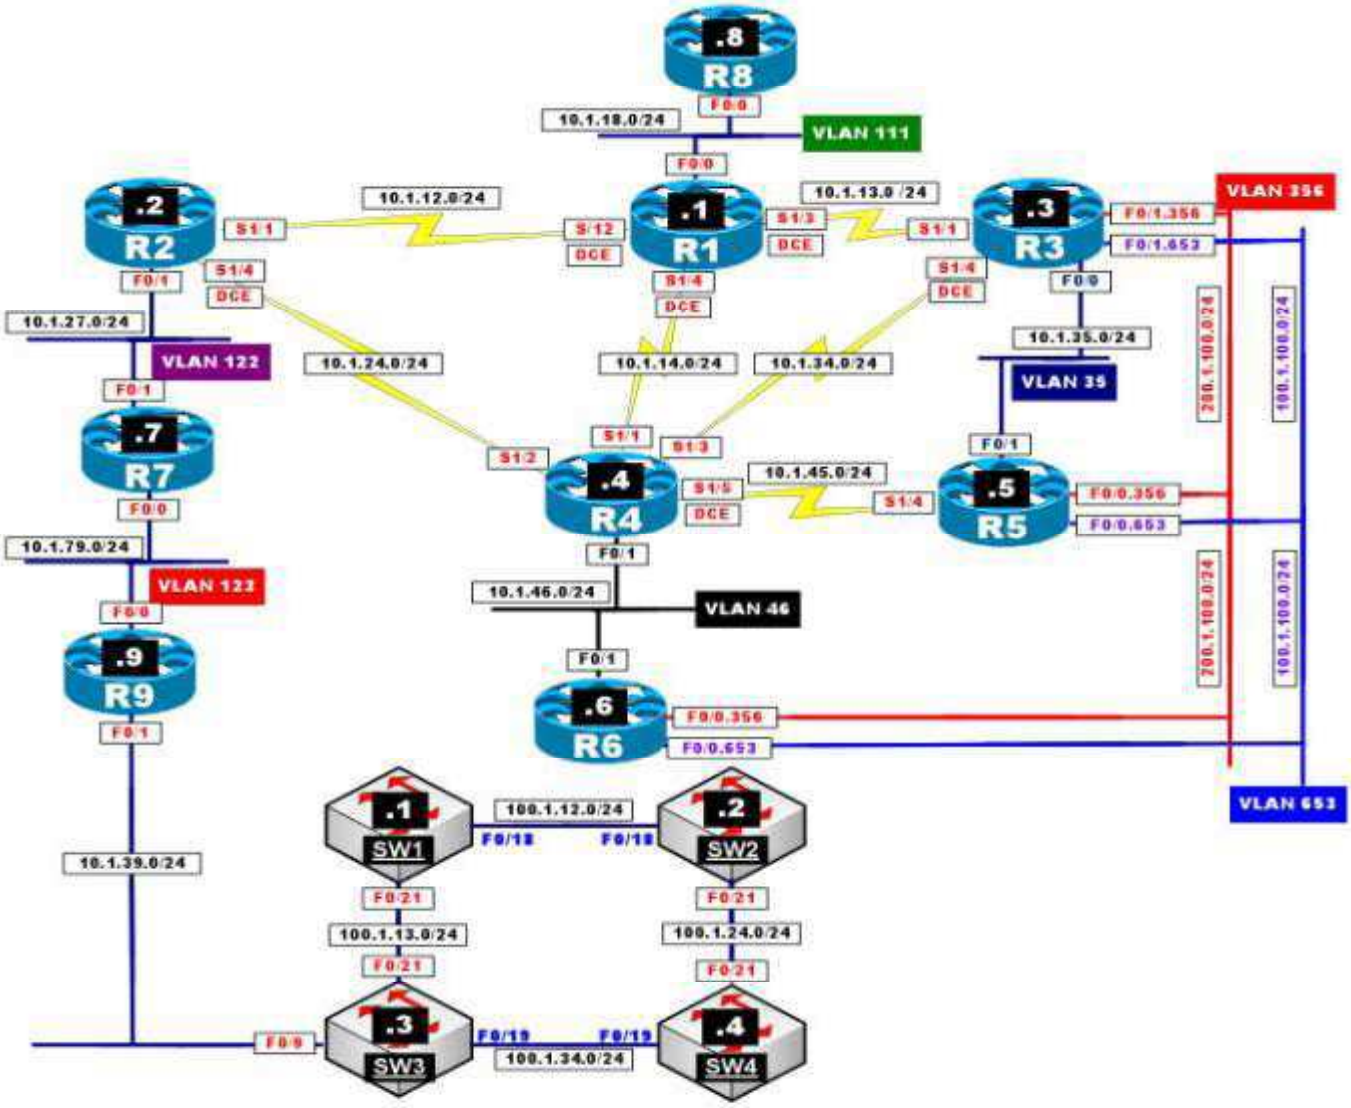
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

TS-Mock-Lab-3

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

IPv4 Topology



Save the running configuration of SW3 and reload the switch before proceeding.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "TS-Mock-Labs" → "TS-Mock-Lab-3".

Ticket 1

Routers R4 and R6 can NOT establish an OSPF adjacency in VLAN 46. DO NOT remove any command/s to fix this problem.

Let's find out if OSPF is running on the correct interfaces of these two routers:

On R4:

```
R4#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Se1/1        41    0               10.1.14.4/24       781   P2P   1/1
Lo0          10    0               44.4.4.4/8         1     LOOP  0/0
Se1/3        10    0               10.1.34.4/24       781   P2P   1/1
Fa0/1        46    0               10.1.46.4/24       1     DR    0/0
```

R4 is running three OSPF processes. PID 46 contains the network address in VLAN 46 and the configuration looks correct so far..... Let's verify R6's information:

On R6:

```
R6#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Fa0/1        64    0               10.1.46.6/24       1     DR    0/0
```

This configuration looks OK as well... Let's verify the problem:

On R4:

```
R4#Show ip ospf 46 neighbor
R4#
```

Let's see if R4 has reachability to R6:

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Ping 10.1.46.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.46.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the ARP table:

On R4:

```
R4#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.46.4              -   0011.5c78.e1a1  ARPA   FastEthernet0/1
Internet  10.1.46.6            159   000e.84b6.88b1  ARPA   FastEthernet0/1
```

Let's verify how these two routers are connected based on the physical diagram, they should be connected to SW2 because they are using their F0/1 interfaces, and they should be configured in VLAN 46:

On SW2:

```
SW2#Show cdp nei | Inc R4|R6

R6               Fas 0/6           130         R S I      2811      Fas 0/1
R4               Fas 0/4           157         R S I      2811      Fas 0/1
```

Hmmmm……let's check the switchport configuration on SW2:

```
SW2#Show run int f0/4 | B inter

interface FastEthernet0/4
 switchport access vlan 46
 switchport mode access
 spanning-tree portfast

SW2#Show run int f0/6 | B inter

interface FastEthernet0/6
 switchport access vlan 46
 switchport mode access
 spanning-tree portfast
```

Let's check and see the VLAN configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show vlan brief | Inc VLAN_|46

VLAN Name                                Status    Ports
46   VLAN0046                            active    Fa0/4, Fa0/6
```

What can be wrong? Let's display the running configuration of SW2 and look for anything that has "46" in it:

On SW2:

```
SW2#Show run | Inc 46

vlan filter tst vlan-list 46
vlan 25,35,46,99,111-112,122-124,216,225,231,356,653
 switchport access vlan 46
 switchport access vlan 46
```

We can see a VLAN Access-map configuration that is applied to VLAN 46, let's check the configuration of this VLAN Access-map:

```
SW2#Sh run | B tst

vlan access-map tst 10
 action drop
 match ip address 1
vlan filter tst vlan-list 46
(The rest of the output is omitted for brevity)
```

Let's see Access-list 1:

```
SW2#Show access-list 1

Standard IP access list 1
    10 permit 209.209.209.209
```

We can clearly see the problem. The configured VLAN access-map called "tst" references access-list 1, which permits host 209.209.209.209, and denys the everything else; basically this vlan access-map states that if the source IP address of any traffic matches the IP address permitted in the access-list, drop the traffic. Since this vlan access-map is NOT configured to instruct the IOS what to do if the traffic does NOT match the configured access-list, IOS will try to protect you, and therefore, it'll drop the rest of the traffic; this behavior is just like a route-map.

Let's fix the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2(config)#vlan access-map tst 90
```

NOTE: The above addition to the existing "Vlan access-map" can fix the problem because it tells the switch that if the traffic does NOT match the access-list 1, do what's norm, which is to forward the traffic. Let's check the console messages on R4 and R6:

On R4:

Once we console to R4, we should se the following console message stating that a neighbor adjacency has been established:

```
%OSPF-5-ADJCHG: Process 46, Nbr 6.6.6.6 on FastEthernet0/1 from LOADING
to FULL, Loading Done
```

On R6:

The same should occur on R6:

```
%OSPF-5-ADJCHG: Process 64, Nbr 44.4.4.4 on FastEthernet0/1 from LOADING
to FULL, Loading Done
```

On R4:

```
R4#Show ip ospf 46 neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
6.6.6.6           1   FULL/BDR        00:00:37    10.1.46.6       FastEthernet0/1

R4#Ping 10.1.46.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.46.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R4#Show ip ospf inter br

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Se1/1        41    0               10.1.14.4/24       781   P2P   1/1
Lo0          10    0               44.4.4.4/8         1     LOOP  0/0
Se1/3        10    0               10.1.34.4/24       781   P2P   1/1
Fa0/1        46    0               10.1.46.4/24       1     DR    1/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 2

R3 and R5 can NOT establish an Eigrp adjacency in Vlan 35 AS 100.

Let's verify the problem, but before we can verify the problem let's see which EIGRP process we are running on the F0/0 interface of R3:

```
R3#Show ip eigrp int
EIGRP-IPv4 Interfaces for AS(100)
                      Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable  Flow Timer   Routes
Fa0/0              0        0/0        0        0/1           0           0
EIGRP-IPv4 Interfaces for AS(653)
                      Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable  Flow Timer   Routes
Lo1                0        0/0        0        0/1           0           0
Fa0/1.653          2        0/0      920        0/1        4604           0
```

On R3:

```
R3#Show ip eigrp 100 neighbors
EIGRP-IPv4 Neighbors for AS(100)
R3#
```

Let's verify the configuration of EIGRP in this routing domain:

```
R3#Show run | S router eigrp 100

router eigrp 100
 Network 10.1.35.3 0.0.0.0
```

```
R3#Show ip int brief | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.35.3       YES NVRAM  up                    up
FastEthernet0/1.356    200.1.100.3     YES manual up                    up
FastEthernet0/1.653    100.1.100.3     YES manual up                    up
Serial1/1              10.1.13.3       YES manual up                    up
Serial1/4              10.1.34.3       YES manual up                    up
Loopback0              33.3.3.3        YES manual up                    up
Loopback1              3.3.3.3         YES manual up                    up
```

Nothing to worry about yet, let's verify R5:

On R5:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show run | s router eigrp 100

router eigrp 100
 network 10.1.35.5 0.0.0.0

R5#Show ip int brief | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0.356    200.1.100.5     YES manual up                    up
FastEthernet0/0.653    100.1.100.5     YES manual up                    up
FastEthernet0/1        10.1.35.5       YES manual up                    up
Serial1/4              10.1.45.5       YES manual up                    up
Loopback1              5.5.5.5         YES manual up                    up
```

They seem to be configured correctly. Do they have reachability?

On R3:

```
R3#Ping 10.1.35.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.35.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the configuration of SW1 and SW2 for this VLAN:

On SW1:

```
SW1#Sh run int F0/3 | B inter

interface FastEthernet0/3
 switchport access vlan 35
 switchport mode access
 spanning-tree portfast
end
```

On SW2:

```
SW2#Show run int F0/5 | B inter

interface FastEthernet0/5
 switchport access vlan 35
 switchport mode access
 spanning-tree portfast
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
end
```

Is there another VLAN access-map configured for this VLAN?

On SW1:

```
SW1#Show vlan access-map
SW1#
```

On SW2:

```
SW2#Show vlan access-map

Vlan access-map "tst"  10
  Match clauses:
    ip  address: 1
  Action:
    drop
Vlan access-map "tst"  90
  Match clauses:
  Action:
    Forward

SW2#Show run | Inc tst

vlan access-map tst 10
vlan access-map tst 90
vlan filter tst vlan-list 46
```

The above VLAN Access-map does not affect these two routers. Let's check and see if the trunk is established between the two switches and VLAN 35 is allowed on the trunk link between the two switches:

On SW1:

```
SW1#Show int trunk

Port        Mode             Encapsulation  Status        Native vlan
Fa0/5       on               802.1q         trunking      1
Fa0/6       on               802.1q         trunking      1
Fa0/19      on               802.1q         trunking      1
Fa0/20      on               isl            trunking      1
Fa0/22      on               802.1q         trunking      1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Vlans allowed on trunk
Fa0/5       1-34,36-4094
Fa0/6       1-34,36-4094
Fa0/19      1-34,36-4094
Fa0/20      35
Fa0/22      1-34,36-4094

Port        Vlans allowed and active in management domain
Fa0/5       1,46,112,122-123,356,653
Fa0/6       1,46,112,122-123,356,653
Fa0/19      1,46,112,122-123,356,653
Fa0/20      35
Fa0/22      1,46,112,122-123,356,653

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/5       1,46,112,122-123,356,653
Fa0/6       1,46,112,122-123,356,653
Fa0/19      1,46,112,122-123,356,653
Fa0/20      35
Fa0/22      none
```

On SW2:

```
SW2#Show int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/3       on           802.1q         trunking      1
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1
Fa0/22      auto         n-isl          trunking      1
Fa0/23      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/3       1-4094
Fa0/19      1-34,36-4094
Fa0/20      35
Fa0/22      1-4094
Fa0/23      1-4094

Port        Vlans allowed and active in management domain
Fa0/3       1,35,46,112,122-123,356,653
Fa0/19      1,46,112,122-123,356,653
Fa0/20      35
Fa0/22      1,35,46,112,122-123,356,653
Fa0/23      1,35,46,112,122-123,356,653
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/3       1,35,46,112,122-123,356,653
Fa0/19      1,46,112,122-123,356,653
Fa0/20      35
Fa0/22      1,35,46,112,122-123,356,653
Fa0/23      1,35,46,112,122-123,356,653
```

WOW..... we can see the problem, one end of the trunk is configured with ISL encapsulation, and the other end is configured with DOT1Q. Let's change one side to match the other, in this case the encapsulation of ISL is changed to DOT1Q:

On SW1:

```
SW1(config)#Int F0/20
SW1(config-if)#swi tru enc do
```

Let's verify the console messages on R3 and R5, this may take 30 seconds:

On R3:

```
%OSPFv3-5-ADJCHG: Process 353, Nbr 5.5.5.5 on FastEthernet0/0 from
LOADING to FULL, Loading Done

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.35.5 (FastEthernet0/0)
is up: new adjacency
```

It seems like few other problems were also fixed, because we can see that OSPFv3 adjacency came up between R3 and R5 as well.

This looked good but few seconds later the adjacency started flapping with the following console messages on both routers:

The Hold-timer expires:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.35.5 (FastEthernet0/0)
is down: holding time expired
```

The neighbor adjacency comes up and the cycle repeats:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.35.5 (FastEthernet0/0)
is up: new adjacency
```

Let's check the hold-timer on these two routers:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show run int f0/0 | B interface

interface FastEthernet0/0
 ip address 10.1.35.3 255.255.255.0
 ip hello-interval eigrp 100 30
 ip hold-time eigrp 100 120
 duplex auto
 speed auto
 ipv6 address FE80::3 link-local
 ipv6 address 35::3/64
 ipv6 ospf 353 area 0
end
```

On R5:

```
R5#Show run int f0/1 | B interface

interface FastEthernet0/1
 ip address 10.1.35.5 255.255.255.0
 ip hello-interval eigrp 100 30
 ip hold-time eigrp 100 20
 duplex auto
 speed auto
 ipv6 address FE80::5 link-local
 ipv6 address 35::5/64
 ipv6 ospf 353 area 0
end
```

We can see that the hold-time's value is configured less than hello-interval and this is why the adjacency between the two routers are flapping. Let's configure R5's Holdtimer to match R3's:

On R5:

```
R5(config)#Int F0/1
R5(config-if)#IP hold-time eigrp 100 120
```

To verify the configuration:

On R3:

```
R3#Show ip eigrp 100 neighbors

IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
0   10.1.35.5               Fa0/0              14 00:02:02    1   200  0  39
```

Ticket 3

R5 and R6 are configured to inject a default route in RIP's routing domain in VLAN 356, but R3 can ONLY see one default route. Fix the problem using a layer three solution.

Let's verify the problem:

On R3:

```
R3#Show ip route rip

R*   0.0.0.0/0 [120/1] via 200.1.100.6, 00:00:23, FastEthernet0/1.356
```

Let's see if both R5 and R6 are configured to inject a default route:

On R5:

```
R5#Show run | s rip

router rip
 version 2
 network 200.1.100.0
 default-information originate
 no auto-summary
```

On R6:

```
R6#Show run | s rip

router rip
 version 2
 network 200.1.100.0
 default-information originate
 no auto-summary
```

We can see that both routers are configured with the "**Default-information originate**" command. Unlike OSPF, in RIP, you don't need to have a default route to inject one, therefore, this should have worked.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The problem in this ticket is an easy and a basic one. When one of the routers (let's say R5) injects a default route, the other two routers receive the default route; Since R6 receives a default route through its F0/0.356 sub-interface, split horizon rules will NOT allow this router to advertise a default route out of that same sub-interface.

Since the ticket states that a layer three solution should be used to fix this problem, the following layer three solution was used:

On R5 and R6:

```
Rx(config)#IP Prefix-list T3 deny 0.0.0.0/0
Rx(config)#IP Prefix-list T3 permit 0.0.0.0/0 LE 32

Rx(config)#Router rip
Rx(config-router)#distribute-list prefix T3 in F0/0.356

Rx#Clear ip route *
```

To verify the configuration:

On R3:

```
R3#Show ip route rip

R*    0.0.0.0/0 [120/1] via 200.1.100.6, 00:00:34, FastEthernet0/1.356
                [120/1] via 200.1.100.5, 00:00:01, FastEthernet0/1.356
```

NOTE: Since both R5 and R6 are configured to filter the default route from getting into the routing table through their F0/0.356 sub-interface, the split horizon rule does NOT apply anymore.

Ticket 4

Remove the previous solution before resolving this ticket.

R5 and R6 are configured to inject a default route in RIP's routing domain in VLAN 356, but R3 can ONLY see one default route. Fix the problem using a layer three solution. DO NOT use the previous solution to fix this ticket.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To remove the solution from the previous ticket:

On R5 and R6:

```
R5(config)#No IP Prefix-list T3

R5(config)#Router rip
R5(config-router)#No distribute-list prefix T3 in F0/0.356
```

On R3, R5 and R6:

```
Rx#Clear ip route *
```

Another way to resolve this problem is to configure R5 and R6 to send their updates as a unicast to R3,

configuring the routers this way, R5 and R6 should NOT see each other's update for the default route.

On Both R5 and R6:

```
Rx(config)#Router rip
Rx(config-router)#Passive-interface F0/0.356
Rx(config-router)#Neighbor 200.1.100.3
```

On R3, R5 and R6:

```
Rx#Clear ip route *
```

To verify the configuration:

On R3:

```
R3#Show ip route rip

R*   0.0.0.0/0 [120/1] via 200.1.100.6, 00:00:01, FastEthernet0/1.356
                [120/1] via 200.1.100.5, 00:00:07, FastEthernet0/1.356
```

Ticket 5

SW1 is configured in AS 100, it's configured to establish a BGP peer session with Switches SW2 and SW3 in AS 200.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Switches SW2, SW3 and SW4 are all configured in BGP AS 200, and they are configured to establish a full mesh IBGP peer session within their AS.

Switches SW2 and SW3 are configured to establish a peer session with SW1 in AS 100.

SW1 is advertising network 1.0.0.0/8 in RIP and BGP. The appropriate router in AS 200 is configured to inject a local preference such that the routers in this AS take the path through SW3 to reach network 1.0.0.0/8.

The customer states that the local preference is working in BGP but a traceroute displays a conflicting story.

DO NOT configure static routes, or add another routing protocol/s, or PBR, EEM, or remove commands from any routing process to fix this problem. DO NOT use Weight attribute.

SW1 – SW4 are running RIPv2 for NLRI.

Let's verify the configuration and the problem:

On SW1:

```
SW1#Show run | B router bgp 100

router bgp 100
 no synchronization
 bgp log-neighbor-changes
 network 1.0.0.0
 neighbor 100.1.12.2 remote-as 200
 neighbor 100.1.12.2 update-source FastEthernet0/18
 neighbor 100.1.13.3 remote-as 200
 neighbor 100.1.13.3 update-source FastEthernet0/21
 no auto-summary
[The rest of the output is omitted]

SW1#Show ip bgp summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
100.1.12.2      4   200      79      79        2    0    0 01:16:12        0
100.1.13.3      4   200      80      79        2    0    0 01:16:13        0

SW1#Show ip int br | Exc unass
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/18       100.1.12.1      YES manual up                    up
FastEthernet0/21       100.1.13.1      YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
```

SW1 uses its F0/18 interface to establish a peer session with SW2 in AS 200, and its F0/21 interface for its peering with SW3 in AS 200.

On SW2:

```
SW2#Show run | B router bgp

router bgp 200
 no synchronization
 bgp log-neighbor-changes
 neighbor 100.1.12.1 remote-as 100
 neighbor 100.1.12.1 update-source FastEthernet0/18
 neighbor 100.1.24.4 remote-as 200
 neighbor 100.1.24.4 update-source FastEthernet0/21
 neighbor 100.1.34.3 remote-as 200
 neighbor 100.1.34.3 update-source FastEthernet0/21
 no auto-summary
[The rest of the output is omitted]

SW2#Show ip bgp summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
100.1.12.1      4   100      91      91        3    0    0 01:28:17        1
100.1.24.4      4   200      94      94        3    0    0 01:31:49        0
100.1.34.3      4   200      94      94        3    0    0 01:30:58        1

SW2#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/18       100.1.12.2      YES NVRAM  up                    up
FastEthernet0/21       100.1.24.2      YES NVRAM  up                    up
```

SW2 is using its F0/18 interface to establish a BGP peer session with SW1 in AS 100, and its F0/21 interface to peer up with SW3 and SW4 in AS 200.

On SW3:

```
SW3#Show run | B router bgp

router bgp 200
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 no synchronization
 bgp default local-preference 400
 bgp log-neighbor-changes
 neighbor 100.1.13.1 remote-as 100
 neighbor 100.1.13.1 update-source FastEthernet0/21
 neighbor 100.1.24.2 remote-as 200
 neighbor 100.1.24.2 update-source FastEthernet0/19
 neighbor 100.1.34.4 remote-as 200
 neighbor 100.1.34.4 update-source FastEthernet0/19
 no auto-summary
[The rest of the output is omitted]

SW3#Show ip bgp summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
100.1.13.1      4   100      96      97        2    0    0 01:33:10        1
100.1.24.2      4   200      99      99        2    0    0 01:35:49        0
100.1.34.4      4   200      99     100        2    0    0 01:36:29        0

SW3#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/9        10.1.39.3       YES NVRAM  up                    up
FastEthernet0/19       100.1.34.3      YES NVRAM  up                    up
FastEthernet0/21       100.1.13.3      YES NVRAM  up                    up
```

SW3 is using its F0/21 interface to establish a BGP peer session with SW1 in AS 100, and its F0/19 interface to peer up with SW2 and SW4 in AS 200.

On SW4:

```
SW4#Show run | B router bgp

router bgp 200
 no synchronization
 bgp log-neighbor-changes
 neighbor 100.1.24.2 remote-as 200
 neighbor 100.1.24.2 update-source FastEthernet0/21
 neighbor 100.1.34.3 remote-as 200
 neighbor 100.1.34.3 update-source FastEthernet0/19
 no auto-summary
[The rest of the output is omitted]

SW4#Show ip bgp summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
100.1.24.2      4   200      104      104         4    0    0 01:41:36        0
100.1.34.3      4   200      105      104         4    0    0 01:41:24        1
```

```
SW4#Show ip int br | Exc unass

Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/19           100.1.34.4      YES NVRAM  up                    up
FastEthernet0/21           100.1.24.4      YES NVRAM  up                    up
```

SW4 has two IBGP peers, one through its F0/19 and another one through its F0/21 interfaces with SW3 and SW2 respectively.

Let's verify the problem:

On SW3:

```
SW3#Show run | Inc bgp

router bgp 200
 bgp default local-preference 400
 bgp log-neighbor-changes
```

The output of the above show command reveals that SW3 is the device that is configured with a local-preference of 400 to influence all prefixes:

```
SW3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          100.1.13.1               0             0 100 i
```

Let's perform a Traceroute and verify the path:

```
SW3#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.13.1 4 msec *  0 msec
```

It looks like SW3 is working fine. Let's check SW4:

On SW4:

```
SW4#Show ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
r>i1.0.0.0          100.1.13.1               0    400      0 100 i
```

We can see a rib failure, but typically rib failures should NOT affect anything, this happens because SW1 has advertised network 1.0.0.0/8 in both BGP and RIP. To see the reason for the rib-failure:

```
SW4#Show ip bgp rib-failure

Network          Next Hop                   RIB-failure   RIB-NH Matches
1.0.0.0          100.1.13.1          Higher admin distance          n/a
```

Let's verify the path to network 1.0.0.0 using a Traceroute:

```
SW4#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.24.2 0 msec 4 msec 0 msec
  2 100.1.12.1 4 msec *  0 msec
```

SW4 received the local preference attached to netwrk 1.0.0.0/8 (Which is 400), but the problem seems to be the next-hop. SW4 should use SW3 as its next-hop and NOT SW2.

On SW2:

```
SW2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
r>i1.0.0.0          100.1.13.1               0    400      0 100 i
r                   100.1.12.1               0             0 100 i
```

SW2 is receiving the local-preference.

```
SW2#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.12.1 0 msec *  0 msec
```

It looks like SW2 also has a next-hop attribute problem, because from SW2's perspect it should go through SW4 → SW3 → SW1, but instead, it is going directly to SW1.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Since the use of PBR and Weight attribute and static routes are prohibited, we can actually use BGP to influence the next-hop; in the following configuration, the border routers (SW2 and SW3) are configured to change the next-hop:

On SW2:

```
SW2(config)#Router bgp 200
SW2(config-router)#Neighbor 100.1.24.4 next-hop-self
SW2(config-router)#Neighbor 100.1.34.3 next-hop-self

SW2#Clear ip bgp *
```

On SW3:

```
SW3(config)#Router bgp 200
SW3(config-router)#Neighbor 100.1.24.2 next-hop-self
SW3(config-router)#Neighbor 100.1.34.4 next-hop-self

SW3#Clear ip bgp *
```

To verify the configuration:

On SW2:

```
SW2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path

*>i1.0.0.0          100.1.34.3               0    400      0 100 i
*                   100.1.12.1               0             0 100 i
```

This looks greate, because SW2's next-hop to reach network 1.0.0.0/8 is SW3, let's see how this switch reaches 100.1.34.3:

```
SW2#Show ip route 100.1.34.3

Routing entry for 100.1.34.0/24
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 100.1.24.4 on FastEthernet0/21, 00:00:07 ago
  Routing Descriptor Blocks:
  * 100.1.24.4, from 100.1.24.4, 00:00:07 ago, via FastEthernet0/21
      Route metric is 1, traffic share count is 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The output of the above show command reveals that SW2 goes through 100.1.24.4 which is SW4 to reach 100.1.34.0/24 network, this should work, let's test this:

On SW2:

```
SW2#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.12.1 0 msec *  0 msec
```

Hmmmmm..... It does NOT work.....Let's try this on the SW4:

On SW4:

```
SW4#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.24.2 4 msec 4 msec 0 msec
  2 100.1.12.1 4 msec *  0 msec
```

WOW..... let's verify the BGP table of SW4:

On SW4:

```
SW4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
r>i1.0.0.0          100.1.34.3               0    400      0 100 i
```

Can it be the rib-failure that is causing this problem?

Let's test this by increasing the administrative distance of RIPv2 to 201, or any value that is higher than 200 (IBGP's Administrative distance).

On SW2, SW3 and SW4:

```
SWx(config)#Router rip
SWx(config-router)#distance 201

SWx#Clear ip route *
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify and test the configuration:

On SW4:

```
SW4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          100.1.34.3               0    400      0 100 i

SW4#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.34.3 0 msec 0 msec 0 msec
  2 100.1.13.1 4 msec *  0 msec
```



On SW2:

```
SW2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          100.1.34.3               0    400      0 100 i
*                   100.1.12.1               0             0 100 i

SW2#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 100.1.24.4 0 msec 0 msec 0 msec
  2 100.1.34.3 0 msec 9 msec 0 msec
  3 100.1.13.1 0 msec *  0 msec
```

NEVER assume that rib-failures are OK.

Since the route was NOT active in BGP, BGP's attributes did NOT take affect. Once the route was active in BGP (The rib failure was resolved), the BGP's next-hop attribute took affect.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 6

The routers in VLAN 653 are configured with Eigrp in AS 653. These routers are configured to advertise their Lo1 interface in this routing domain. These routers should be configured such that existing and future internal routes advertised in this VLAN are all reachable via R3. This policy should NOT affect the directly connected routes. The administrator is totally lost and he/she needs your help.DO NOT use PBR, Static routes, another routing process, EEM to accomplish this task.

Let's verify the configuration:

On R3:

```
R3#Show ip eigrp 653 interfaces
IP-EIGRP interfaces for process 653

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT  Un/Reliable  Flow Timer  Routes
Lo1                0        0/0        0       0/1           0          0
Fa0/1.653          2        0/0      920       0/1        4600          0

R3#Show run | S router eigrp 653

router eigrp 653
 network 3.3.3.3 0.0.0.0
 network 100.1.100.3 0.0.0.0

R3#Show ip route eigrp 653

D    5.0.0.0/8 [90/156160] via 100.1.100.5, 00:34:49, FastEthernet0/1.653
D    6.0.0.0/8 [90/156160] via 100.1.100.6, 00:34:49, FastEthernet0/1.653
```

On R5:

```
R5#Show ip eigrp 653 interfaces
IP-EIGRP interfaces for process 653

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT  Un/Reliable  Flow Timer  Routes
Lo1                0        0/0        0       0/1           0          0
Fa0/0.653          2        0/0        3       0/1          50          0

R5#Show run | s router eigrp 653

router eigrp 653
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 network 5.5.5.5 0.0.0.0
 network 100.1.100.5 0.0.0.0

R5#Show ip route eigrp 653 | B Gate
Gateway of last resort is not set

D    3.0.0.0/8 [90/156160] via 100.1.100.3, 00:36:00, FastEthernet0/0.653
D    6.0.0.0/8 [90/156160] via 100.1.100.6, 00:36:00, FastEthernet0/0.653
```

On R6:

```
R6#Show ip eigrp 653 interfaces
IP-EIGRP interfaces for process 653

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable  Flow Timer   Routes
Lo1                0        0/0       0        0/1           0          0
Fa0/0.653          2        0/0    1021        0/1        5104          0

R6#Show run | s router eigrp 653

router eigrp 653
 network 6.6.6.6 0.0.0.0
 network 100.1.100.6 0.0.0.0

R6#Show ip route eigrp 653

D    3.0.0.0/8 [90/156160] via 100.1.100.3, 00:37:40, FastEthernet0/0.653
D    5.0.0.0/8 [90/156160] via 100.1.100.5, 00:37:40, FastEthernet0/0.653
```

One way to accomplish this task without breaking the rules of this ticket is to configure R5 and R6 to go through R3 for existing and future internal routes, R5 should have R3 as the next hop router to reach networks 3.0.0.0 /8 and 6.0.0.0 /8, and R6 should have R3 as the next hop to reach networks 5.0.0.0 /8 and 3.0.0.0 /8.

To accomplish this task R5 and R6 can NOT have a neighbor adjacency, or else they will exchange routes and advertise themselves as the next hop, therefore, R5 and R6 should both be configured with a neighbor command pointing to R3, and R3 should also be configured with two neighbor commands, one for R5 and another for R6:

On R5 and R6:

```
Rx(config)#Router eigrp 653
Rx(config-router)#Neighbor 100.1.100.3 F0/0.653
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3(config)#Router eigrp 653
R3(config-router)#Neighbor 100.1.100.5 F0/1.653
R3(config-router)#Neighbor 100.1.100.6 F0/1.653
```

Unlike RIPv2, Eigrp can establish a neighbor adjacency in two possible ways, EITHER through Multicast or through Unicast BUT NOT BOTH; meaning that the local router can NOT establish an adjacency using Unicast with one neighbor and multicast with another on a multi-access networks.

To verify the configuration:

On R3:

```
R3#Show ip route eigrp 653
Gateway of last resort is 200.1.100.5 to network 0.0.0.0

D    5.0.0.0/8 [90/156160] via 100.1.100.5, 00:00:35, FastEthernet0/1.653
D    6.0.0.0/8 [90/156160] via 100.1.100.6, 00:00:28, FastEthernet0/1.653
```

On R5:

```
R5#Show ip route eigrp 653
Gateway of last resort is not set

D    3.0.0.0/8 [90/156160] via 100.1.100.3, 00:01:31, FastEthernet0/0.653
```

On R6:

```
R6#Show ip route eigrp 653
Gateway of last resort is 200.1.100.5 to network 0.0.0.0

D    3.0.0.0/8 [90/156160] via 100.1.100.3, 00:01:50, FastEthernet0/0.653
```

The output of the above show commands reveal that R3 can see all the routes advertised by R5 and R6, but R5 and R6 can ONLY see the routes advertised by R3.

The problem is the "Split-Horizon" used by Eigrp, basically when R3 receives networks 5.0.0.0/8 and 6.0.0.0/8 through its F0/1.653 sub-interface, it does NOT advertise them out of that same sub-interface (F0/1.653); therefore, in order for R5 and R6 to see all the routes advertised in this routing domain, the split horizon MUST be disabled on R3's sub-interface:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config)#Int F0/1.653
R3(config-subif)#No ip split-horizon Eigrp 653
```

To verify the configuration:

On R5:

```
R5#Show ip route eigrp 653
Gateway of last resort is not set

D    3.0.0.0/8 [90/156160] via 100.1.100.3, 00:21:32, FastEthernet0/0.653
D    6.0.0.0/8 [90/158720] via 100.1.100.3, 00:01:04, FastEthernet0/0.653

R5#Traceroute 6.6.6.6 numeric

Type escape sequence to abort.
Tracing the route to 6.6.6.6
VRF info: (vrf in name/id, vrf out name/id)
  1 100.1.100.3 0 msec 4 msec 0 msec
  2 100.1.100.6 4 msec *  0 msec
```

On R6:

```
R6#Show ip route eigrp 653
Gateway of last resort is 200.1.100.5 to network 0.0.0.0

D    3.0.0.0/8 [90/156160] via 100.1.100.3, 00:21:51, FastEthernet0/0.653
D    5.0.0.0/8 [90/158720] via 100.1.100.3, 00:01:27, FastEthernet0/0.653

R6#Traceroute 5.5.5.5 numeric

Type escape sequence to abort.
Tracing the route to 5.5.5.5
VRF info: (vrf in name/id, vrf out name/id)
  1 100.1.100.3 4 msec 0 msec 4 msec
  2 100.1.100.5 0 msec *  0 msec
```

Ticket 7

SW3's F0/9 and R9's F0/1 are configured in area 0; R7's G0/0, and R9's F0/0 interfaces are both configured in OSPF area 1; R7's G0/1 and R2's F0/1 interfaces are configured in area 1; R2's S1/1 and R1's S1/2 are configured in area 2.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R1 is advertising prefix 88.8.8.8/32 in OSPF area 2, but the routers in area 0 and some routers in area 1 can NOT see this prefix. DO NOT configure SW3 to accomplish this task.

Let's verify OSPF neighbor adjacency between SW3 and R9:

On SW3:

```
SW3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.39.9         1   EXSTART/BDR     00:00:39    10.1.39.9       FastEthernet0/9
```

Let's check R9:

On R9:

```
R9#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
100.1.34.3        1   EXSTART/DR      00:00:36    10.1.39.3       FastEthernet0/1
10.1.79.7         1   FULL/DR         00:00:34    10.1.79.7       FastEthernet0/0
```

Let's use the "debug ip ospf adj" and may be we can see the problem:

On R9:

```
R9#Debug ip ospf adj

OSPF-8 ADJ   Fa0/1: Rcv DBD from 100.1.34.3 seq 0x23E3 opt 0x52 flag 0x7
len 32  mtu 1504 state EXSTART

OSPF-8 ADJ   Fa0/1: Nbr 100.1.34.3 has larger interface MTU
```

Let's disable the debug:

```
R9#Undebug all
All possible debugging has been turned off
```

On SW3:

```
SW3#Show system mtu
System MTU size is 1504 bytes
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R9:

```
R9#Show int f0/1 | Inc MTU
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
```

Since the ticket states that SW3 can NOT be configured, and on FastEthernet interfaces a user settable MTU is NOT supported, our ONLY choice is to configure the device with the lower MTU to ignore MTU, in this case, R9 should be configured to ignore the MTU:

On R9:

```
R9(config)#Int F0/1
R9(config-if)#IP OSPF mtu-ignore

R9#Clea ip ospf proc
Reset ALL OSPF processes? [no]: Y
```

We should see the following console message on R9:

```
%OSPF-5-ADJCHG: Process 8, Nbr 100.1.34.3 on FastEthernet0/1 from LOADING
to FULL, Loading Done
```

Let's verify the routes:

On SW3:

```
SW3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.39.9         1   FULL/BDR        00:00:31    10.1.39.9       FastEthernet0/9

SW3#Show ip route ospf

     10.0.0.0/24 is subnetted, 3 subnets
O IA    10.1.27.0 [110/3] via 10.1.39.9, 00:01:25, FastEthernet0/9
O IA    10.1.79.0 [110/2] via 10.1.39.9, 00:01:25, FastEthernet0/9
```

Based on the ticket, area 2 is NOT connected to area 0, therefore, a Virtual-link or a GRE tunnel MUST be configured to resolve this ticket; let's verify the configuration of a virtual-link or a GRE tunnel:

On R9:

```
R9#Show run | S router ospf 8
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
router ospf 8
 router-id 10.1.39.9
 area 1 virtual-link 10.1.122.2
 network 10.1.39.9 0.0.0.0 area 0
 network 10.1.79.9 0.0.0.0 area 1

R9#Show ip ospf | Inc ID

Routing Process "ospf 8" with ID 10.1.39.9
```

On R2:

```
R2#Show ip ospf | Inc ID

Routing Process "ospf 8" with ID 10.1.122.2

R2#Sh run | s router ospf

router ospf 8
 router-id 10.1.122.2
 area 1 virtual-link 10.1.39.9
 network 10.1.12.2 0.0.0.0 area 2
 network 10.1.27.2 0.0.0.0 area 1
```

The configuration of the virtual-link looks correct, why is it still down? Let's check the OSPF database of R9 and R2 and focus on area 1:

On R2:

```
R2#Sh ip ospf database


            OSPF Router with ID (10.1.122.2) (Process ID 8)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.122.2      10.1.122.2      338         0x80000003 0x00EB46 0

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.122.2      10.1.122.2      1849        0x80000005 0x008771 1

                Router Link States (Area 2)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.122.2      10.1.122.2      338         0x80000004 0x00C988 2
88.8.8.8        88.8.8.8        575         0x80000004 0x007078 3
```

On R9:

```
R9#Show ip ospf database

            OSPF Router with ID (10.1.39.9) (Process ID 8)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.39.9       10.1.39.9       293         0x80000002 0x00A99A 1
100.1.34.3      100.1.34.3      294         0x80000001 0x004B62 1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.39.9       10.1.39.9       293         0x80000001 0x00A855

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.27.0       10.1.39.9       578         0x80000001 0x003A9F
10.1.79.0       10.1.39.9       578         0x80000001 0x00F1B4

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.39.9       10.1.39.9       582         0x80000007 0x00FEF1 1
10.1.79.7       10.1.79.7       49          0x80000007 0x00125E 2

                Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.79.7       10.1.79.7       297         0x80000003 0x00FADD

                Summary Net Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.39.0       10.1.39.9       578         0x80000001 0x00AB23
```

WOW……These two routers have different databases and they are NOT identical. Did they configure a filter? Let's check:

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9#Show run | Inc ospf

 ip ospf mtu-ignore
router ospf 8

R9#Show run | S router ospf

router ospf 8
 router-id 10.1.39.9
 area 1 virtual-link 10.1.122.2
 network 10.1.39.9 0.0.0.0 area 0
 network 10.1.79.9 0.0.0.0 area 1
```

On R7:

```
R7#Show run | Inc ospf

 ip ospf database-filter all out
router ospf 8

R7#Show run | s router ospf

router ospf 8
 no capability transit
 network 10.1.27.7 0.0.0.0 area 1
 network 10.1.79.7 0.0.0.0 area 1
```

On R7:

```
R7#Show run | Inc ospf|interface

interface Embedded-Service-Engine0/0
interface GigabitEthernet0/0
interface GigabitEthernet0/1
 ip ospf database-filter all out
router ospf 8
```

The output of the above show command reveals that every entry in the database was filtered. In OSPF, the routers within the same area should have an identical database, "may be this is one reason that they are saying that". Let's remove the command and verify the virtual-link:

The filter command is configured under the F0/1 interface.

```
R7(config)#Int g0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7(config-if)#No ip ospf database-filter

R7#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

Let's verify the status on the virtual-link:

On R7:

Within 15 seconds you should see the following console message on R7:

```
%OSPF-5-ADJCHG: Process 8, Nbr 10.1.122.2 on GigabitEthernet0/1 from
LOADING to FULL, Loading Done
```

On R9:

```
R9#Show ip ospf virtual-links

Virtual Link OSPF_VL0 to router 10.1.122.2 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface FastEthernet0/0
 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 00:00:01
    Adjacency State FULL (Hello suppressed)
    Index 2/3, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
R9#
```

To verify the solution:

On SW3:

```
SW3#Show ip route ospf

     10.0.0.0/24 is subnetted, 4 subnets
O IA    10.1.12.0 [110/784] via 10.1.39.9, 00:02:11, FastEthernet0/9
O IA    10.1.27.0 [110/3] via 10.1.39.9, 00:02:11, FastEthernet0/9
O IA    10.1.79.0 [110/2] via 10.1.39.9, 00:02:11, FastEthernet0/9
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
     88.0.0.0/32 is subnetted, 1 subnets
O IA    88.8.8.8 [110/785] via 10.1.39.9, 00:02:11, FastEthernet0/9
```

Ticket 8

R6 and R4 are running IPv6 and OSPFv3 area 0 on their F0/1 interfaces; R3 is running OSPFv3 area 0 on its F0/0 interface, R5 is running OSPFv3 area 0 on its F0/1 interface.

R3 and R4 are configured with 6to4 automatic tunnel over their serial interfaces used to connect isolated IPv6 networks (R4 and R6, and R3 and R5) over IPv4 network.
The firm's policy does not allow this company to run static routes, therefore, OSPFv3 is running on all links, but these routers do not have NLRI to each other.

Let's check OSPFv3's neighbor adjacency between routers R6 and R4, and R3 and R5:

On R6:

```
R6#Show ipv6 ospf neighbor

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
0.0.0.4           1   FULL/DR         00:00:35    4               FastEthernet0/1
```

On R5:

```
R5#Show ipv6 ospf neighbor

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
0.0.0.3           1   FULL/BDR        00:00:31    3               FastEthernet0/1
```

Let's check the configuration of R3 and R4:

On R3:

```
R3#Show ip int br | Exc unass

Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/0            10.1.35.3       YES manual up                    up
FastEthernet0/1.356        200.1.100.3     YES manual up                    up
FastEthernet0/1.653        100.1.100.3     YES manual up                    up
Serial1/1                  10.1.13.3       YES manual up                    up
Serial1/4                  10.1.34.3       YES manual up                    up
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Loopback0                  33.3.3.3        YES manual up                    up
Loopback1                  3.3.3.3         YES manual up                    up
```

Let's verify the configuration of Tunnel 34:

On R3:

```
R3#Show run int tunnel34 | B interface

interface Tunnel34
 no ip address
 no ip redirects
 ipv6 address 2002:2103:303::/128
 ipv6 ospf 343 area 0
 tunnel source Loopback0
 tunnel mode ipv6ip 6to4
end
```

Since Loopback0 is used as the source IP address of this tunnel, let's verify its IP address and see if the translation was done properly:

```
R3#Show run int lo0 | B interface

interface Loopback0
 ip address 33.3.3.3 255.0.0.0
```

Let's do the conversion from Dec to Hex:

Dec	33	3	3	3
Hex	21	03	03	03

The conversion is done properly. Let's check R4's configuration:

On R4:

```
R4#Show ip int br | Inc Tunn

Tunnel43                   unassigned      YES manual up                    up
```

It looks like Tunnel 43 interface is used on R4:

```
R4#Show run int tunnel43 | B interface

interface Tunnel43
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 no ip address
 no ip redirects
 ipv6 address 2002:2C04:404::/128
 ipv6 ospf 343 area 0
 tunnel source Loopback0
 tunnel mode ipv6ip 6to4
end
```

```
R4#Show run int Lo0 | B interface

interface Loopback0
 ip address 44.4.4.4 255.0.0.0
```

Let's do the conversion from Dec to Hex:

Dec	44	4	4	4
Hex	2c	04	04	04

The conversion is done properly. 6to4 tunnels require special addressing on the edge routers, this addressing is concatenation of a reserved IPv6 address of 2002 followed by the translated source IPv4 address of the tunnel. This is configured properly.

The tunnel encapsulation and mode is also configured properly. An automatic 6to4 tunnel is NOT flexible as far as routing protocols are concerned, ONLY static and/or BGP routing protocol can be used to provide NLRI. Therefore, OSPFv3 WILL NOT WORK. Since the ticket states that static routes are prohibited, the ONLY choice is BGP. Let's configure BGP on R3 and R4:

To provide routing, R3 and R4 should both have a static route for 2002 prefix, let's verify:

On R3:

```
R3#Show ipv6 route static

IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
S   2002::/16 [1/0]
     via ::, Tunnel34
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show ipv6 route static

IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
S   2002::/16 [1/0]
     via ::, Tunnel43
```

Let's configure BGP on the edge routers (R3 and R4), in this configuration R3 is configured in AS 300 and R4 is configured in AS 400:

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#Neighbor 2002:2103:303:: remote-as 300
R4(config-router)#Neighbor 2002:2103:303:: ebgp-multihop

R4(config-router)#Address-family ipv6 Unicast
R4(config-router-af)#Neighbor 2002:2103:303:: Act
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Neighbor 2002:2c04:404:: remote-as 400
R3(config-router)#Neighbor 2002:2c04:404:: ebgp-multihop

R3(config-router)#Address-family ipv6 Unicast
R3(config-router-af)#Neighbor 2002:2c04:404:: Act
```

You should see the following console message:

On R3:

```
%BGP-5-ADJCHANGE: neighbor 2002:2C04:404:: Up
```

Once BGP is configured between the two routers, OSPFv3 is redistributed into BGP and BGP is redistributed into OSPF:

On R3:

```
R3(config)#Router bgp 300
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-router)#address-family ipv6 unicast
R3(config-router-af)#Redistribute ospf 353 include-connected

R3(config)#IPv6 router ospf 353
R3(config-rtr)#Redistribute BGP 300
```

On R4:

```
R4(config)#router bgp 400
R4(config-router)#address-family ipv6 unicast
R4(config-router-af)#Redistribute ospf 464 include-connected

R4(config)#ipv6 router ospf 464
R4(config-rtr)#Redistribute bgp 400
```

We should remove the OSPF configuration from the tunnel interfaces:

On R3:

```
R3(config)#Int tunnel 34
R3(config-if)#No IPv6 ospf 343 area 0
```

On R4:

```
R3(config)#Int tunnel 43
R3(config-if)#No IPv6 ospf 464 area 0
```

To verify the configuration:

On R5:

```
R5#Show ipv6 route ospf

IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OE2  46::/64 [110/1]
     via FE80::3, FastEthernet0/1
```

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Show ipv6 route ospf

IPv6 Routing Table - 4 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
OE2  35::/64 [110/1]
     via FE80::4, FastEthernet0/1
```

Ticket 9

Routers R1 and R8 in VLAN 111 can NOT establish an Eigrp peer session in AS 112. DO NOT remove any command/s to fix this ticket.

Let's verify the configuration:

On R1:

```
R1#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(112)
```

Let's see if the correct interface is configured in eigrp:

```
R1#Show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(112)

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable  Flow Timer   Routes
Fa0/0              0        0/0        0        0/1           0           0
```

Let's verify the F0/0 interface of R1:

```
R1#Show run int F0/0 | B interface

interface FastEthernet0/0
ip address 10.1.18.1 255.255.255.0
 duplex auto
 speed auto
end
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check R8's configuration:

On R8:

```
R8#Show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(112)

                   Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT  Un/Reliable  Flow Timer  Routes
Gi0/0              0        0/0        0       0/0          0          0

R8#Show run int G0/0 | B interface

interface GigabitEthernet0/0
 ip address 10.1.18.8 255.255.255.0
 duplex auto
 speed auto
end
```

Let's see if these two routers have reachability to each other:

On R8:

```
R8#Ping 10.1.18.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.18.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Show arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.18.1              113  0013.7f10.7590  ARPA   GigabitEthernet0/0
Internet  10.1.18.8                -  b838.611c.0700  ARPA   GigabitEthernet0/0
```

Let's check their switch connection:

On SW1:

```
SW1#Show run int f0/1 | b inter

interface FastEthernet0/1
 switchport access vlan 111
 switchport mode access
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 spanning-tree portfast
end

SW1#Show run int f0/8 | b inter

interface FastEthernet0/8
 switchport access vlan 111
 switchport mode access
 spanning-tree portfast
end

SW1#Show inter status err-disabled
SW1#

SW1#Show inter status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/1                        connected    111        a-full  a-100 10/100BaseTX
Fa0/2                        notconnect   112          auto   auto 10/100BaseTX
Fa0/3                        connected    35         a-full  a-100 10/100BaseTX
Fa0/4                        notconnect   1            auto   auto 10/100BaseTX
Fa0/5                        connected    trunk      a-full  a-100 10/100BaseTX
Fa0/6                        connected    trunk      a-full  a-100 10/100BaseTX
Fa0/7                        connected    123        a-full  a-100 10/100BaseTX
Fa0/8                        connected    111        a-full  a-100 10/100BaseTX
(The rest of the output is omitted for brevity)

SW1#Show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/4, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/23, Fa0/24, Fa0/25
                                                Fa0/26, Fa0/27, Fa0/28, Fa0/29
                                                Fa0/30, Fa0/31, Fa0/32, Fa0/33
                                                Fa0/34, Fa0/35, Fa0/36, Fa0/37
                                                Fa0/38, Fa0/39, Fa0/40, Fa0/41
                                                Fa0/42, Fa0/43, Fa0/44, Fa0/45
                                                Fa0/46, Fa0/47, Fa0/48, Gi0/1
                                                Gi0/2, Gi0/3, Gi0/4
35   VLAN0035                         active    Fa0/3
46   VLAN0046                         active
111  VLAN0111                         act/unsup Fa0/1, Fa0/8
112  VLAN0112                         active    Fa0/2
122  VLAN0122                         active
123  VLAN0123                         active    Fa0/7, Fa0/9
356  VLAN0356                         active
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
653  VLAN0653                         active
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

WOW…… it displays the VLAN as active but unsupported. Let's check VLAN 111:

```
SW1#Show vlan id 111

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
111  VLAN0111                         act/unsup Fa0/1, Fa0/5, Fa0/6, Fa0/12
                                                Fa0/19, Fa0/22

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
111  fddi  100111     1500  -      -      -        -    -        0      0

Remote SPAN VLAN
----------------
Disabled

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
```

We can see the problem, let's change the media type back to Ethernet:

On SW1:

```
SW1(config)#Vlan 111
SW1(config-vlan)#media Ethernet
```

To verify the configuration:

On SW1:

```
SW1#Show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/4, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/23, Fa0/24, Fa0/25
                                                Fa0/26, Fa0/27, Fa0/28, Fa0/29
                                                Fa0/30, Fa0/31, Fa0/32, Fa0/33
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                                                Fa0/34, Fa0/35, Fa0/36, Fa0/37
                                                Fa0/38, Fa0/39, Fa0/40, Fa0/41
                                                Fa0/42, Fa0/43, Fa0/44, Fa0/45
                                                Fa0/46, Fa0/47, Fa0/48, Gi0/1
                                                Gi0/2, Gi0/3, Gi0/4
35   VLAN0035                         active    Fa0/3
46   VLAN0046                         active
111  VLAN0111                         active    Fa0/1, Fa0/8
112  VLAN0112                         active    Fa0/2
122  VLAN0122                         active
123  VLAN0123                         active    Fa0/7, Fa0/9
356  VLAN0356                         active
653  VLAN0653                         active
```

Let's check the console message of R1:

On R1:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 112: Neighbor 10.1.18.8 (FastEthernet0/0)
is up: new adjacency

R1#Show ip eigrp neighbors

IP-EIGRP neighbors for process 112
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.18.8               Fa0/0             11 00:01:17    1  4500  0  5

R1#Ping 10.1.18.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.18.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Ticket 10

The administrator of R9 is complaining that he can not send any traffic that is 1500 Bytes or larger in size to R2's F0/1 IP address.

Let's verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R9:

```
R9#Ping 10.1.27.2 Size 1500

Type escape sequence to abort.
Sending 5, 1500-byte ICMP Echos to 10.1.27.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see if a policy is configured on R9:

On R9:

```
R9#Show policy-map interface
R9#

R9#Sh route-map
R9#
```

Let's try the same on R7:

On R7:

```
R7#Show policy-map interface
 GigabitEthernet0/1

  Service-policy output: tst

    Class-map: class-default (match-any)
      4102 packets, 426469 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      police:
          cir 1000000 bps, bc 1500 bytes
        conformed 149 packets, 15110 bytes; actions:
          transmit
        exceeded 5 packets, 7570 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps
```

If 7570 is divided by 5 (5 Packets) = 1514. But we pinged with 1500 Byte packets.

Let's clear the counters and try again:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R7:

```
R7#Clear counters
Clear "show interface" counters on all interfaces [confirm]
PRESS ENTER

R9#Ping 10.1.27.2 Size 1500 rep 10

Type escape sequence to abort.
Sending 10, 1500-byte ICMP Echos to 10.1.27.2, timeout is 2 seconds:
..........
Success rate is 0 percent (0/10)
```

Let's verify the ping:

```
R7#Show policy-map interface
GigabitEthernet0/1

  Service-policy output: tst

    Class-map: class-default (match-any)
      28 packets, 17188 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      police:
          cir 1000000 bps, bc 1500 bytes
        conformed 0 packets, 0 bytes; actions:
          transmit
        exceeded 10 packets, 15140 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps
```

In Class based Policing the layer 2 overhead is NOT considered, to verify this, R9 should successfully ping with a packet size of 1500 minus 14 (The overhead) which is 1486. Let's verify this:

On R9:

```
R9#Ping 10.1.27.2 Size 1486 rep 1

Type escape sequence to abort.
Sending 1, 1486-byte ICMP Echos to 10.1.27.2, timeout is 2 seconds:
!
Success rate is 100 percent (1/1), round-trip min/avg/max = 1/1/1 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To fix this ticket, the "Bc" value of the class-based Policing configured on R7 should be adjusted to 1514 or higher.

On R7:

```
R7(config)#policy-map tst
R7(config-pmap)#class class-default
R7(config-pmap-c)#police 1000000
R7(config-pmap-c-police)#conform-action transmit
R7(config-pmap-c-police)#exceed-action drop
R7(config-pmap-c-police)#end
```

NOTE: When configuring CB-Policing, if the Bc is not configured, the IOS will use (CIR/32) or 1500 Bytes or which ever one is higher formula to assign one to the policy. To verify this:

On R7:

```
R7#Sh policy-map inter
 GigabitEthernet0/1

  Service-policy output: tst

    Class-map: class-default (match-any)
      143 packets, 48235 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
      police:
          cir 1000000 bps, bc 31250 bytes
        conformed 15 packets, 14192 bytes; actions:
          transmit
        exceeded 17 packets, 22806 bytes; actions:
          drop
        conformed 0 bps, exceed 0 bps
```

NOTE: 1000,000 divided by 32 is equal 31250

To test the configuration:

On R9:

```
R9#Ping 10.1.27.2 size 1514 rep 10

Type escape sequence to abort.
Sending 10, 1514-byte ICMP Echos to 10.1.27.2, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 4/4/4 ms
```

Ticket 11

Erase the startup configuration on the Routers and the Switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

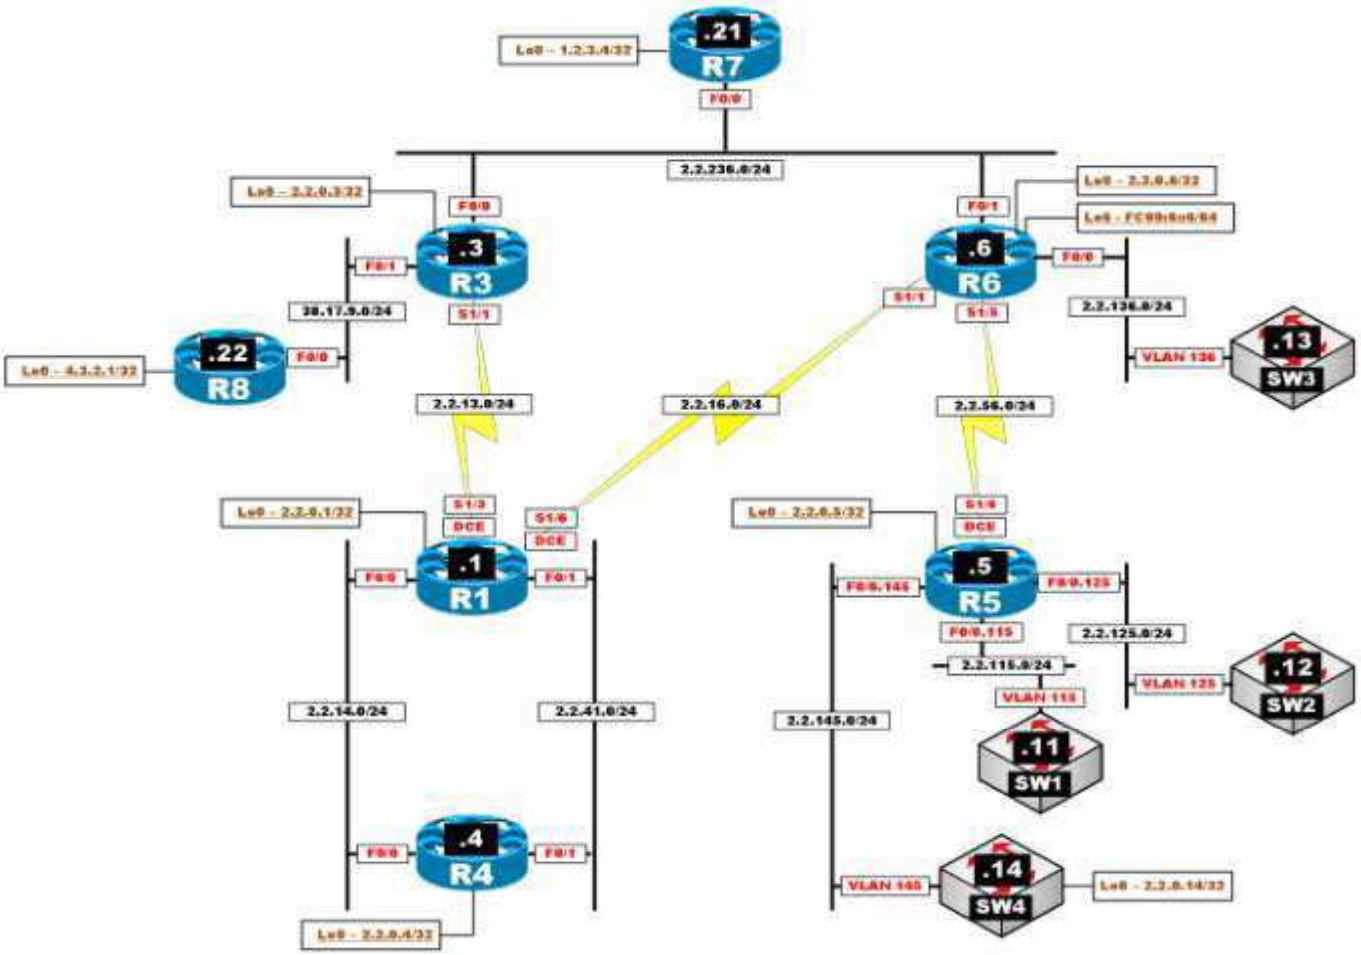
www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

TS-Mock-Lab-4

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "TS-Mock-Labs" →"TS-Mock-Lab-4".

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Rules:

- No Static routes, No default routes, NO Policy Based Routing is allowed, unless otherwise specified.

Ticket 1

R3 and R6 cannot establish an OSPF neighbor adjacency.

Let's check and see which interfaces are in OSPF routing domain:

On R3:

```
R3#Show ip ospf inter brie

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     36               2.2.0.3/32         1     LOOP  0/0
Se1/1        1     36               2.2.13.3/24        781   DOWN  0/0
Fa0/0        1     36               2.2.236.3/24       1     DR    0/0
```

On R6:

```
R6#Show ip ospf interface brie

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     36              2.2.0.6/32         1     LOOP  0/0
Se1/5        1     36              2.2.56.1/31        781   P2P   1/1
Se1/1        1     36              2.2.16.6/24        1562  P2P   1/1
Fa0/1        1     36              2.2.236.6/24       1     DR    0/0
```

The output of the above Show commands reveals that OSPF area 36 is running on the correct interfaces; R3 and R6 are both claiming to be the "DR" which means that they don't see each other, let's verify reachability between these two routers using Ping:

```
R6#Ping 2.2.236.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.236.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Well....we do have another router (R7) on this VLAN, let's check R6's reachability to R7:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Ping 2.2.236.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.236.21, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the ARP table of R6:

```
R6#Show ip arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  2.2.236.3              0   Incomplete      ARPA
Internet  2.2.236.6              -   001a.6ca1.ca19  ARPA   FastEthernet0/1
Internet  2.2.236.21             0   Incomplete      ARPA
```

Based on the output of the above show command R6 does NOT have any reachability to R3 or R7 that are configured in the same VLAN.

Is this a one-way connectivity? Let's Ping R6 and R7 from R3 and verify the ARP table of R3:

On R3:

```
R3#Ping 2.2.236.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.236.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 2.2.236.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.236.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

R3#Show ip arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  2.2.236.3              -   001d.46de.4d18  ARPA   FastEthernet0/0
Internet  2.2.236.6              0   Incomplete      ARPA
Internet  2.2.236.21             0   0017.5a2e.9ca8  ARPA   FastEthernet0/0
Internet  38.17.9.3              -   001d.46de.4d19  ARPA   FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R3 does NOT have reachability to R6 but it has reachability to R7. Well...the F0/0 interface of R7 and R3 are connected to SW1, whereas, the F0/1 interface of R6 is connected to SW2. Could there be a connectivity problem between the two switches?
Let's check the connection between SW1 and SW2, let's start with checking the CDP table of SW1:

On SW1:

```
SW1#Show cdp neighbor | Inc SW

SW4              Fas 0/23          141          R S I      WS-C3550- Fas 0/23
```

So far it seems like SW1's F0/23 is connected to SW4's F0/23, let's check the cdp table of SW4:

On SW4:

```
SW4#Show cdp neighbors | Inc SW

SW1              Fas 0/23          153          R S I      WS-C3560-2Fas 0/23
SW2              Fas 0/21          168          R S I      WS-C3560-2Fas 0/21
```

So far, SW1's F0/23 is connected to SW4's F0/23, and SW4's F0/21 interface is connected to SW2's F0/21. Let's check the spanning-tree status of these links, starting with F0/23 interface of SW1:

On SW1:

```
SW1#Show spanning-tree interface F0/23

Vlan             Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -------
VLAN0001         Root FWD 19        128.25   P2p
```

WOW...ONLY VLAN 1 is active on this link. The link should be a trunk, let's check to see if this link is trunking:

On SW1: Note the link is NOT trunking

```
SW1#Show interface F0/23 trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/23      auto         negotiate      not-trunking  1

Port        Vlans allowed on trunk
Fa0/23      1

Port        Vlans allowed and active in management domain
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/23      1

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      1
```

The output of the above show command reveals that F0/23 interface is configured in AUTO state, and it is waiting for DTP negotiations. Let's check the other end of this link which is SW4:

On SW4:

```
SW4#Show interface F0/23 trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/23      auto         802.1q         not-trunking  1

Port        Vlans allowed on trunk
Fa0/23      1

Port        Vlans allowed and active in management domain
Fa0/23      1

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      1
```

We can see that the F0/23 interface of SW4 is also in AUTO mode, and once again, it is waiting for DTP negotiations. Well...this seems to be the problem here, since both ends of the link are configured in AUTO mode, none of them start the negotiation process, let's change one of the switches in desirable mode:

On SW4:

```
SW4(config)#Int f0/23
SW4(config-if)#Switchport mode dynamic desirable
```

You should see the following console messages on SW4:

```
Line protocol on Interface FastEthernet0/23, changed state to down
Line protocol on Interface FastEthernet0/23, changed state to up
```

To verify the configuration:

On SW4:

```
SW4#Show interface F0/23 trunk
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Mode             Encapsulation  Status        Native vlan
Fa0/23      desirable        802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/23      1-4094

Port        Vlans allowed and active in management domain
Fa0/23      1,14,41,115,125,136,145,223,236

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      none

SW4#Show spanning-tree interface F0/23

Vlan             Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
VLAN0001         Desg FWD 19        128.23   P2p
VLAN0014         Desg FWD 19        128.23   P2p
VLAN0041         Desg FWD 19        128.23   P2p
VLAN0115         Desg FWD 19        128.23   P2p
VLAN0125         Desg FWD 19        128.23   P2p
VLAN0136         Desg FWD 19        128.23   P2p
VLAN0145         Desg FWD 19        128.23   P2p
VLAN0223         Desg FWD 19        128.23   P2p
VLAN0236         Desg FWD 19        128.23   P2p
```

NOTE: SW1 and SW4 are trunking and spanning-tree is starting to learn MAC addresses for all the VLANs, let's check the console messages of R3 and R6:

To verify the configuration:

On R3:

Once we console in to R3, we should see the following console messages:

```
%PIM-5-NBRCHG: neighbor 2.2.236.6 UP on interface FastEthernet0/0

%PIM-5-DRCHG: DR change from neighbor 2.2.236.3 to 2.2.236.6 on interface
FastEthernet0/0

%OSPF-5-ADJCHG: Process 1, Nbr 2.2.0.6 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

You may have to wait few seconds to see the following output.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.0.1           0   FULL/  -        00:00:37    2.2.13.1        Serial1/1
2.2.0.6           1   FULL/DR         00:00:34    2.2.236.6       FastEthernet0/0

R3#Show ip route ospf | B Gate
Gateway of last resort is not set

      2.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
O        2.2.0.1/32 [110/782] via 2.2.13.1, 00:17:58, Serial1/1
O        2.2.0.5/32 [110/783] via 2.2.236.6, 00:01:20, FastEthernet0/0
O        2.2.0.6/32 [110/2] via 2.2.236.6, 00:01:20, FastEthernet0/0
O E2     2.2.0.12/32 [110/20] via 2.2.236.6, 00:01:20, FastEthernet0/0
O        2.2.16.0/24 [110/1563] via 2.2.236.6, 00:01:20, FastEthernet0/0
O        2.2.56.0/31 [110/782] via 2.2.236.6, 00:01:20, FastEthernet0/0
O E2     2.2.115.0/24 [110/20] via 2.2.236.6, 00:01:20, FastEthernet0/0
O E2     2.2.125.0/24 [110/20] via 2.2.236.6, 00:01:20, FastEthernet0/0
```

Ticket 2

SW1 does NOT have any reachability to the Lo0 interface of SW2 or R5. DO NOT remove any command/s, or enable routing, or using static routes or run/modify any routing protocol/s to fix this ticket.

Let's verify the problem by pinging R5's Lo0 interface:

On SW1:

```
SW1#Ping 2.2.0.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Well.....Let's check the routing table of SW1:

```
SW1#Show ip route | B Gateway
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
     2.0.0.0/24 is subnetted, 1 subnets
C       2.2.115.0 is directly connected, Vlan115
S*   0.0.0.0/0 is directly connected, Vlan115
```

Let's try to ping R5's F0/0.115:

```
SW1#Ping 2.2.115.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.115.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms
```

The Ping proves that SW1 does have reachability to R5's interface in VLAN 115, let's check R5's Lo0, may be it's shutdown:

On R5:

```
R5#Sh ip rou connected | B Gate
Gateway of last resort is 2.2.125.12 to network 0.0.0.0

      2.0.0.0/8 is variably subnetted, 14 subnets, 3 masks
C        2.2.0.5/32 is directly connected, Loopback0
C        2.2.56.0/31 is directly connected, Serial1/6
L        2.2.56.0/32 is directly connected, Serial1/6
C        2.2.115.0/24 is directly connected, FastEthernet0/0.115
L        2.2.115.5/32 is directly connected, FastEthernet0/0.115
C        2.2.125.0/24 is directly connected, FastEthernet0/0.125
L        2.2.125.5/32 is directly connected, FastEthernet0/0.125
```

Well it seems to be fine, let's check and see how the default route was configured on SW1:

On SW1:

```
SW1#Show run | Inc ip route

ip route 0.0.0.0 0.0.0.0 Vlan115
```

They have configured a static route pointing to a broadcast interface. When a static route is configured to point to a broadcast interface, ARP can be a problem:

```
SW1#Ping 2.2.0.5 rep 2

Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 2.2.0.5, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
..
Success rate is 0 percent (0/2)

SW1#Show ip arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  2.2.0.5                0   Incomplete      ARPA
Internet  2.2.115.5              4   0007.0e1f.bf60  ARPA   Vlan115
Internet  2.2.115.11             -   0016.46f6.e9c1  ARPA   Vlan115
```

We can see the "Incomplete" ARP entry for 2.2.0.5 IP address outside of VLAN 115. This could be "proxy-arped", let's check the F0/0.115 interface of R5:

On R5:

```
R5#Show ip interface F0/0.115 | Inc Proxy

  Proxy ARP is disabled
  Local Proxy ARP is disabled
```

Let's see R5's F0/0.115's configuration:

```
R5#Show run int f0/0.115 | B interface

interface FastEthernet0/0.115
 encapsulation dot1Q 115
 ip address 2.2.115.5 255.255.255.0
 no ip proxy-arp
 ip pim dr-priority 100
end
```

Let's enable Proxy-ARP and test:

On R5:

```
R5(config)#Int F0/0.115
R5(config-subif)#IP Proxy-arp
```

To verify and test the configuration:

On SW1:

```
SW1#Ping 2.2.0.5

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 2.2.0.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8 ms

SW1#Ping 2.2.0.12

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.12, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Ticket 3

Your client keeps on complaining about R6 loosing the RIP routes that it receives from SW3. Fix this problem so R6 keeps the routes that it receives from SW3, unless SW3 or its interface to R6 is down. DO NOT remove any command/s to fix this problem.

Whenever you troubleshoot flapping routes, "Debug ip routing" is the BEST tool to use. But before enabling the debug, let's verify reachability between R6 and SW3:

On R6:

```
R6#Ping 2.2.136.13

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.136.13, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's inspect R6's ARP table:

```
R6#Show ip arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  2.2.236.3              8   001d.46de.4d18  ARPA   FastEthernet0/1
Internet  2.2.236.6              -   001a.6ca1.ca19  ARPA   FastEthernet0/1
Internet  2.2.236.21             8   0017.5a2e.9ca8  ARPA   FastEthernet0/1
```

We should have seen an entry for 2.2.136.13, but we do not have any entry in the ARP table, NOT even an "Incomplete" entry. Maybe the interface facing SW3 is down. Let's check the status of all interfaces on R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Show ip int brief | Exc unass

Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/0            2.2.136.6       YES manual up                    up
FastEthernet0/1            2.2.236.6       YES manual up                    up
Serial1/1                  2.2.16.6        YES manual up                    up
Serial1/5                  2.2.56.1        YES manual up                    up
Loopback0                  2.2.0.6         YES manual up                    up
NVI0                       2.2.0.6         YES unset  up                    up
```

The F0/0 interface of R6 is UP and running. Could it be the mask?

```
R6#Show ip route connected | B Gate

      2.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
C        2.2.0.6/32 is directly connected, Loopback0
C        2.2.16.0/24 is directly connected, Serial1/1
L        2.2.16.6/32 is directly connected, Serial1/1
C        2.2.56.0/31 is directly connected, Serial1/5
L        2.2.56.1/32 is directly connected, Serial1/5
C        2.2.236.0/24 is directly connected, FastEthernet0/1
L        2.2.236.6/32 is directly connected, FastEthernet0/1

R6#Show run int F0/0 | B inter

interface FastEthernet0/0
 ip vrf forwarding 34
 ip address 2.2.136.6 255.255.255.0
 ip rip advertise 5
 duplex auto
 speed auto
end
```

The output of the above show command reveals that the F0/0 interface is running in VRF 34. Let's Ping and display the ARP table for this VRF:

```
R6#Ping vrf 34 2.2.136.13

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.136.13, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R6#Show ip arp VRF 34

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Internet  2.2.136.6               -   001a.6ca1.ca18  ARPA   FastEthernet0/0
Internet  2.2.136.13             10   000c.302d.2c80  ARPA   FastEthernet0/0
```

Now that we have found that the interface is configured in a VRF, let's see if a "debug ip routing vrf 34" command can indicate what the problem is:

```
R6#Debug ip routing VRF 34
IP routing debugging is on for VRF 34

RT(34): delete subnet route to 2.2.0.13/32

RT(34): updating rip 2.2.0.13/32 (0x1):
    via 2.2.136.13 Fa0/0

RT(34): add 2.2.0.13/32 via 2.2.136.13, rip metric [120/1]

RT(34): delete route to 2.2.0.13 via 2.2.136.13, rip metric [120/1]

RT(34): no routes to 2.2.0.13, entering holddown

RT(34): delete subnet route to 2.2.0.13/32
```

The output of the debug command verifies the problem, but what could be the cause of this problem? Let's disable the debug:

```
R6#Undebug All
All possible debugging has been turned off
```

Let's check the configuration of RIPv2 on R6:

```
R6#Show run | S router rip
router rip
 !
 address-family ipv4 vrf 34
  timers basic 5 10 10 15
  redistribute bgp 56 metric 1
  network 2.0.0.0
  no auto-summary
  version 2
 exit-address-family
ipv6 router rip RIPNG
 redistribute ospf 1 metric 1
```

The only change here are the timers, they are NOT the default. Let's check the configuration of RIP on SW3 and compare it to R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Show ip protocol vrf 34 | Inc Send|Inva

  Sending updates every 5 seconds, next due in 4 seconds
  Invalid after 10 seconds, hold down 10, flushed after 15
    Interface              Send  Recv  Triggered RIP  Key-chain
```

On SW3:

```
SW3#Show ip protocol | Inc Send|Inva

  Sending updates every 30 seconds, next due in 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
```

R6 invalidates the routes after 10 seconds and flushes them after 15 seconds, whereas SW3 sends an update every 30 seconds. This means every 15 seconds the routes are flushed and then, 15 seconds later they are advertised by SW3.

To fix the problem, the "timers basic" command is changed on SW3 to match R6:

```
SW3(config)#router rip
SW3(config-router)#Timers basic 5 10 10 15
```

To verify and test the configuration:

On R6:

```
R6#Show ip route vrf 34 rip | B Gate
Gateway of last resort is not set

      2.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R        2.2.0.13/32 [120/1] via 2.2.136.13, 00:00:04, FastEthernet0/0
```

Ticket 4

R5 is configured to filter the default route it receives from SW2, but for some strange reason it is still receiving the default route.

Let's verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R5:

```
R5#Show ip route eigrp | B Gate
Gateway of last resort is 2.2.125.12 to network 0.0.0.0

D*EX  0.0.0.0/0 [170/2560002816] via 2.2.125.12, 00:59:33, FastEthernet0/0.125
      2.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
D        2.2.0.12/32 [90/156160] via 2.2.125.12, 00:59:33, FastEthernet0/0.125
```

Let's check R5's configuration:

```
R5#Show run | s router eigrp

router eigrp 125
 distribute-list prefix pl125 in FastEthernet0/0
 network 2.0.0.0
 redistribute ospf 1 metric 1 1 1 1 1
```

Let's check the prefix-list called "pl125":

```
R5#Show ip prefix-list

ip prefix-list pl12: 2 entries
   seq 5 deny 0.0.0.0/0
   seq 10 permit 0.0.0.0/0 le 32
ip prefix-list pl125: 0 entries
```

It looks like the name of the prefix was misspelled, let's try it and see the result:

```
R5(config)#Router eigrp 125
R5(config-router)#No distribute-list prefix pl125 in F0/0
R5(config-router)#distribute-list prefix pl12 in F0/0
```

To verify the configuration:

On R5:

```
R5#Show ip route eigrp | B Gate
Gateway of last resort is 2.2.125.12 to network 0.0.0.0

D*EX  0.0.0.0/0 [170/2560002816] via 2.2.125.12, 01:08:03, FastEthernet0/0.125
      2.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
D        2.2.0.12/32 [90/156160] via 2.2.125.12, 01:08:03, FastEthernet0/0.125
```

The default route is still there. One thing that you should have noticed is the fact that EIGRP did NOT

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Resync its neighbor adjacency when we changed the policy. But why?

The reason is that the "distribute-list" command was configured on the wrong interface. Remember that R5's F0/0 is configured as a dot1q trunk, and any layer 3 configuration applied to the major interface will NOT be inherited by the sub-interfaces, therefore, the "distribute-list" command should be applied to the sub-interface (F0/0.125) and NOT the major interface (F0/0). Let's fix this problem:

```
R5(config)#router eigrp 125
R5(config-router)#No distribute-list prefix pl12 in F0/0
R5(config-router)#distribute-list prefix pl12 in F0/0.125
```

You should see the following console message within few seconds:

```
IP-EIGRP(0) 125: Neighbor 2.2.125.12 (FastEthernet0/0.125) is resync:
route configuration changed
```

To verify the configuration:

On R5:

```
R5#Show ip route eigrp | B Gate
Gateway of last resort is not set

 2.0.0.0/8 is variably subnetted, 13 subnets, 3 masks
D    2.2.0.12/32 [90/156160] via 2.2.125.12, 00:10:17, FastEthernet0/0.125
```

Ticket 5

R3 and R4 can't reach each other's Lo0. DO NOT remove any commands to fix this problem.

Let's verify the problem:

On R3:

```
R3#Ping 2.2.0.4 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.4, timeout is 2 seconds:
Packet sent with a source address of 2.2.0.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

Let's use a "Traceroute" and see where the communication fails, the following Traceroute is generated with a minimum TTL of 1 and a maximum TTL of 4:



On R3:

```
R3#Traceroute 2.2.0.4 source Loopback 0 ttl 1 4

Type escape sequence to abort.
Tracing the route to 2.2.0.4

  1  *  *  *
  2  *  *  *
  3  *  *  *
  4  *  *  *
```

It looks like no one answered R3's call. Let's look at R3's routing entry for R4's Lo0:

```
R3#Show ip route 2.2.0.4

% Subnet not in table
```

Since there is no specific route for 2.2.0.4, maybe a default route is used, let's check the routing table of R3:

```
R3#Show ip route | Inc 0.0.0.0
R3#
```

No default routes either. Let's move closer to R4 and try to ping from R1:

On R1:

```
R1#Ping 2.2.0.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 2.2.0.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.3, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

R1 can Ping both R3 and R4's Lo0. But why doesn't R1 advertise the route to R3?

Maybe it's NOT an OSPF route. Let's check:

On R1:

```
R1#Show ip route 2.2.0.3

Routing entry for 2.2.0.3/32
  Known via "ospf 1", distance 110, metric 1564, type intra area
  Last update from 2.2.16.6 on Serial1/6, 01:30:51 ago
  Routing Descriptor Blocks:
  * 2.2.16.6, from 2.2.0.3, 01:30:51 ago, via Serial1/6
      Route metric is 1564, traffic share count is 1

R1#Show ip route 2.2.0.4

Routing entry for 2.2.0.4/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 2.2.41.4 on FastEthernet0/1, 01:53:08 ago
  Routing Descriptor Blocks:
    2.2.41.4, from 2.2.0.4, 01:53:08 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1
  * 2.2.14.4, from 2.2.0.4, 01:53:21 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```

The output of the above show commands reveals the following:

- Both routes are OSPF routes
- Both routes belong to the same OSPF process
- Both routes are intra-area OSPF routes

If both interfaces are in the same area then, what could prevent flooding of this OSPF route? Let's check the OSPF interface table:

```
R1#Show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs F/C
Fa0/1	1	4	2.2.41.1/24	1	DR	1/1
Fa0/0	1	4	2.2.14.1/24	1	DR	1/1
Lo0	1	36	2.2.0.1/32	1	LOOP	0/0
Se1/3	1	36	2.2.13.1/24	781	DOWN	0/0

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Se1/6          1      36              2.2.16.1/24         1562  P2P    1/1
```

By looking at the output of the above show command, we can determine the following:

1. R1 is an ABR
2. R1 and R4 are NOT in the same area
3. R1's connection to R3 is in area 36
4. R1's connection to R4 is in area 4

May be routes are filtered between areas? Let's check OSPF's configuration:

```
R1#Show run | s router ospf 1

router ospf 1
```

We don't see any filtering, let's have a closer look at the OSPF interface table:

```
R1#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Fa0/1        1     4               2.2.41.1/24        1     DR    1/1
Fa0/0        1     4               2.2.14.1/24        1     DR    1/1
Lo0          1     36              2.2.0.1/32         1     LOOP  0/0
Se1/3        1     36              2.2.13.1/24        781   DOWN  0/0
Se1/6        1     36              2.2.16.1/24        1562  P2P   1/1
```

R1 is connected to two different areas, 36 and 4. A router that is connected to two or more areas is NOT necessarily an ABR, an ABR is a router that connects at least one NON-ZERO area to area ZERO. As we can see, R1 has NO connection to area 0, therefore, it's NOT an ABR. Since R1 is NOT an ABR, it will NOT send inter-area routes. Let's configure R1 as an ABR, since Loopback 0 is NOT connected to any other router, it's an ideal interface to place in area 0:

```
R1(config)#int lo0
R1(config-if)#ip ospf 1 area 0
```

To verify the configuration:

On R3:

```
R3#Ping 2.2.0.4 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.4, timeout is 2 seconds:
Packet sent with a source address of 2.2.0.3
!!!!!
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/52 ms

R3#Traceroute 2.2.0.4 source Loopback 0 ttl 1 4
Type escape sequence to abort.

Tracing the route to 2.2.0.4
VRF info: (vrf in name/id, vrf out name/id)
  1 2.2.13.1 12 msec 12 msec 16 msec
  2 2.2.14.4 16 msec *  52 msec
```

Ticket 6

Your company is about to change R1 and R4's F0/1 to a 1G link. To prepare for this transition, the BW parameter was configured to 1,000,000 Kb/sec, this was configured to force the traffic flow Via F0/1 interface and NOT F0/0 interface. But the routers (R1 and R4) are still load sharing between the two links. Fix this problem without using an interface configuration command.

Let's check R4's routing table and verify the problem:

On R4:

```
R4#Show ip route 2.2.0.1

Routing entry for 2.2.0.1/32
  Known via "ospf 1", distance 110, metric 2, type inter area
  Last update from 2.2.41.1 on FastEthernet0/1, 00:02:33 ago
  Routing Descriptor Blocks:
    2.2.41.1, from 2.2.0.1, 00:02:33 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1
  * 2.2.14.1, from 2.2.0.1, 00:02:33 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```

The output of the above show command verifies the problem; there are two paths, one through F0/0 and the other through F0/1 interface. Let's check the interfaces and see if the correct BW is configured:

```
R4#Show int F0/0 | Inc BW

  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show int F0/1 | Inc BW

  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 100 usec,
```

We can see that the interfaces are configured correctly. 100 Mb/sec on F0/0 and 1000 Mb/sec on the F0/1 interface. But if the BW of these interfaces are different, then, why do they have a cost of 2?

There are two ways to manipulate OSPF's cost:

- Using the "**IP OSPF cost**" interface configuration command
- Using the "**Auto-cost reference-bandwidth**" router configuration command

By default, OSPF calculates the cost of an interface by dividing the bandwidth of the interface into 100 million. Using the default value, when your network has interfaces with a bandwidth greater than 100 million is not recommended, because OSPF will not be able to differentiate between 100 mbps interface and an interface with a bandwidth that is greater than 100 mbps, "IP OSPF COST" command enables you to change the OSPF cost for an interface, but a better way to accomplish this is to change the default reference value used to calculate the OSPF cost of a given interface. This value can be modified using the "auto-cost reference-bandwidth"router configuration command. If you are planning to use this command, every router in the OSPF routing domain must be configured as well.

Since we are NOT allowed to use an interface configuration command, our ONLY choice is the "Auto-cost reference-bandwidth" router configuration command, as follows:

On R1, R3, R4, R5 and R6

```
Rx(config)#Router ospf 1
Rx(config-router)#auto-cost reference-bandwidth 1000

% OSPF: Reference bandwidth is changed.
        Please ensure reference bandwidth is consistent across all
routers.
```

To verify the configuration:

On R4:

```
R4#Show ip route 2.2.0.1

Routing entry for 2.2.0.1/32
  Known via "ospf 1", distance 110, metric 2, type inter area
  Last update from 2.2.41.1 on FastEthernet0/1, 00:10:50 ago
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing Descriptor Blocks:
* 2.2.41.1, from 2.2.0.1, 00:10:50 ago, via FastEthernet0/1
    Route metric is 2, traffic share count is 1
```

Ticket 7

R1 is configured to load balance traffic for R7's Lo0 using R3 and R6 in BGP, but this is NOT working. DO NOT configure tunnel/s, another routing protocol or remove any commands to accomplish this task.

Let's check the routing table of R1:

On R1:

```
R1#Show ip route bgp | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
B        1.2.3.4 [20/0] via 2.2.16.6, 01:52:03
```

We see one BGP route in R1's routing table, let's check the BGP table of this router:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  1.2.3.4/32       2.2.13.3                               0 3 21 i
*>                  2.2.16.6                               0 56 21 i
```

We do see two BGP paths for 1.2.3.4 prefix. Let's examine the path closely:

```
R1#Show ip bgp 1.2.3.4

BGP routing table entry for 1.2.3.4/32, version 4
Paths: (2 available, best #2, table Default-IP-Routing-Table)
Multipath: eBGP
  Advertised to update-groups:
        1
  3 21
    2.2.13.3 from 2.2.13.3 (2.2.0.3)
      Origin IGP, localpref 100, valid, external
  56 21
    2.2.16.6 from 2.2.16.6 (2.2.0.6)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Origin IGP, localpref 100, valid, external, best
```

The output of the above show command reveals that the prefix was originated in AS 21, which is R7's AS, but the prefix was received from AS 56 from R6 and AS 3 from R3. We also see that the origin, localpref, and origin-type are all the same on both paths, NOT to mention the AS-PATH which is 2 through both paths.

May be BGP is configured to allow ONLY one path in R1's routing table. Let's check R1's BGP configuration:

```
R1#Show run | s router bgp

router bgp 1
 bgp log-neighbor-changes
 neighbor 2.2.13.3 remote-as 3
 neighbor 2.2.16.6 remote-as 56
 maximum-paths 2
```

Well.... that looks fine.
Going through BGP Bestpath process we can see that BGP should have installed two routes in R1's routing table, but it did NOT.
By default, BGP will load balance across EBGP paths ONLY if the AS-PATH is exactly the same.

Since the solution is NOT documented. Let's see if our friend the "?" can help us find the solution:

```
R1(config)#router bgp 1
R1(config-router)#BGP Bestpath as-path ?
% Unrecognized command
R1(config-router)#BGP Bestpath as-path multipath-relax ?
% Unrecognized command
R1(config-router)#BGP Bestpath as-path multipath-relax
```

Well.....The "?" is NOT always our friend, as you can see a hidden undocumented IOS command is used, and Yes it's a fair game. The command instructs IOS to consider ONLY the AS-PATH length and NOT to worry about the actual AS numbers.

For the changes to take effect:

```
R1#Clear ip bgp *
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip route 1.2.3.4
Routing entry for 1.2.3.4/32

  Known via "bgp 1", distance 20, metric 0
  Tag 56, type external
  Last update from 2.2.13.3 00:00:18 ago
  Routing Descriptor Blocks:
  * 2.2.16.6, from 2.2.16.6, 00:00:18 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 56
      MPLS label: none
    2.2.13.3, from 2.2.13.3, 00:00:18 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 56
      MPLS label: none

R1#Show ip route bgp | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
B        1.2.3.4 [20/0] via 2.2.16.6, 00:01:13
                 [20/0] via 2.2.13.3, 00:01:13
```

And to see if Multipath is used in the BGP table:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.2.3.4/32       2.2.16.6                               0 56 21 i
*m                   2.2.13.3                               0 3 21 i
```

Ticket 8

R1's Lo0 interface is configured to join 239.1.1.1 using the "IP Igmp join-group" command. SW1 can NOT ping the group address.

Let's verify the problem:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Ping 239.1.1.1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.1.1.1, timeout is 2 seconds:
.
```

Note the ping is NOT successful

Let's check to see if Multicast-routing is enabled on SW1:

```
SW1#Show ip multicast

  Multicast Routing: enabled
  Multicast Multipath: disabled
  Multicast Route limit: No limit
  Multicast Triggered RPF check: enabled
  Multicast Fallback group mode: Sparse

SW1#Show ip mroute | B Outgoing

Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 02:32:21/00:02:09, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan115, Forward/Sparse, 02:32:19/00:02:09
```

We should verify if PIM is enabled on V115:

On SW1:

```
SW1#Show ip pim interface

Address          Interface            Ver/   Nbr    Query  DR     DR
                                      Mode   Count  Intvl  Prior
2.2.115.11       Vlan115              v2/S   0      30     1      2.2.115.11
```

Since PIM is enabled on V115, let's check and see if R5 is a PIM neighbor:

```
SW1#Show ip pim neighbor

PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Address                                                    Prio/Mode
```

SW1 does NOT have any PIM neighbors. Let's check to see if PIM is enabled on R5's interface facing SW1 (F0/0.115):

On R5:

```
R5#Show run | Inc interface|pim

interface Loopback0
interface FastEthernet0/0
 ip pim sparse-mode
interface FastEthernet0/0.115
 ip pim dr-priority 100
interface FastEthernet0/0.125
interface FastEthernet0/0.145
interface FastEthernet0/1
interface Serial1/0
interface Serial1/1
interface Serial1/2
interface Serial1/3
interface Serial1/4
interface Serial1/5
interface Serial1/6
 ip pim sparse-mode
interface Serial1/7
ip pim autorp listener
```

Note PIM was enabled on the physical interface (F0/0) but NOT on the logical interface (F0/0.115), the "ip pim sparse-mode" command is NOT inherited by the sub-interfaces. The following will enable PIM on F0/0.115:

```
R5(config)#Int F0/0.115
R5(config-subif)#IP PIM sparse-mode
```

We should see the following console messages:

```
%PIM-5-NBRCHG: neighbor 2.2.115.11 UP on interface FastEthernet0/0.115

%PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 2.2.115.5 on interface
FastEthernet0/0.115
```

To verify the configuration:

On R5:

```
R5#Ping 239.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.1.1.1, timeout is 2 seconds:
.
```

Since we are using "Sparse-mode" we should check the rp mapping of SW1:

On SW1:

```
SW1#Show ip pim rp mapping

PIM Group-to-RP Mappings
SW1#
```

Nothing..... What about R5? Does it have any rp mapping? Let's check:

On R5:

```
R5#Show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 2.2.0.4 (?), v2v1
    Info source: 2.2.0.4 (?), elected via Auto-RP
         Uptime: 00:45:30, expires: 00:02:14
```

Let's check how R5's interfaces are configured:

```
R5#Show run | Inc interface|sparse

interface Loopback0
interface FastEthernet0/0
 ip pim sparse-mode
interface FastEthernet0/0.115
 ip pim sparse-mode
interface FastEthernet0/0.125
interface FastEthernet0/0.145
interface FastEthernet0/1
interface Serial1/0
interface Serial1/1
interface Serial1/2
interface Serial1/3
interface Serial1/4
interface Serial1/5
interface Serial1/6
 ip pim sparse-mode
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
interface Serial1/7
```

Note the F0/0.115 interface is NOT configured as "sparse-dense-mode", let's see if the "**IP Pim autorp listener**" is configured:

On R5:

```
R5#Show run | Inc ip pim aut

ip pim autorp listener
```

Since everything looks good, let's check and see if we are experiencing RPF Failures:

On SW1:

```
SW1#Show ip rpf 2.2.0.4

RPF information for ? (2.2.0.4)
  RPF interface: Vlan115
  RPF neighbor: ? (0.0.0.0)
  RPF route/mask: 0.0.0.0/0
  RPF type: unicast (static)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
```

We see a valid RPF entry. Let's check R4's configuration:

On R4:

```
R4#Show run | Inc pim

ip pim sparse-mode
 ip pim sparse-mode
 ip pim sparse-mode
ip pim autorp listener
ip pim send-rp-announce Loopback0 scope 4
ip pim send-rp-discovery Loopback0 scope 4
```

We can see that R4 is sending the RP-discovery messages with a TTL of 4. Let's see the number of hops SW1 takes to reach R4's Lo0:

On SW1:

```
SW1#Traceroute 2.2.0.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Tracing the route to 2.2.0.4

  1 2.2.115.5 0 msec 0 msec 0 msec
  2 2.2.56.1 51 msec 50 msec 50 msec
  3 2.2.236.3 34 msec 25 msec 25 msec
  4 2.2.13.1 25 msec 34 msec 25 msec
  5 2.2.41.4 25 msec *  25 msec
```

Since the "traceroute" states that the lo0 of R4 is 5 hops away, let's change the TTL of the RP-discovery messages to 5 and test:

On R4:

```
R4(config)#ip pim send-rp-discovery lo0 scope 5
```

To verify the configuration:

On SW1:

```
SW1#Show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s) 224.0.0.0/4
  RP 2.2.0.4 (?), v2v1
    Info source: 2.2.0.4 (?), elected via Auto-RP
         Uptime: 00:00:39, expires: 00:02:18

SW1#Ping 239.1.1.1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.1.1.1, timeout is 2 seconds:

Reply to request 0 from 2.2.0.1, 75 ms
```

Ticket 9

R6's Lo6 is configured with an IPv6 address but it's NOT reachable from R5.

Let's find out the IPv6 address of R6's Lo6 interface and verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Show run int lo6 | B interface

interface Loopback6
 no ip address
 ipv6 address FC00:6::6/64
 ipv6 rip RIPNG enable
end
```

On R5:

```
R5#Ping FC00:6::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00:6::6, timeout is 2 seconds:
% No valid source address for destination
Success rate is 0 percent (0/1)
```

Let's check R5's IPv6 routing table:

```
R5#Show ipv6 route

IPv6 Routing Table - 3 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external

C   FC00:56::/64 [0/0]
     via ::, Serial0/0.56
L   FC00:56::5/128 [0/0]
     via ::, Serial0/0.56
L   FF00::/8 [0/0]
     via ::, Null0
```

The route is NOT there. Let's check the connectivity between R5 and R6:

On R6:

```
R6#Show run int S1/5 | B interface

interface Serial1/5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 ip address 2.2.56.1 255.255.255.254
 ip pim sparse-mode
 ip ospf 1 area 36
 ipv6 address FC00:56::6/64
 ipv6 ospf 1 area 56
 mpls ip
end
```

On R5:

```
R5#Ping FC00:56::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00:56::6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

Let's see what routing protocol is used between these two routers:

```
R5#Show ipv6 protocol

IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 1"
  Interfaces (Area 56):
    Serial1/6
  Redistribution:
    None
IPv6 Routing Protocol is "bgp 56"
IPv6 Routing Protocol is "bgp multicast"

R5#Show ipv6 ospf neighbor

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
2.2.0.6           0   FULL/  -        00:00:30    10              Serial1/6
```

Well……R5 and R6 are OSPFv3 neighbors. Let's check R5's OSPFv3 database, maybe the route is in the Database and NOT the routing table:

On R5:

```
R5#Show ipv6 ospf database

            OSPFv3 Router with ID (2.2.0.5) (Process ID 1)

                Router Link States (Area 56)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
ADV Router      Age         Seq#         Fragment ID  Link count  Bits
 2.2.0.5        1908        0x80000006   0            1           None
 2.2.0.6        1361        0x80000009   0            1           E

               Link (Type-8) Link States (Area 56)

ADV Router      Age         Seq#         Link ID    Interface
 2.2.0.5        1908        0x80000005   11         Se1/6
 2.2.0.6        1862        0x80000005   10         Se1/6

               Intra Area Prefix Link States (Area 56)

ADV Router      Age         Seq#         Link ID    Ref-lstype  Ref-LSID
 2.2.0.5        1908        0x80000004   0          0x2001      0
 2.2.0.6        1862        0x80000004   0          0x2001      0
```

The route is NOT in the database. We should check R6's configuration next:

On R6:

```
R6#Show run interface lo6 | B interface

interface Loopback6
 no ip address
 ipv6 address FC00:6::6/64
 ipv6 rip RIPNG enable
end

R6#Show run interface S1/5 | B interface

interface Serial1/5
 ip address 2.2.56.1 255.255.255.254
 ip pim sparse-mode
 ip ospf 1 area 36
 ipv6 address FC00:56::6/64
 ipv6 ospf 1 area 56
 mpls ip
end
```

Lo6 of R6 is advertised in RIPng. Is RIPng redistributed into OSPFv3? Let's check and see:

```
R6#Show run | S ipv6 router

ipv6 router ospf 1
 redistribute rip RIPNG
ipv6 router rip RIPNG
 redistribute ospf 1 metric 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Even though a mutual redistribution was configured, the connected routes were NOT redistributed, this is the default behavior of IPv6 routing protocols and this is the reason that they have included the "include-connected" keyword, let's configure this keyword:

```
R6(config)#ipv6 router ospf 1
R6(config-rtr)#redistribute rip RIPNG include-connected
```

To verify the configuration:

On R5:

```
R5#Show ipv6 route ospf

IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       D - EIGRP, EX - EIGRP external, NM - NEMO, ND - Neighbor Discovery
       l - LISP
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OE2 FC00:6::/64 [110/20]
     via FE80::21C:58FF:FEF6:F7C8, Serial1/6

R5#Ping FC00:6::6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00:6::6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms
```

Ticket 10

When trying to Telnet from R1 to 1.2.3.4 IP address, you notice the following message **on R6**:

```
*Jul 25 22:22:06.325: %IP-4-DUPADDR: Duplicate address
2.2.236.100
```

Let's verify the problem:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Telnet 1.2.3.4
Trying 1.2.3.4 ... Open

Password required, but none set
[Connection to 1.2.3.4 closed by foreign host]
```

Let's check R6's console:

On R6:

The MAC address in your RACK maybe different to the one displayed here:

```
%IP-4-DUPADDR: Duplicate address 2.2.236.100 on FastEthernet0/1, sourced
by 0007.0e1f.db50
```

Sure enough..... we see the message on R6. Let's check the configuration of R6 and look for the 2.2.236.100 IP address:

```
R6#Show run | Inc 2.2.236

ip address 2.2.236.6 255.255.255.0
 neighbor 2.2.236.21 remote-as 21
  primary 2.2.236.6
   peer 2.2.236.3
ip nat pool TST-P 2.2.236.100 2.2.236.100 prefix-length 24
```

We can see that 2.2.236.100 is used in an IP NAT pool called "TST-P". Let's have a look at the NAT configuration on R6:

```
R6#Show run | S nat

ip nat outside
 ip nat inside
ip nat Stateful id 36
  primary 2.2.236.6
   peer 2.2.236.3
   mapping-id 36
ip nat pool TST-P 2.2.236.100 2.2.236.100 prefix-length 24
ip nat inside source route-map TST pool TST-P mapping-id 36 overload
```

The output of the above show command reveals that the Nat pool called "TST-P" was used in a Stateful NAT. Let's check the MAC address that was displayed in the console message of R6:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show mac address-table | Inc 0007.0e1f.db50

 236    0007.0e1f.db50    DYNAMIC     Fa0/3
```

This is R3, let's check R3's NAT configuration:

On R3:

You should see the following message on R3's console as well:

```
%IP-4-DUPADDR: Duplicate address 2.2.236.100 on FastEthernet0/0, sourced
by 001c.58f6.f7c9

R3#Show run | s nat

 ip nat inside
 ip nat outside
 ip nat inside
ip nat Stateful id 36
  primary 2.2.236.3
   peer 2.2.236.6
   mapping-id 36
ip nat pool TST-P 2.2.236.100 2.2.236.100 prefix-length 24
ip nat inside source route-map TST pool TST-P mapping-id 36 overload
```

We can see a similar error message on R3. We also see that 2.2.236.100 IP address is also used in a NAT pool on R3, and the final important point here is that R3 is configured with Stateful NAT.

In Stateful NAT ONLY one router should allocate IP addresses from the Pool, let's check the NAT configuration of R3 and R6:

On R3:

```
R3#Sh run | S State

ip nat Stateful id 36
  primary 2.2.236.3
   peer 2.2.236.6
   mapping-id 36
```

On R6:

```
R6#Show run | S State

ip nat Stateful id 36
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
primary 2.2.236.6
 peer 2.2.236.3
 mapping-id 36
```

The configuration of R3 and R6 reveals that both routers are configured as primary. Let's correct this configuration:

On R6:

```
R6(config)#No ip nat stateful id 36
R6(config)#ip nat stateful id 36
R6(config-ipnat-snat)#backup 2.2.236.6
R6(config-ipnat-snat-bkp)#peer 2.2.236.3
R6(config-ipnat-snat-bkp)#mapping-id 36
```

We should see the following console messages immediately:

```
%SNAT-5-PROCESS: Id 36, System starts converging

SNAT (Receive): CONVERGENCE Message from Router-Id: 36 for Router-Id: 0's
entries

%SNAT-5-PROCESS: Id 36, System fully converged
```

Let's try to Telnet and see if this configuration change fixed the problem:

On R3 and R6:

```
Rx#Clear ip nat trans *
```

On R1

```
R1#Telnet 1.2.3.4
Trying 1.2.3.4 ... Open
Password required, but none set

[Connection to 1.2.3.4 closed by foreign host]
```

You may have to wait for a while for the error messages to go away

On R6:

No error messages

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

No error messages

Ticket 11

SW3 and SW4 are configured to exchange RIP routes, SW3 receives the RIP routes from SW4, but SW4 does NOT receive any RIP routes from SW3.

Based on ticket 3, we know that SW3 and R6 are connected through a VRF configuration; let's check R5's VRF configuration to see if it's connected to SW4:

On R5:

```
R5#Show ip vrf
  Name                              Default RD          Interfaces
  34                                56:5                Fa0/0.145
```

SW4 is connected to a VRF on R5, we can also see that R5 has ONLY one interface configured as a VRF interface, therefore, the ONLY way to connect between the two VRFs on different routers without having a special interface connecting the VRFs (VRF-Lite) is MPLS VPN.

Let's check SW3 and SW4's routing table:

On SW3:

```
SW3#Show ip route rip

     2.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
R       2.2.0.14/32 [120/1] via 2.2.136.6, 00:00:03, Vlan136
R       2.2.145.0/24 [120/1] via 2.2.136.6, 00:00:03, Vlan136
```

On SW4:

```
SW4#Show ip route rip
SW4#
```

We can see RIP routes on SW3 but not on SW4. Let's check SW4's configuration:

On SW4:

```
SW4#Show ip protocols
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Sending updates every 30 seconds, next due in 3 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface              Send  Recv  Triggered RIP  Key-chain
    Vlan145                2     2
    Loopback0              2     2
  (The rest of the output is omitted for brevity)
```

The above output shows that RIPv2 is enabled on all Layer 3 interfaces. Let's check and see the connectivity between SW4 and R5:

On SW4:

```
SW4#Ping 2.2.145.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.145.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R5:

```
R5#Show ip route vrf 34 | B Gateway
Gateway of last resort is not set

      2.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R        2.2.0.14/32 [120/1] via 2.2.145.14, 00:00:02, FastEthernet0/0.145
C        2.2.145.0/24 is directly connected, FastEthernet0/0.145
L        2.2.145.5/32 is directly connected, FastEthernet0/0.145
```

We do not see any IBGP route/s from R6 to advertise to SW4. Let's check R5's BGP table for VPNv4 routes:

```
R5#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 56:5 (default for vrf 34)
*> 2.2.0.14/32      2.2.145.14               1         32768 ?
*> 2.2.145.0/24     0.0.0.0                  0         32768 ?
```

We can ONLY see the routes coming from SW4 and nothing from SW3. Let's check R6's VPNv4 table:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 56:5
*>i2.2.0.14/32      2.2.0.5                  1    100      0 ?
*>i2.2.145.0/24     2.2.0.5                  0    100      0 ?
Route Distinguisher: 56:6 (default for vrf 34)
*> 2.2.0.13/32      2.2.136.13               1          32768 ?
*>i2.2.0.14/32      2.2.0.5                  1    100      0 ?
*> 2.2.136.0/24     0.0.0.0                  0          32768 ?
*>i2.2.145.0/24     2.2.0.5                  0    100      0 ?
```

R6 has the routes for SW3. Let's enable BGP debugging on R5 and configure R6 to resend its routes to R5 using the "**Clear ip bgp * in**" command.

On R5:

```
R5#Debug ip bgp vpnv4 unicast updates
BGP updates debugging is on for address family: VPNv4 Unicast

R5#Clear ip bgp * in

BGP(4): 2.2.0.6 rcvd UPDATE w/ attr: nexthop 2.2.0.6, origin ?, localpref
100, metric 1, extended community RT:56:6
BGP(4): 2.2.0.6 rcvd 56:6:2.2.0.13/32, label 30 -- DENIED due to:
extended community not supported;
BGP(4): 2.2.0.6 rcvd UPDATE w/ attr: nexthop 2.2.0.6, origin ?, localpref
100, metric 0, extended community RT:56:6
BGP(4): 2.2.0.6 rcvd 56:6:2.2.136.0/24, label 16 -- DENIED due to:
extended community not supported;
```

The output of the above debug command reveals that R5 receives the routes advertised by SW3 via R6, but they are rejected because of unsupported extended community attribute of 56:6. Let's disable the debug and check R6 and R5's VRF configuration:

```
R5#Undebug all
All possible debugging has been turned off
```

On R6:

```
R6#Show run | S vrf

ip vrf 34
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 rd 56:6
 route-target export 56:6
 route-target import 56:5
 ip vrf forwarding 34
 address-family ipv4 vrf 34
  timers basic 5 10 10 15
  redistribute bgp 56 metric 1
  network 2.0.0.0
  no auto-summary
  version 2
 address-family ipv4 vrf 34
  redistribute rip
```

On R5:

```
R5#Show run | S vrf

ip vrf 34
 rd 56:5
 route-target export 56:5
 route-target import 56:5
 ip vrf forwarding 34
 address-family ipv4 vrf 34
  redistribute bgp 56 metric 1
  network 2.0.0.0
  no auto-summary
  version 2
 address-family ipv4 vrf 34
  redistribute rip
```

The output of the above show command revealed the problem; R5 is importing an RT of 56:5 and NOT 56:6. Let's fix this problem:

On R5:

```
R5(config)#ip vrf 34
R5(config-vrf)#route-target import 56:6
```

To verify the configuration:

On SW4:

```
SW4#Show ip route rip

     2.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        2.2.0.13/32 [120/1] via 2.2.145.5, 00:00:04, Vlan145
R        2.2.136.0/24 [120/1] via 2.2.145.5, 00:00:04, Vlan145
```

To test the configuration:

On SW4:

```
SW4#Ping 2.2.0.13

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.0.13, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms
```

Ticket 12

R3 was configured to ONLY allow authenticated users to Telnet to R8's Lo0's IP address (4.3.2.1 /32). Users should first authenticate to R3, if successful, then, they should be able to Telnet to R8's Lo0. Users are authenticating to R3 successfully, but still can't Telnet to R8.

Let's verify the problem:

On R4:

```
R4#Telnet 4.3.2.1

Trying 4.3.2.1 ...
% Destination unreachable; gateway or host down
```

Let's try to Telnet to R3 and then to R8:

On R4:

```
R4#Telnet 2.2.0.3

Trying 2.2.0.3 ... Open

User Access Verification

Username: CCIE
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Password: CCIE

R3>Exit

R4#Telnet 4.3.2.1

Trying 4.3.2.1 ...
% Destination unreachable; gateway or host down
```

This method of first authenticating to R3 using Telnet so we can Telnet to R8 is called "Lock and Key" or Dynamic Access-lists. Let's check R3's configuration:

On R3:

```
R3#Show run | Inc access|interface

interface Loopback0
interface FastEthernet0/0
 ip access-group 133 in
interface FastEthernet0/1
interface Serial1/0
interface Serial1/1
 ip access-group 133 in
interface Serial1/2
interface Serial1/3
interface Serial1/4
interface Serial1/5
interface Serial1/6
interface Serial1/7
access-list 133 dynamic A permit tcp any host 4.3.2.1 eq telnet
access-list 133 deny   tcp any host 4.3.2.1 eq telnet
access-list 133 permit ip any any
access-list 136 permit ip 2.2.0.0 0.0.255.255 host 1.2.3.4
```

We can see that the dynamic Access-list was configured to allow Telnet to R8 and the Access-list was applied to all R3's transit interfaces. Next, let's check the configuration of "username" and the vty lines:

On R3:

```
R3#Show run | S username|line vty

username CCIE password 0 CCIE
line vty 0 4
 login local
 transport input all
```

We do NOT see the command that triggers the Lock & Key to open a hole in the dynamic access-list. Let's

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

configure this for the user called "CCIE":

On R3:

```
R3(config)#Username CCIE autocommand access-enable host
```

To verify and test the configuration:

On R4:

```
R4#Telnet 2.2.0.3

Trying 2.2.0.3 ... Open

User Access Verification

Username: CCIE
Password: CCIE

[Connection to 2.2.0.3 closed by foreign host]

R4#Telnet 4.3.2.1

Trying 4.3.2.1 ... Open

User Access Verification

Username: CCIE
Password: CCIE
R8>exit

[Connection to 4.3.2.1 closed by foreign host]
```

To see the dynamically configured access-list on R3:

On R3:

```
R3#Show ip access-list 133

Extended IP access list 133
    10 Dynamic A permit tcp any host 4.3.2.1 eq telnet
       permit tcp host 2.2.41.4 host 4.3.2.1 eq telnet (51 matches)
    20 deny tcp any host 4.3.2.1 eq telnet (6 matches)
    30 permit ip any any (34685 matches)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 13

Users complain that telnet sessions from R4's Lo0 to R7 are very slow. Using a sniffer you discover that packets are reaching R7 out of sequence. Fix the problem, while ensuring that R1 continues to load Balance sessions between it's S1/6 and S1/3 interfaces for different services.

Based on the problem and the topology, R1 is using the paths through R3 and R6 to reach R7, and this could be the problem because maybe they have different BWs. But before removing or adding any configuration, let's test and verify how Telnet traffic is load balanced:

On R1:

```
R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/6

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/6
```

We can see that R1 is performing per packet load balancing, because in the last two tests the same source port was used. Let's check the configuration on R1:

On R1:

```
R1#Show run int S1/3 | Inc ip load

 ip load-sharing per-packet

R1#Show run int S1/6 | Inc ip load

 ip load-sharing per-packet
```

We can see the problem, the two interfaces on R1 are configured to perform per packet load balancing, let's remove the commands and see the result:

```
R1(config)#Int s1/3
R1(config-if)#No ip load-sharing per-packet
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Int s1/6
R1(config-subif)#No ip load-sharing per-packet
```

On R1:

```
R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3
```

Obviously the solution worked, because packets of the same session will go through the same path, and as a result they will NOT arrive out of order. But remember that the ticket's requirement was that R1 should still perform load balancing, let's check the configured load balancing on R1:

```
R1#Show run | Inc cef

ip cef
ip cef load-sharing algorithm include-ports destination
```

Since R1 is configured to perform load-sharing per destination port, But all Telnet sessions use the same destination port (23), and that is why R1 chooses ONLY one path. The following shows how different destination ports use different paths:

```
R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 23
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/3

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 80
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/6

R1#Show ip cef exact-route 2.2.0.4 src-port 1111 1.2.3.4 dest-port 80
2.2.0.4 -> 1.2.3.4 => IP adj out of Serial1/6
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 14

Erase the startup configuration on the Routers and the Switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Advanced Troubleshooting CCIE Routing & Switching

v5.0

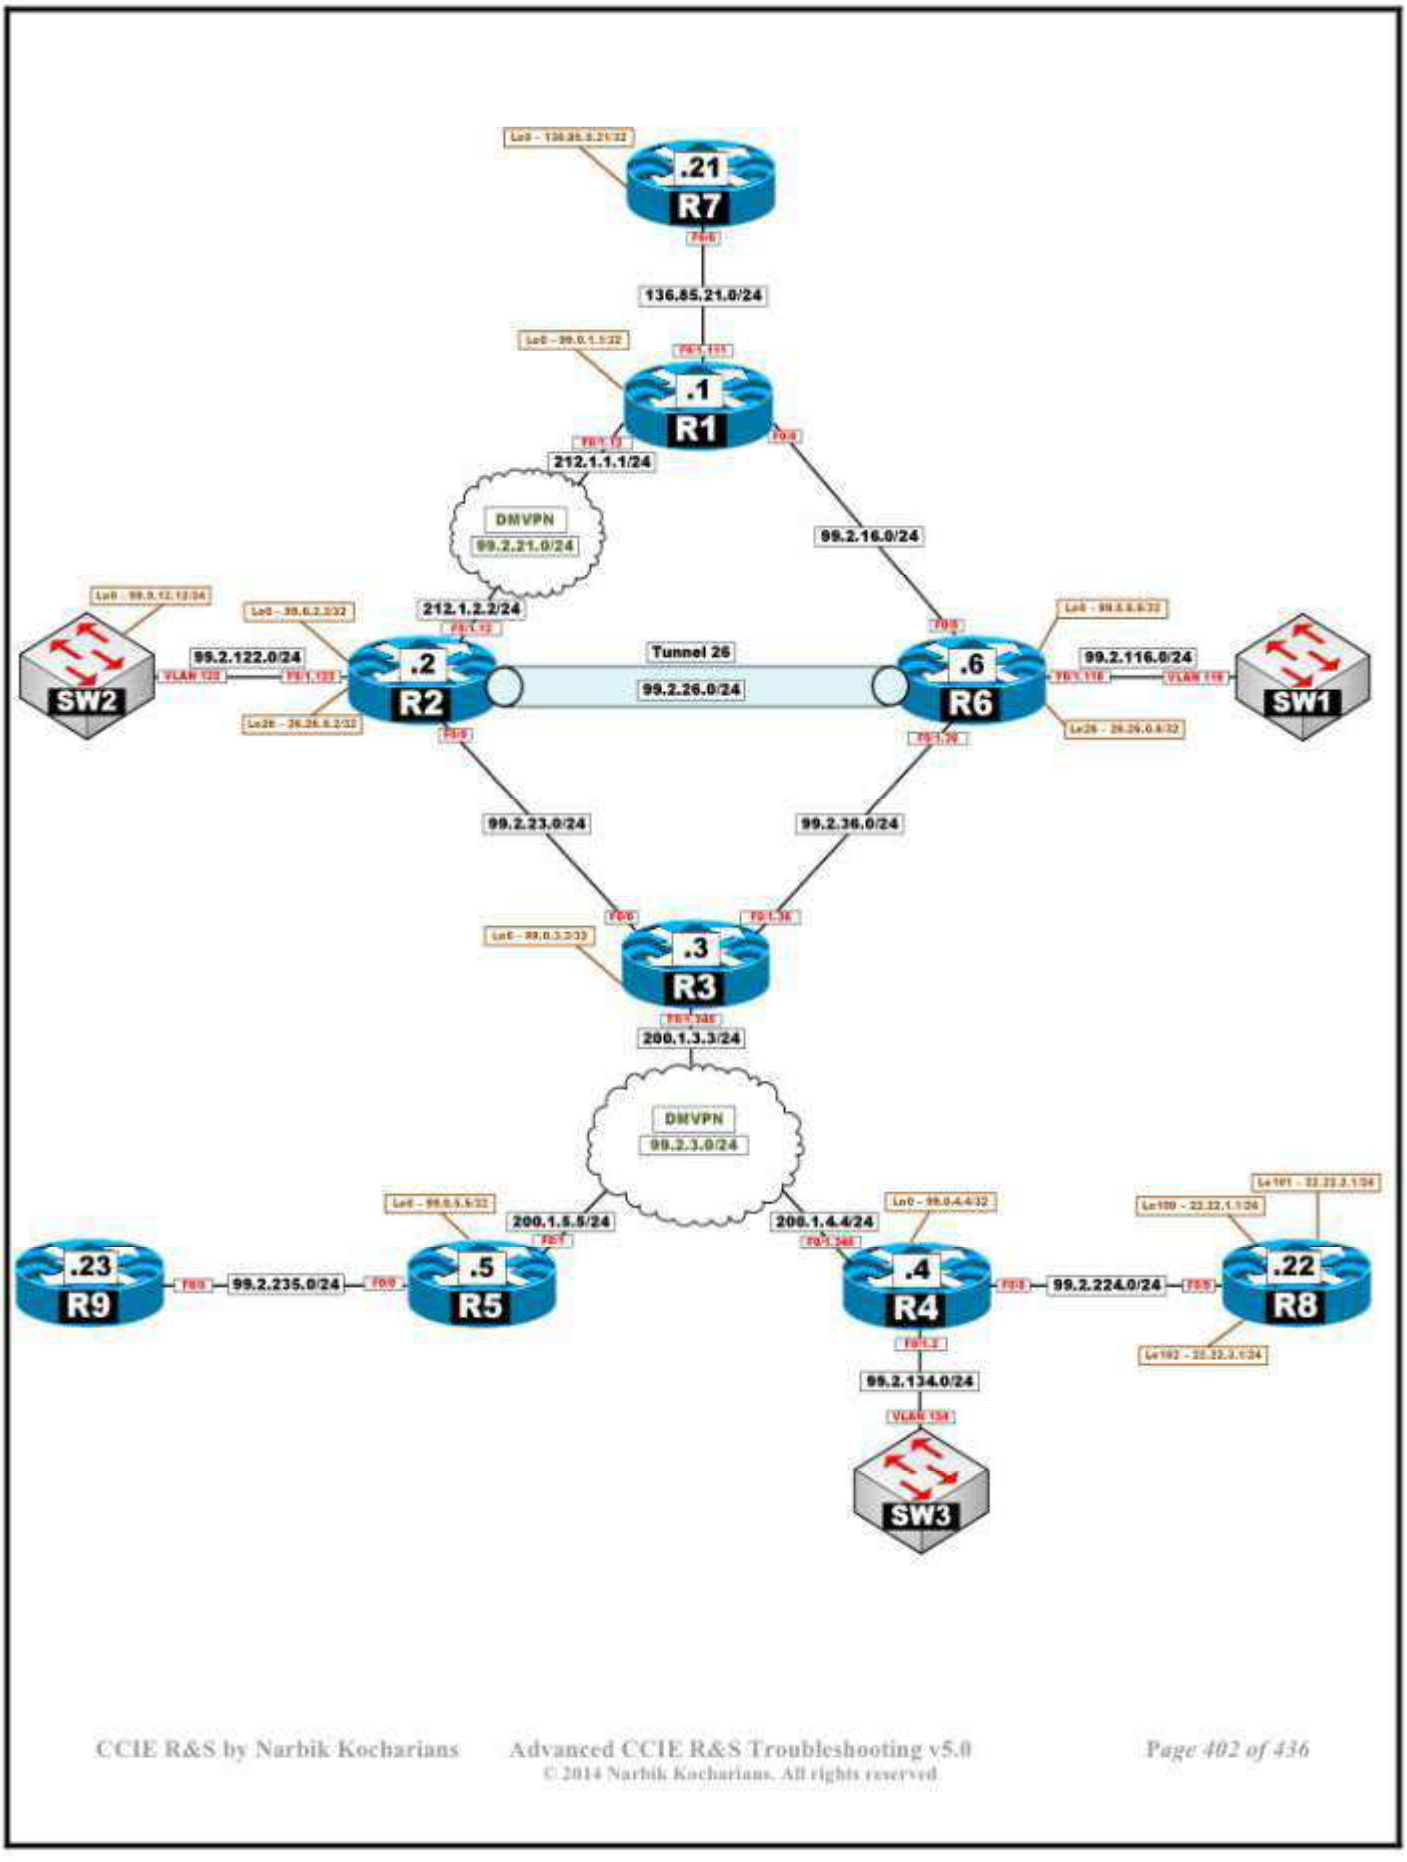
www.MicronicsTraining.com

Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

TS-Mock-Lab-5

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Rules:

- No Static routes, No default routes, NO Policy Based Routing is allowed, unless otherwise specified.

Lab Setup:

To copy and paste the initial configurations, go to "CCIE-TS-Initial-configurations" folder → "TS-Mock-Labs" →"TS-Mock-Lab-5".

Ticket 1

The CPU utilization of the Switches is exceeding 80 percent. HELP.

Let's verify the problem:

On SW1:

```
SW1#Show process cpu | Inc utilization

CPU utilization for five seconds: 99%/13%; one minute: 99%; five minutes:
98%
```

On SW2:

```
SW2#Show process cpu | Inc utilization

CPU utilization for five seconds: 73%/43%; one minute: 72%; five minutes:
75%
```

On SW3:

```
SW3#Show process cpu | Inc utilization

CPU utilization for five seconds: 94%/92%; one minute: 94%; five minutes:
95%
```

We can see the problem, typically this is caused by Spanning-tree loops, let's check the port status of VLAN 1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1#Show spanning-tree vlan 1 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/10           Desg FWD 19        128.12   P2p
Fa0/17           Desg FWD 19        128.19   P2p
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/21           Desg FWD 19        128.23   P2p
Fa0/23           Root FWD 19        128.25   P2p
```

On SW2:

```
SW2#Show spanning-tree vlan 1 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/19           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.25   P2p
```

On SW3:

```
SW3#Show spanning-tree vlan 1 | B Interface

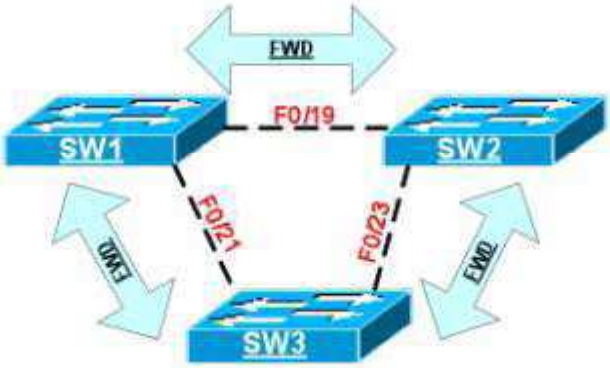
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/21           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.23   P2p
```

As we can see all ports are in forwarding state; SW1 is connected to SW2 through F0/19 and these ports are in forwarding state on both switches. SW2 is connected to SW3 through F0/23, and these ports are also in forwarding state, and SW3 is connected to SW1 through its F0/21 interface and it is also in forwarding state.

Let's draw a diagram so we can se it better:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



Let's check the switches for Spanning-tree configuration:

On SW1:

```
SW1#Show run | Inc spanning

spanning-tree mode pvst
spanning-tree extend system-id
```

We don't see any configuration on SW1 that could cause this problem.

On SW2:

```
SW2#Show run | Inc spanning

spanning-tree mode pvst
spanning-tree extend system-id
```

SW2 seems to be ok.

On SW3

```
SW3#Show run | Inc spanning

spanning-tree mode pvst
spanning-tree extend system-id
 spanning-tree bpdufilter enable
```

We can see that "BPDUFilter" is configured on this switch, let's find the port that this configuration is applied to:

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3#Show run | Inc interface|spanning

spanning-tree mode pvst
spanning-tree extend system-id
interface FastEthernet0/1
interface FastEthernet0/2
interface FastEthernet0/3
interface FastEthernet0/4
interface FastEthernet0/5
interface FastEthernet0/6
interface FastEthernet0/7
interface FastEthernet0/8
interface FastEthernet0/9
interface FastEthernet0/10
interface FastEthernet0/11
interface FastEthernet0/12
interface FastEthernet0/13
interface FastEthernet0/14
interface FastEthernet0/15
interface FastEthernet0/16
interface FastEthernet0/17
interface FastEthernet0/18
interface FastEthernet0/19
interface FastEthernet0/20
interface FastEthernet0/21
interface FastEthernet0/22
interface FastEthernet0/23
 spanning-tree bpdufilter enable
interface FastEthernet0/24
interface GigabitEthernet0/1
interface GigabitEthernet0/2
interface Vlan1
interface Vlan134
```

It's configured on the F0/23 interface, which is the trunk link between SW2 and SW3. It's the "BPDUFILTER" command that is causing this problem; we have few options to resolve this problem, one easy way to resolve the problem is to disable "BPDUFilter" on the F0/23 interface of SW3:

On SW3:

```
SW3(config)#int f0/23
SW3(config-if)#Spanning-tree bpdufilter disable

SW3#Show process cpu | Inc utilization
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
CPU utilization for five seconds: 0%/0%; one minute: 21%; five minutes:71%
```

Ticket 2

SW1 was configured to isolate R1 and R6 from the rest of the ports in VLAN 16; however, R1 can NOT establish an EIGRP session with R6. DO NOT remove any command/s to resolve this ticket.

Let's verify the problem:

On R1:

```
R1#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(3)
```

On R6:

```
R6#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(3)
```

Let's check and see if these two routers have basic reachability:

On R1:

```
R1#Ping 99.2.16.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 99.2.16.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Since SW1 was configured to isolate these two routers from the rest of the ports in Vlan 16, let's see how SW1 was configured to do that:

On SW1:

```
SW1#Show run int f0/1 | B inter

interface FastEthernet0/1
 switchport access vlan 16
 switchport private-vlan host-association 16 161
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
switchport mode private-vlan host
```

```
SW1#Show run int f0/6 | B inter

interface FastEthernet0/6
 switchport access vlan 16
 switchport private-vlan host-association 16 161
 switchport mode private-vlan host
```

We can see that PVLAN is configured on these ports, with VLAN 16 as the primary and VLAN 161 as the secondary. Both ports are configured as host ports, which mean that they are part of a secondary VLAN. Maybe VLAN 161 is an isolated secondary VLAN, which means that the hosts that are member of this VLAN can NOT communicate with each other.

Let's check:

On SW1:

```
SW1#Show vlan private-vlan

Primary Secondary Type              Ports
------- --------- ----------------- -----------------
16      161       isolated          Fa0/1, Fa0/6
```

Sure enough, they are in Isolated VLAN. Let's fix this problem by changing them to a Community secondary VLAN:

On SW1:

```
SW1(config)#vlan 161
SW1(config-vlan)#private-vlan Community
SW1(config-vlan)#Exit
```

Once the community Secondary VLAN is configured, you should see the following console messages on R1 and/or R6:

On R1 or R6:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 3: Neighbor 99.2.16.6 (FastEthernet0/0) is
up: new adjacency
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip eigrp neighbors

IP-EIGRP neighbors for process 3
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   99.2.16.6               Fa0/0             10 00:02:33    9   200  0  14
```

Ticket 3

R6 can't ping 99.2.26.2. You must configure R6 to fix this problem.

Let's verify the problem:

On R6:

```
R6#Ping 99.2.26.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 99.2.26.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check R6's routing table:

On R6:

```
R6#Show ip route 99.2.26.2

Routing entry for 99.2.26.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Tunnel26
      Route metric is 0, traffic share count is 1
```

99.2.26.2 is directly connected via tunnel 26 interface. Let's check the status of this tunnel:

```
R6#Show int tunnel 26

Tunnel26 is up, line protocol is up
  Hardware is Tunnel
  Internet address is 99.2.26.6/24
  MTU 17912 bytes, BW 100 Kbit/sec, DLY 50000 usec,
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Tunnel TTL 255, Fast tunneling enabled
  Tunnel transport MTU 1472 bytes
  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps)
```

Based on the output of the above show command we can see that the tunnel key values don't match, the Tunnel key value of R2 is set to a hex value of 0X328 which translates to a decimal value of 808.

Let's fix this problem by configuring the tunnel key value of R6:

On R6:

```
R6(config)#Int tunnel 26
R6(config-if)#Tunnel key 808
```

You should see the following console message:

```
%PIM-5-NBRCHG: neighbor 99.2.26.2 UP on interface Tunnel26
```

To test and verify the configuration:

On R6:

```
R6#Ping 99.2.26.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 99.2.26.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Ticket 4

R1 can't ping R7's G0/0 interface using its Lo0 as the source. You should configure R1 to fix this problem.

Let's verify the problem:

On R1:

```
R1#Ping 136.85.21.21 source Lo0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.85.21.21, timeout is 2 seconds:
Packet sent with a source address of 99.0.1.1
.....
Success rate is 0 percent (0/5)
```

Since these two routers are directly connected, let's check and see if they have reachability to each other:

On R1:

```
R1#Ping 136.85.21.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.85.21.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Looking at the diagram, we can see that these two routers should be exchanging routes using RIPv2 routing protocol. Let's check and see if the correct networks are advertised in RIP:

On R1:

```
R1#Show ip protocol | s rip

Redistributing: 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 25 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: eigrp 3, rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/1.111   1     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    136.85.0.0
    136.85.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    136.85.21.21         200      00:00:25
  Distance: (default is 200)
  Routing Information Sources:
    Gateway         Distance      Last Update
    136.85.21.21         200      00:00:20
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Distance: (default is 200)
```

We can see that RIP is enabled on the F0/1.111 interface of R1 and NOT the Lo0. Let's check the configuration of RIP to see if it was redistributed from EIGRP:

```
R1#Show run | S router rip

router rip
 redistribute eigrp 3 metric 1

 network 136.85.0.0
 distance 200
 no auto-summary
 version 2
```

Since it's redistributed into RIP, let's check RIP's database:

On R1:

```
R1#Show ip rip database 99.0.1.1 255.255.255.255

99.0.1.1/32    redistributed
    [1] via 0.0.0.0,
```

This should work. Let's enable debugging.

On R1:

```
R1#Debug ip rip
RIP protocol debugging is on

RIP: sending v1 update to 255.255.255.255 via FastEthernet0/1.111
(136.85.21.1)
RIP: build update entries
network 99.0.0.0 metric 1
```

We can see that R1 is sending Version 1 updates to 255.255.255.255, let's check RIP's configuration again:

On R1:

```
R1#Show run | S router rip

router rip
 version 2
 redistribute eigrp 3 metric 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 network 136.85.0.0
 distance 200
 no auto-summary

R1#Show run | Inc interface|version

version 15.1
interface Loopback0
interface Tunnel12
interface FastEthernet0/0
interface FastEthernet0/1
interface FastEthernet0/1.12
interface FastEthernet0/1.111
 ip rip send version 1
interface Serial1/0
interface Serial1/1
interface Serial1/2
interface Serial1/3
interface Serial1/4
interface Serial1/5
interface Serial1/6
interface Serial1/7
 version 2
```

We can see the problem. R1 is sending RIP version 1 updates out of its F0/1.111 interface. Let's remove this command and test reachability using ping:

On R1:

```
R1(config)#int f0/1.111
R1(config-if)#IP rip send version 2
```

To test the configuration:

On R1:

```
R1#Ping 136.85.21.21 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.85.21.21, timeout is 2 seconds:
Packet sent with a source address of 99.0.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

If the ping is unsuccessful, do a "**Clear ip route ***" to implement the changes.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Ticket 5

R1 and R2 can't establish an EIGRP session; you should use an EIGRP command to fix this problem.

Let's verify the problem:

On R1:

```
R1#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(3)

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   99.2.16.6               Fa0/0             14 20:21:43    5   200  0  18
```

On R2:

```
R2#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(3)
```

R2 has no neighbors at all. Let's see if EIGRP is enabled on the tunnel interfaces of R1 and R2:

On R1:

```
R1#Show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(3)

                    Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable  Flow Timer   Routes
Fa0/0          1        0/0        1        0/1           50          0
Lo0            0        0/0        0        0/1            0          0
Tu12           0        0/0        0        6/6            0          0
```

On R2:

```
R2#Show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(3)

                    Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable  Flow Timer   Routes
Tu12           0        0/0        0        6/6            0          0
```

Since everything is configured correctly, let's see if they have reachability:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Ping 99.2.21.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 99.2.21.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms
```

Let's check the configuration of these routers, maybe authentication is configured incorrectly on these interfaces:

On R1:

```
R1#Show run int tunnel 12 | B interface

interface Tunnel12
 ip address 99.2.21.1 255.255.255.0
 no ip redirects
 ip nhrp map 99.2.21.2 212.1.2.2
 ip nhrp network-id 111
 tunnel source FastEthernet0/1.12
 tunnel mode gre multipoint
end
```

No authentication.

On R2:

```
R2#Show run int tunnel 12 | B interface

interface Tunnel12
 ip address 99.2.21.2 255.255.255.0
 no ip redirects
 ip nhrp map 99.2.21.1 212.1.1.1
 ip nhrp network-id 222
 tunnel source FastEthernet0/1.12
 tunnel mode gre multipoint
end
```

No authentication here either. But we can see that the tunnel 12 interfaces of R1 and R2 are missing the mapping for Multicast, and we know that the EIGRP spaeaking routers communicate using 224.0.0.10; to fix this problem we can simply map "**Multicast**", but the task specifically states that an EIGRP command must be used to fix this problem. Therefore, the "**Neighbor**" command can be used to change the default behavior which is Multicast Hellos to Unicast Hellos:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#router eigrp 3
R1(config-router)#Neighbor 99.2.21.2 tunnel 12
```

On R2:

```
R2(config)#router eigrp 3
R2(config-router)#Neighbor 99.2.21.1 tunnel 12
```

We should see the following console message:

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 3: Neighbor 99.2.12.1 (Serial0/0) is up:
new adjacency
```

To test and verify the configuration:

On R2:

```
R2#Show ip eigrp neighbors

IP-EIGRP neighbors for process 3
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   99.2.12.1               Tu12            131 00:00:51   36   486  0  24
```

Ticket 6

R2 was configured to filter SW1 and R3's Lo0 from being sent to SW2; however SW2 still sees SW1's Lo0.

Let's verify the problem:

On SW2:

```
SW2#Show ip route ospf | B Gate
SW2#
```

No routes at all, let's verify the neighbor table:

```
SW2#Sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address         Interface
99.0.2.2          1   EXSTART/DR      00:00:38    99.2.122.2      Vlan122
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#router eigrp 3
R1(config-router)#Neighbor 99.2.21.2 tunnel 12
```

On R2:

```
R2(config)#router eigrp 3
R2(config-router)#Neighbor 99.2.21.1 tunnel 12
```

We should see the following console message:

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 3: Neighbor 99.2.12.1 (Serial0/0) is up:
new adjacency
```

To test and verify the configuration:

On R2:

```
R2#Show ip eigrp neighbors

IP-EIGRP neighbors for process 3
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   99.2.12.1               Tu12            131 00:00:51   36   486  0  24
```

Ticket 6

R2 was configured to filter SW1 and R3's Lo0 from being sent to SW2; however SW2 still sees SW1's Lo0.

Let's verify the problem:

On SW2:

```
SW2#Show ip route ospf | B Gate
SW2#
```

No routes at all, let's verify the neighbor table:

```
SW2#Sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address         Interface
99.0.2.2          1   EXSTART/DR      00:00:38    99.2.122.2      Vlan122
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's enable debugging and see what's going on:

On R2:

```
R2#Debug ip ospf adj
OSPF adjacency debugging is on

OSPF-1 ADJ   Fa0/1.122: Elect BDR 99.0.2.2
OSPF-1 ADJ   Fa0/1.122: Elect DR 99.0.12.12
OSPF-1 ADJ   Fa0/1.122: DR: 99.0.12.12 (Id)   BDR: 99.0.2.2 (Id)
OSPF-1 ADJ   Fa0/1.122: Rcv DBD from 99.0.12.12 seq 0xACA opt 0x52 flag
0x7 len 32  mtu 1504 state EXSTART
OSPF-1 ADJ   Fa0/1.122: Nbr 99.0.12.12 has larger interface MTU
OSPF-1 ADJ   Fa0/1.122: Reset old DR
```

It is the MTU, the output of the debug states that SW2 has a larger MTU. Let's fix this problem:

On SW2:

```
SW2#Show system mtu

System MTU size is 1504 bytes
System Jumbo MTU size is 1504 bytes
Routing MTU size is 1504 bytes

SW2(config)#System mtu routing 1500
```

We should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 99.0.2.2 on Vlan122 from LOADING to FULL,
Loading Done
```

Now, let's verify the routing table of SW2:

```
SW2#Show ip route ospf

     136.85.0.0/16 is variably subnetted, 2 subnets, 2 masks
O E2    136.85.21.0/24 [110/20] via 99.2.122.2, 00:00:58, Vlan122
O E2    136.85.0.21/32 [110/20] via 99.2.122.2, 00:00:58, Vlan122
     99.0.0.0/8 is variably subnetted, 13 subnets, 2 masks
O IA    99.2.116.0/24 [110/4] via 99.2.122.2, 00:00:58, Vlan122
O IA    99.2.36.0/24 [110/3] via 99.2.122.2, 00:00:58, Vlan122
O IA    99.2.3.3/32 [110/2] via 99.2.122.2, 00:00:58, Vlan122
O IA    99.0.11.11/32 [110/5] via 99.2.122.2, 00:00:58, Vlan122
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O IA     99.0.6.6/32 [110/4] via 99.2.122.2, 00:00:59, Vlan122
O IA     99.0.2.2/32 [110/2] via 99.2.122.2, 00:00:59, Vlan122
O E2     99.0.1.1/32 [110/20] via 99.2.122.2, 00:00:59, Vlan122
O E2     99.0.6.0/24 [110/20] via 99.2.122.2, 00:00:59, Vlan122
O E2     99.2.16.0/24 [110/20] via 99.2.122.2, 00:00:59, Vlan122
O E2     99.2.21.0/24 [110/20] via 99.2.122.2, 00:00:59, Vlan122
O IA     99.2.23.0/24 [110/2] via 99.2.122.2, 00:00:59, Vlan122
```

Let's check the filtering configured on R2:

On R2:

```
R2#Show run | s router ospf

router ospf 1
 log-adjacency-changes

 area 0 range 99.0.3.3 255.255.255.255 not-advertise
 area 0 range 99.0.11.11 255.255.255.255 not-advertise
 redistribute eigrp 3 subnets
```

The "area range Not-advertise" command is a filtering mechanism used to filter intra-area and NOT inter-area routes. To verify this information:

On R2:

```
R2#Show ip route 99.0.3.3

Routing entry for 99.0.3.3/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Redistributing via eigrp 3
  Advertised by eigrp 3 metric 1 1 1 1 1
  Last update from 99.2.23.3 on FastEthernet0/0, 22:22:22 ago
  Routing Descriptor Blocks:
  * 99.2.23.3, from 99.0.3.3, 22:22:22 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```

This is the reason that the "area range Not-advertise" command worked for 99.0.3.3/32 prefix, because it's a intra-area route. Let's check the 99.0.11.11/32 prefix:

```
R2#Show ip route 99.0.11.11

Routing entry for 99.0.11.11/32
  Known via "ospf 1", distance 110, metric 4, type inter area
  Redistributing via eigrp 3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Advertised by eigrp 3 metric 1 1 1 1 1
  Last update from 99.2.23.3 on FastEthernet0/0, 20:39:57 ago
  Routing Descriptor Blocks:
  * 99.2.23.3, from 99.0.6.6, 20:39:57 ago, via FastEthernet0/0
      Route metric is 4, traffic share count is 1
```

Now we know why the "area range Not-advertise" command did not work for this prefix, because this prefix is an inter-area route. To filter inter-area prefixes LSA type 3 filtering mechanism must be used, which is implemented using the "area filter-list" command. Let's fix this problem:

On R2:

```
R2(config)#IP prefix-list TST seq 5 deny 99.0.11.11/32
R2(config)#IP prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R2(config)#router ospf 1
R2(config-router)#Area 0 filter-list prefix TST out
```

To test and verify the configuration:

On SW2:

```
SW2#Sh ip route 99.0.11.11

% Subnet not in table
```

Ticket 7

R3 can't ping R7's Lo0 (136.85.0.21/32).

Let's verify the problem:

On R3:

```
R3#Ping 136.85.0.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.85.0.21, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check R1's routing table for this prefix:

On R1:

```
R1#Show ip route 136.85.0.21

Routing entry for 136.85.0.21/32
  Known via "eigrp 3", distance 170, metric 2561280256, type external
  Redistributing via eigrp 3, rip
  Advertised by rip metric 1
  Last update from 99.2.21.2 on Tunnel12, 00:33:20 ago
  Routing Descriptor Blocks:
  * 99.2.21.2, from 99.2.21.2, 00:33:20 ago, via Tunnel12
      Route metric is 2561280256, traffic share count is 1
      Total delay is 50010 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

The route should have been known by RIP and NOT EIGRP and R1 is pointing to R2 as the next-hop.

NOTE: You may see R6 as the next-hop, it all depends which router (R2, R6) received the EIGRP advertisement regarding this prefix first.

Let's follow the path:

On R2:

```
R2#Show ip route 136.85.0.21

Routing entry for 136.85.0.21/32
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 2
  Redistributing via eigrp 3
  Advertised by eigrp 3 metric 1 1 1 1 1
  Last update from 99.2.23.3 on FastEthernet0/0, 01:31:13 ago
  Routing Descriptor Blocks:
  * 99.2.23.3, from 99.0.6.6, 01:31:13 ago, via FastEthernet0/0
      Route metric is 20, traffic share count is 1
```

Let's follow the path to R3

On R3:

```
R3#Show ip route 136.85.0.21

Routing entry for 136.85.0.21/32
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1
  Last update from 99.2.36.6 on FastEthernet0/1.36, 01:32:18 ago
  Routing Descriptor Blocks:
  * 99.2.36.6, from 99.0.6.6, 01:32:18 ago, via FastEthernet0/1.36
      Route metric is 20, traffic share count is 1
```

Let's follow the path to R6

On R6:

```
R6#Show ip route 136.85.0.21

Routing entry for 136.85.0.21/32
  Known via "eigrp 3", distance 170, metric 2561282816, type external
  Redistributing via ospf 1, eigrp 3
  Advertised by ospf 1 subnets
  Last update from 99.2.16.1 on FastEthernet0/0, 00:38:05 ago
  Routing Descriptor Blocks:
  * 99.2.16.1, from 99.2.16.1, 00:38:05 ago, via FastEthernet0/0
      Route metric is 2561282816, traffic share count is 1
      Total delay is 50110 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 2
```

It's pointing back to R1, and we know that R1 is pointing to R2. R1 should have known this prefix through RIP with R7 as the next-hop, and NOT through EIGRP with R2 as the next-hop.

RIP has a lower Administrative distance than EIGRP external routes, 120 versus 170. Do we have a routing loop? Let's Traceroute on R3 for this destination:

On R3:

```
R3#Traceroute 136.85.0.21
Type escape sequence to abort.

Tracing the route to 136.85.0.21
VRF info: (vrf in name/id, vrf out name/id)
  1 99.2.36.6 [MPLS: Label 26 Exp 0] 0 msec 0 msec 0 msec
  2 99.2.16.1 0 msec 4 msec 0 msec
  3 99.2.21.2 4 msec 0 msec 4 msec
  4 99.2.23.3 [MPLS: Label 27 Exp 0] 4 msec 0 msec 4 msec
  5 99.2.36.6 [MPLS: Label 26 Exp 0] 0 msec 4 msec 4 msec
  6 99.2.16.1 0 msec 4 msec 0 msec
  7 99.2.21.2 4 msec 4 msec 0 msec
  8 99.2.23.3 [MPLS: Label 27 Exp 0] 4 msec 4 msec 0 msec
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 9 99.2.36.6 [MPLS: Label 26 Exp 0] 4 msec 4 msec 0 msec
10 99.2.16.1 4 msec 4 msec 0 msec
11 99.2.21.2 4 msec 4 msec 4 msec
12 99.2.23.3 [MPLS: Label 27 Exp 0] 0 msec 4 msec 4 msec
13 99.2.36.6 [MPLS: Label 26 Exp 0] 0 msec 4 msec 4 msec
14 99.2.16.1 4 msec 0 msec 4 msec
15 99.2.21.2 4 msec 4 msec 0 msec
16 99.2.23.3 [MPLS: Label 27 Exp 0] 4 msec 4 msec 4 msec
17 99.2.36.6 [MPLS: Label 26 Exp 0] 0 msec 0 msec 4 msec
18 99.2.16.1 4 msec 4 msec 0 msec
19 99.2.21.2 0 msec 4 msec 4 msec
20 99.2.23.3 [MPLS: Label 27 Exp 0] 4 msec 4 msec 4 msec
21 99.2.36.6 [MPLS: Label 26 Exp 0] 0 msec 0 msec 4 msec
22 99.2.16.1 4 msec 4 msec 4 msec
23 99.2.21.2 0 msec 4 msec 4 msec
24 99.2.23.3 [MPLS: Label 27 Exp 0] 4 msec 4 msec 4 msec
25 99.2.36.6 [MPLS: Label 26 Exp 0] 4 msec 4 msec 4 msec
26 99.2.16.1 0 msec 0 msec 4 msec
27 99.2.21.2 4 msec 4 msec 4 msec
28 99.2.23.3 [MPLS: Label 27 Exp 0] 4 msec 4 msec 4 msec
29 99.2.36.6 [MPLS: Label 26 Exp 0] 4 msec 4 msec 4 msec
30 99.2.16.1 4 msec 4 msec 4 msec
```

Let's check RIP's configuration on R1

On R1:

```
R1#Show run | S router rip

router rip
 version 2
 redistribute eigrp 3 metric 1
 network 136.85.0.0
 distance 200
 no auto-summary
```

We can see the problem here, RIP was configured with an AD of 200; let's fix this problem by changing the AD back to 120:

On R1:

```
R1(config)#router rip
R1(config-router)#distance 120
```

To test and verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Show ip route 136.85.0.21

Routing entry for 136.85.0.21/32
  Known via "rip", distance 120, metric 1
  Redistributing via eigrp 3, rip

  Advertised by eigrp 3 metric 1 2222 1 1 1
  Last update from 136.85.21.21 on FastEthernet0/1.111, 00:00:24 ago
  Routing Descriptor Blocks:
  * 136.85.21.21, from 136.85.21.21, 00:00:24 ago, via
FastEthernet0/1.111
      Route metric is 1, traffic share count is 1
```

On R2:

```
R2#Ping 136.85.0.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.85.0.21, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/52 ms
```

Ticket 8

R5 can't ping 22.22.1.1 (R8's Lo100 interface).

Let's verify the ticket:

On R5:

```
R5#Ping 22.22.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 22.22.1.1, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

R5 receives ICMP unreachable messages. Let's enable debugging of ICMP and ping again to see which router is sending it the ICMP unreachable messages.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Debug ip icmp
ICMP packet debugging is on

R5#Ping 22.22.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 22.22.1.1, timeout is 2 seconds:

U.U.U
Success rate is 0 percent (0/5)

ICMP: dst (99.2.3.5) host unreachable rcv from 99.2.3.3
```

Let's check R5's routing table for this prefix:

On R5:

```
R5#Show ip route 22.22.1.1

Routing entry for 22.22.1.0/24
  Known via "bgp 45", distance 200, metric 0
  Tag 22, type internal
  Last update from 99.2.224.22 00:05:12 ago
  Routing Descriptor Blocks:
  * 99.2.224.22, from 99.0.4.4, 00:05:12 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 22
      MPLS label: none
```

We can see that R5 knows about 22.22.1.1 prefix via IBGP advertised by R4 with a next-hop IP address of 99.2.224.22. Let's check R3's routing table:

On R3:

```
R3#Sh ip route 22.22.1.1

% Network not in table
```

The reason the ping failed from R5 is because R3 does NOT have NLRI to 22.22.1.1 prefix.
R3 should have known this route through OSPF through R4. Let's see if R4 is redistributing the BGP routes into OSPF:

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show run | S router ospf

router ospf 1
 log-adjacency-changes
ipv6 router ospf 1
 log-adjacency-changes
```

R4 is NOT redistributing the BGP routes into OSPF. Let's configure redistribute BGP into OSPF:

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#redistribute bgp 45 subnets
```

To test and verify the configuration:

On R5:

```
R5#Ping 22.22.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 22.22.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 112/115/120 ms
```

Ticket 9

The following routes are sent from R8 to R4:

22.22.1.0 /24, 22.22.2.0 /24, and 22.22.3.0 /24

R4 was configured with the following policy for routes coming from R8:

- R5 should NOT advertise Network 22.22.1.0 /24 to R9.
- Filter Network 22.22.2.0 /24 and allow the rest.

However, R4 doesn't see Network 22.22.3.0 /24 coming from R8

To verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R4:

```
R4#Show ip bgp | B Network

Network          Next Hop            Metric LocPrf Weight Path
   0.0.0.0          0.0.0.0                              0 i
*> 22.22.1.0/24     99.2.224.22              0           0 22 i
```

Let's check and see how this BGP policy is implemented:

On R4:

```
R4#Show run | S router bgp

router bgp 45
 no synchronization
 bgp log-neighbor-changes
 neighbor 99.0.5.5 remote-as 45
 neighbor 99.0.5.5 update-source Loopback0
 neighbor 99.0.5.5 send-label
 neighbor 99.2.224.22 remote-as 22
 neighbor 99.2.224.22 default-originate
 neighbor 99.2.224.22 distribute-list 11 in
 neighbor 99.2.224.22 route-map fromR8 in
 no auto-summary
```

We can see that a distribute-list and a route-map is configured on R4. Let's see the distribute-list and route-map's configuration:

```
R4#Sh run | Inc access-list

access-list 11 deny   22.22.2.0
access-list 11 permit any

R4#Sh run | S route-map fromR8

 neighbor 99.2.224.22 route-map fromR8 in
route-map fromR8 permit 10
 match ip address prefix-list plR8
 set community no-export

R4#Sh run | Inc ip prefix-list

ip prefix-list plR8 seq 5 permit 22.22.1.0/24
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The route-map and the prefix-list are configured to advertise prefix 22.22.1.0 /24 to R5 with a community attribute of "no-export", this is done so R5 does NOT advertise this prefix to R9.

The route-map and the distribute-list are configured to comply with the policy, but they DO NOT work together.

In BGP the order of operation is important to remember:

Outbound policy evaluation order:

- Distribute-list
- Prefix-list
- Filter-list (Used for AS-Path filtering)
- Route-map

Inbound policy evaluation order:

- Route-map
- Filter-list (Used for AS-Path filtering)
- Prefix-list
- Distribute-list

One easy way to remember this is to use mnemonics:

- For Outbound : Don't Plan Fix Results
- For Inbound : reverse order

In this task, we can see that the distribute-list is configured properly and if this was the ONLY policy, it would have been implemented with no problems. But the route-map kind of overrides the distribute-list, because in the route-map we are telling the router if it matches the network identified in the prefix-list called "plR8", then, set the community attribute to "no-export", but since we did not configure a catch-all route-map statement, if the prefix does NOT match the one identified in the prefix-list it will be denied/dropped. Let's fix this problem:

On R4:

```
R4(config)#route-map fromR8 permit 90
```

To enforce the policy:

On R4

```
R4#Clear ip bgp *
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To test the first policy:

On R4

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 22.22.1.0/24     99.2.224.22              0             0 22 i
*> 22.22.3.0/24     99.2.224.22              0             0 22 i
```

To successfully implement the second policy, R4 MUST be configured to send-community to its IBGP neighbor 99.0.5.5 (R5), if R4 does NOT configure the following, it will drop the community and R5 will advertise the route to R9:

```
R4(config)#router bgp 45
R4(config-router)#Neighbor 99.0.5.5 send-community
```

To test and verify the second Policy:

On R5:

```
R5#Show ip bgp 22.22.1.0 255.255.255.0 | Inc Community

     Community: no-export
```



On R9:

```
BB3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          99.2.235.5               0             0 45 i
*> 22.22.3.0/24     99.2.235.5                             0 45 22 i
```

Ticket 10

R2 Lo0 is configured to join 239.2.2.2 Multicast group address. R6 can NOT reach this group address using pings sourced from it's Lo0 interface. You should fix this problem without using a PIM command.

Let's verify the problem:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Ping 239.2.2.2 source Lo0

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 99.0.6.6

Reply to request 0 from 99.2.26.2, 4 ms
```

This works, why would we have a trouble ticket? Let's check other aspects of the configuration:

On R6:

```
R6#Sh ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
99.0.6.6         Loopback0                v2/S   0      30     1      99.0.6.6
99.2.26.6        Tunnel26                 v2/S   1      30     1      0.0.0.0
```

This means that R6 sent two Echo requests, one sourcing from its Lo0 and another sourcing from its Tunnel26 interface; but it only received a single reply, let's see why?

Let's map the PIM topology:

On R6:

```
R6#Sh ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
99.0.6.6         Loopback0                v2/S   0      30     1      99.0.6.6
99.2.26.6        Tunnel26                 v2/S   1      30     1      0.0.0.0

R6#Show ip pim neighbor

PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
99.2.26.2         Tunnel26                 1d00h/00:01:19    v2    1 / S

R6#Show ip pim rp mapping

PIM Group-to-RP Mappings
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Group(s): 224.0.0.0/4, Static
    RP: 99.0.2.2 (?)
```

On R2:

```
R2#Show ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
99.0.2.2         Loopback0                v2/S   0      30     1      99.0.2.2
99.2.26.2        Tunnel26                 v2/S   1      30     1      0.0.0.0

R2#Show ip pim neighbor

PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
99.2.26.6         Tunnel26                 1d00h/00:01:36    v2    1 / S

R2#Show ip pim rp mapping

PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 99.0.2.2 (?)
```

We can see that PIM is configured properly on the appropriate interfaces and both routers have established a PIM neighbor adjacency and they both agree on the RP mapping, which is statically configured.

The next step in troubleshooting is to check for RPF failures:

On R2:

```
R2#Show ip rpf 99.0.6.6

failed, no route exists
```

The output of the above show command reveals that R2 has no RPF information for R6's Lo0. There can be two reasons:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

1. R2 has no route to R6's Lo0
2. R2 has a route to R6's Lo0, but it's pointing to a non PIM interface

Let's check and see if R2 has a route to R6's Lo0:

On R2:

```
R2#Show ip route 99.0.6.6

Routing entry for 99.0.6.6/32
  Known via "ospf 1", distance 110, metric 3, type intra area
  Redistributing via eigrp 3
  Advertised by eigrp 3 metric 1 1 1 1 1
  Last update from 99.2.23.3 on FastEthernet0/0, 1d00h ago
  Routing Descriptor Blocks:
  * 99.2.23.3, from 99.0.6.6, 1d00h ago, via FastEthernet0/0
      Route metric is 3, traffic share count is 1
```

We can see that R2 has a route to R6's Lo0, BUT it's pointing to a non PIM interface (F0/0).
To fix this problem, a static mroute can be configured:

On R2:

```
R2(config)#ip mroute 0.0.0.0 0.0.0.0 Tunnel 26
```

To verify the configuration:

On R2:

```
R2#Show ip rpf 99.0.6.6

RPF information for ? (99.0.6.6)
  RPF interface: Tunnel26
  RPF neighbor: ? (99.2.26.6)
  RPF route/mask: 0.0.0.0/0
  RPF type: static
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
```

Now that RPF checks are no longer failing, let's generate a ping to R6's Lo0

On R2:

```
R2#Ping 239.2.2.2 source Lo0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 99.0.2.2

Reply to request 0 from 99.0.2.2, 1 ms
```

Ticket 11

Lo26 interface of R2 and R6 are connected to each other via MPLS L3 VPN. R2's Lo26 is NOT reachable from VRF V26 on R6.

Let's verify the problem:

On R6:

```
R6#Ping vrf V26 26.26.0.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 26.26.0.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check R6's routing table for this prefix:

On R6

```
R6#Show ip route vrf V26 26.26.0.2

outing Table: V26
Routing entry for 26.26.0.2/32
  Known via "bgp 26", distance 200, metric 0, type internal
  Last update from 99.0.2.2 03:54:54 ago
  Routing Descriptor Blocks:
  * 99.0.2.2 (default), from 99.0.2.2, 03:54:54 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: 20
      MPLS Flags: MPLS Required
```

We can see R6 has a valid route, which is known via BGP with a next-hop IP address of R2's Lo0.

Let's check R2's routing table:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Show ip route vrf V26 | B Gate
Gateway of last resort is not set

     26.0.0.0/32 is subnetted, 2 subnets
C       26.26.0.2 is directly connected, Loopback26
B       26.26.0.6 [200/0] via 99.0.6.6, 1d02h
```

We can see that R2 has a route back to R6. Let's check and see if an LSP exists starting from Lo0 interface of R6 to Lo0 interface of R2:

On R6:

```
R6#Sh mpls forwarding-table 99.0.2.2

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
21         16         99.0.2.2/32      0             Fa0/1.36   99.2.36.3
```

R6 has an outgoing label for R2's Lo0. The next hop is R3. Let's go to R3:

On R3:

```
R3#Show mpls forwarding-table 99.0.2.2

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  99.0.2.2/32      92423         Fa0/0      99.2.23.2
```

The output of the above show command reveals that when the local router receives a labeled packet with the top label of 16, it will POP the top label and send it out of it's F0/0 interface toward R2. This shows that we have an LSP from R6 to R2; let's check and see if we have an LSP from R2 to R6:

On R2:

```
R2#Show mpls forwarding-table 99.0.6.6

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
21         20         99.0.6.6/32      0             Fa0/0      99.2.23.3
```

We can see that there is an outgoing label for R6's Lo0. Let's go to R3:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show mpls forwarding-table 99.0.6.6

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
20         No Label   99.0.6.6/32      99760         Fa0/1.36    99.2.36.6
```

Let's check and see the LDP advertisement from R6:

On R3:

```
R3#Show mpls ldp binding neighbor 99.0.6.6 99.0.6.6 32
R3#
```

The above show command proves that R6 never sent a label to the local LSR for this prefix. Let's check the routing table of R6:

On R6:

```
R6#Show ip route | Inc 99.0

99.0.0.0/8 is variably subnetted, 25 subnets, 2 masks
D        99.0.1.1/32 [90/156160] via 99.2.16.1, 03:49:12, FastEthernet0/0
O        99.0.2.2/32 [110/3] via 99.2.36.3, 06:52:21, FastEthernet0/1.36
O        99.0.3.3/32 [110/2] via 99.2.36.3, 06:52:21, FastEthernet0/1.36
O        99.0.4.4/32 [110/1002] via 99.2.36.3, 00:50:21, FastEthernet0/1.36
O        99.0.5.5/32 [110/1002] via 99.2.36.3, 01:21:28, FastEthernet0/1.36
C        99.0.6.0/24 is directly connected, Loopback0
L        99.0.6.6/32 is directly connected, Loopback0
O        99.0.11.11/32 [110/2] via 99.2.116.11, 06:52:16, FastEthernet0/1.116
O IA     99.0.12.12/32 [110/4] via 99.2.36.3, 02:30:33, FastEthernet0/1.36
```

Note R6 sent a label for 99.0.6.0 /24, let's verify:

On R3:

```
R3#Show mpls ldp binding neighbor 99.0.6.6 99.0.6.0 24

  tib entry: 99.0.6.0/24, rev 36
        remote binding: tsr: 99.0.6.6:0, tag: imp-null
```

BUT OSPF advertised a 99.0.6.6 /32. Which means that, LDP advertised a 99.0.6.0 /24, R3 receives a label advertised by R6 with NO matching route in it's routing table. The reason for this behavior is because OSPF advertises it's Loopback interfaces as /32. Let's fix this problem.

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6(config)#int Lo0
R6(config-if)#ip ospf network point-to-point
```

Let's check R3's LFIB:

On R3:

```
R3#Show mpls forwarding-table 99.0.6.6

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
28         Pop Label  99.0.6.0/24      135           Fa0/1.36    99.2.36.6
```

It looks much better, we actually have a label.

To test and verify the configuration:

On R6:

```
R6#Ping vrf V26 26.26.0.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 26.26.0.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Ticket 12

Erase the startup configuration of the Routers and the Switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets