azarpara.vahid@gmail.com

Advanced
CCIE Routing & Switching
v5.0

www.MicronicsTraining.com

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

VOL-I

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets
mcse2012.blogfa.com

azarpara.vahid@gmail.com

Table of content

Subject	Page
Topology	5
Switching	
Lab 1 Configuring Trunks	10
Lab 2 Configuring EtherChannels	48
Lab 3 DHCP Snooping	68
Lab 4 IP Source Gaurd	74
Lab 5 Dynamic ARP Inspection	86
Lab 6 MAC ACLs	95
Lab 7 VLAN Access-maps-I	99
Lab 8 VLAN Access-maps-II	103
Lab 9 Storm Control –I	111
Lab 10 Storm Control –II	116
Lab 11 Private VLANs	126
Lab 12 Basic STP configuration (802.1D)	153
Lab 13 Advanced STP configuration	169
Lab 14 Rapid Spanning-tree Protocol (802.1w)	192
Lab 15 Multiple Instance of STP (802.1s)	209
Lab 16 Spanning-tree Port Fast	220
Lab 17 Configuring Uplink Fast	231
Lab 18 Configuring BPDU Guard	245
Lab 19 Configuring BPDU Filter	252
Lab 20 Spanning-tree BackBone Fast	265
Lab 21 Spanning-tree Root Gaurd	273
Lab 22 Spanning-tree Loop Gaurd	281
DMVPN	
Lab 1 Configuring DMVPN Phase –I Using Static mappings	288
Lab 2 Configuring DMVPN Phase –I Using Dynamic mapping	298
Lab 3 Configuring DMVPN Phase –II Using Static mappings	306
Lab 4 Configuring DMVPN Phase –II Using Dynamic mapping	314
Lab 5 Running Routing Protocols on Phase #I	321
Lab 6 Running Routing Protocols on Phase #II and III	345
RIPv2	
Lab 1 Configuring RIPv2	371

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets
mcse2012.blogfa.com

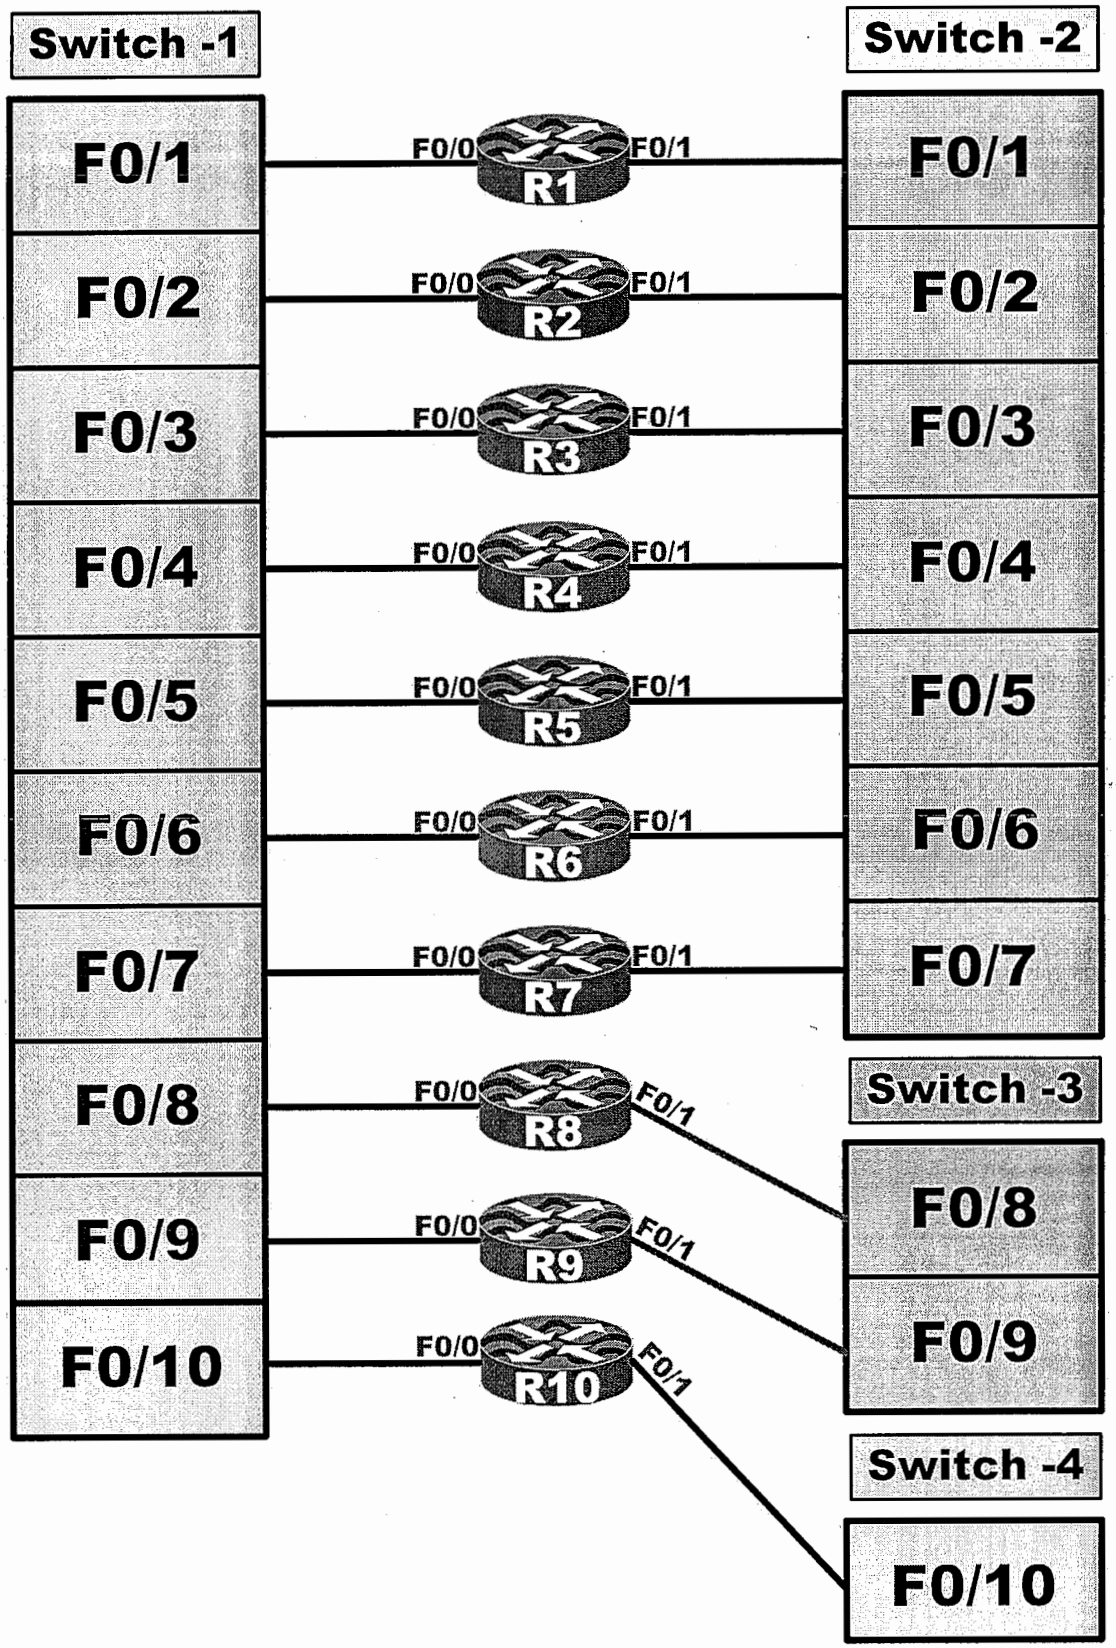
EIGRP	
Lab 1 EIGRP Named Configuration	434
Lab 2 EIGRP Stub	455
Lab 3 Advanced EIGRP Stub Configuration	467
Lab 4 EIGRP Filtering	476
Lab 5 EIGRP Mock Lab	483
Lab 6 Configuring EIGRP Authentication	557
OSPF	
Lab 1 Running OSPF on the interfaces	563
Lab 2 Establishing OSPF Adjacency	579
Lab 3 OSPF Broadcast Networks	604
Lab 4 OSPF Non-Broadcast Networks	618
Lab 5 OSPF Point-To-Point Networks	629
Lab 6 OSPF Point-To-Multipoint & Point-To-Multipoint Non-Broadcast Networks	633
Lab 7 OSPF Area Types	
Lab 8 OSPF Filtering - I	
Lab 9 OSPF Filtering - II	
Lab 10 OSPF Database Overload Protection	
Lab 11 OSPF Traffic Redirection	
Lab 12 Virtual-Links and GRE Tunnels	
Lab 13 OSPF Summarization	
Lab 14 OSPF Cost	
Lab 15 OSPF Authentication	
Lab 16 Default Route Injection	
Lab 17 Forward Address Suppression - I	
Lab 18 Forward Address Suppression – II	
Lab 19 Forward Address Suppression - IIII	
Redistribution	
Lab 1 Basics of Redistribution - I	641
Lab 2 Basics of Redistribution - II	662
Lab 3 Redistribution of RIPv2 & EIGRP	680
Lab 4 EIGR Redistribution of RIPv2 & OSPF	706
BGP	
Lab 1 Establishing Neighbor Adjacency	719

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 2 Route Reflectors	736
Lab 3 Conditional Advertisement and BGP Backdoors	754
Lab 4 Route Dampening	774
Lab 5 BGP Aggregation	783
Lab 6 The community Attribute	795
Lab 7 BGP Cost Community	812
Lab 8 BGP Load Balancing	821
Lab 9 Local Preference -I	838
Lab 10 Local Preference -II	848
Lab 11 The AS-Path Attribute	855
Lab 12 The Weight Attribute	863
Lab 13 Multi Exist Discriminator Attribute	872
Lab 14 BGP Filtering	889
Lab 15 Administrative Distance	933
Lab 16 Advanced BGP Configuration	941
Lab 17 BGP Confederation	953
Lab 18 BGP Hiding Local AS Number	965
Lab 19 BGP Allow-AS	973
Lab 20 BGP Automatic-Tag, and AS-Path tag & Table-map	978
Lab 21 4 Byte AS Numbers	985

© 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

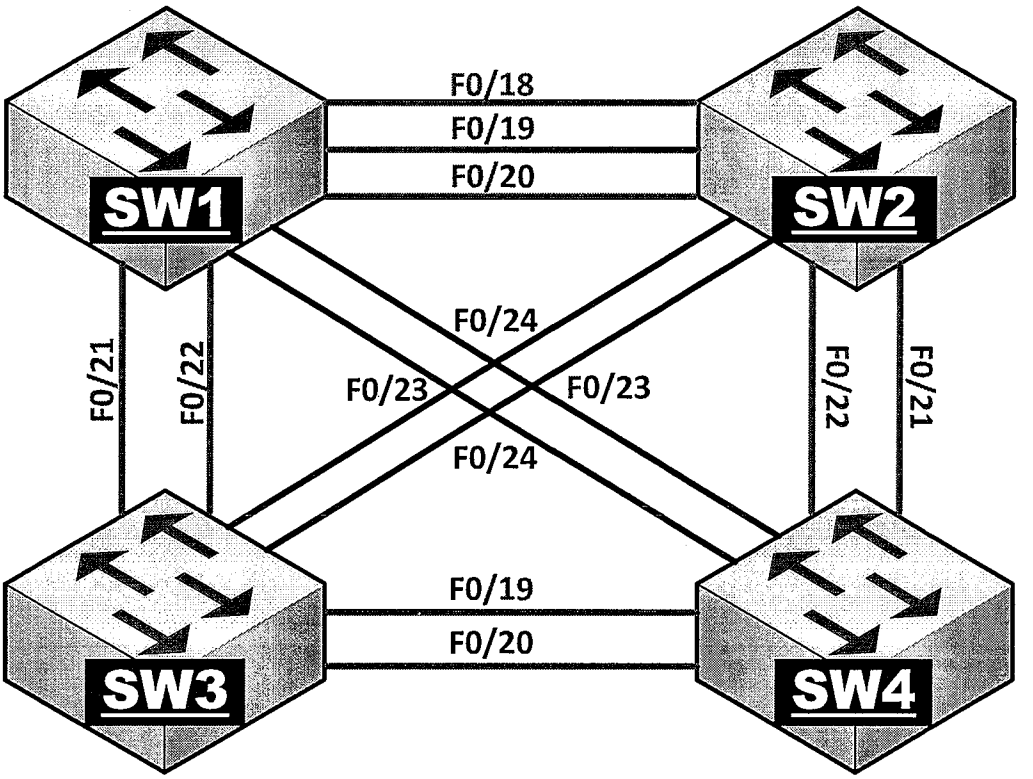
© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

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
CCIE Routing & Switching
v5.0

www.MicronicsTraining.com

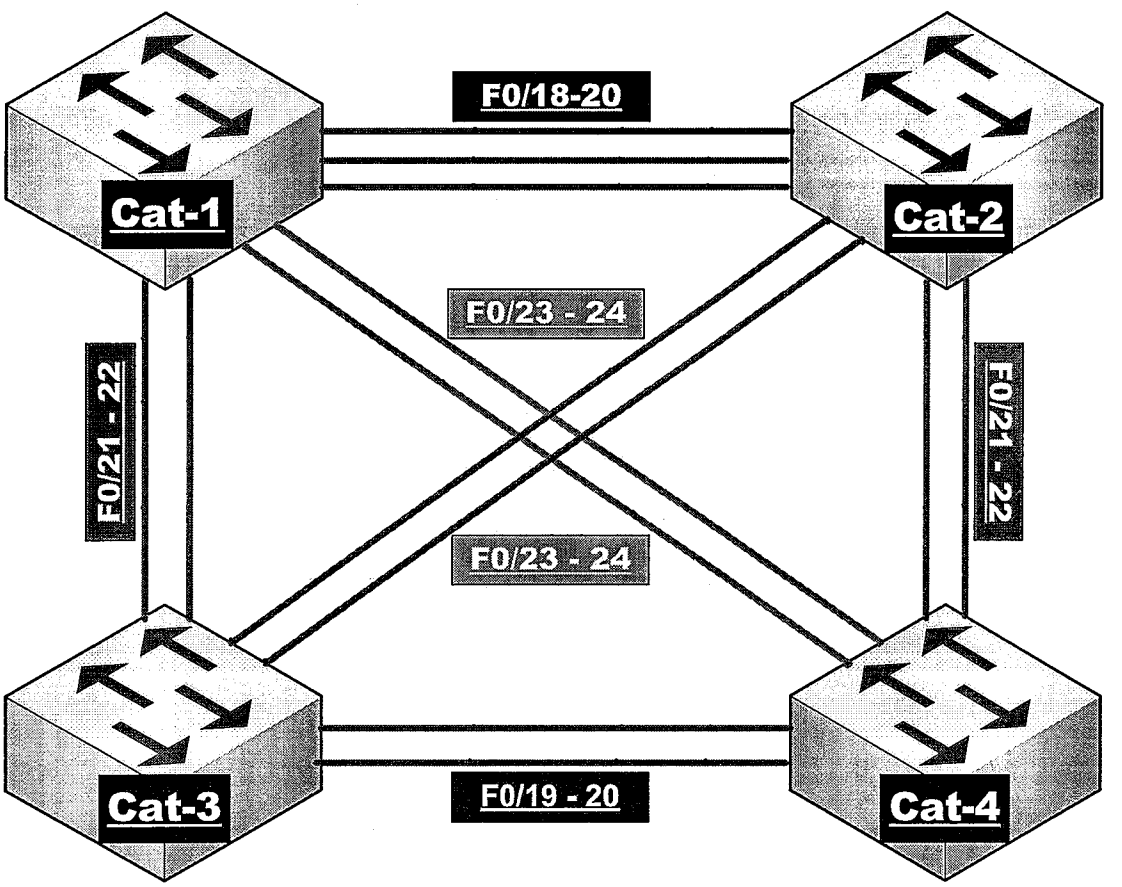
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

Switching

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Configuring Trunks



Task 1

Shutdown all ports on all four switches and configure the VTP domain name to be "TST".

On All Switches:

```
SWx(config)#int range f0/1 - 24
SWx(config-if-range)#shut

SWx(config)#vtp domain TST
```


© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

Task 2

Configure the following hostnames:

The first switch – Cat-1
The second switch – Cat-2
The third switch – Cat-3
The forth switch – Cat-4

On the first switch:

```
Switch(config)#hostname Cat-1
```

On the second switch:

```
Switch(config)#hostname Cat-2
```

On the third switch:

```
Switch(config)#hostname Cat-3
```

On the forth switch:

```
Switch(config)#hostname Cat-4
```

Task 3

Configure a Dot1q trunk between Cat-2 and Cat-4 using their F0/21 interfaces based on the following policy:

Cat-2 - F0/21 → this port should be configured into a permanent Trunking mode and it should negotiate to convert the neighboring interface into a trunk.
Cat-4 – F0/21 → this port should be configured to actively attempt to convert the link to a trunk. You should not configure "switchport trunk encapsulation dot1q" on this port.

On Cat-2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-2(config)#int f0/21
Cat-2(config-if)#switchport trunk encapsulation dot1q
Cat-2(config-if)#switch mode trunk
Cat-2(config-if)#No shutdown
```

On Cat-4:

```
Cat-4(config)#int f0/21
Cat-4(config-if)#switchport mode dynamic desirable
Cat-4(config-if)#No shut
```

To verify the configuration:

On Cat-2:

```
Cat-2#Sh int trunk | inc 802

Fa0/21      on              802.1q          trunking      1
```

On Cat-4:

```
Cat-4#Show int trunk | inc 802

Fa0/21      desirable    n-802.1q        trunking      1
```

Task 4

Configure a trunk between Cat-2 and Cat-4 using their F0/22 interfaces; you should use an industry standard protocol for the trunking encapsulation based on the following policy:

Cat-2 - F0/22 → this port should be configured into permanent Trunking mode and it should negotiate to convert the neighboring interface into a Trunk.

Cat-4 – F0/22 → this port should be configured to negotiate a trunk only if it receives negotiation packets from a neighboring port; this port should never start the negotiation process. You should not configure "switchport trunk encapsulation dot1q" on this port.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On Cat-2:

```
Cat-2(config)#int f0/22
Cat-2(config-if)#switchport trunk encap dot1q
Cat-2(config-if)#switchport mode trunk
Cat-2(config-if)#No shut
```

On Cat-4:

```
Cat-4(config)#int f0/22
Cat-4(config-if)#swi mode dynamic auto
Cat-4(config-if)#No shut
```

To verify the configuration:

On Cat-2:

```
Cat-2#Show int trunk | inc 802

Fa0/21      on          802.1q        trunking      1
Fa0/22      on          802.1q        trunking      1
```

On Cat-4:

In the output of the following show command you should note the difference between "dynamic desirable" and "Dynamic auto"

```
Cat-4#Sh int trunk | inc 802

Fa0/21      desirable   n-802.1q      trunking      1
Fa0/22      auto        n-802.1q      trunking      1
```

Task 5

Configure a Trunk link between Cat-3 and Cat-4 using their F0/19 interfaces. These ports should be configured to negotiate to convert the neighboring interface into a dot1q trunk, but they should not be in permanent trunking mode.

On Both Switches:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk encapsulation dot1q
Cat-x(config-if)#switchport mode dynamic desirable
Cat-x(config-if)#No shut
```

To verify the configuration:

On Cat-3:

```
Cat-3#Show int trunk | inc 802

Fa0/19      desirable    802.1q          trunking      1
```

On Cat-4:

```
Cat-4#Show int trunk | inc 802

Fa0/19      desirable    802.1q          trunking      1
Fa0/21      desirable    n-802.1q        trunking      1
Fa0/22      auto         n-802.1q        trunking      1
```

Task 6

Configure a Dot1q trunk between Cat-3 and Cat-4 using their F0/20 interfaces based on the following policy:

Cat-3 – F0/20 → this port should be configured to actively attempt to convert the link to a Trunk. This port should not be in permanent trunking mode.

Cat-4 – F0/20 → this port should be configured to negotiate a trunk only if it receives negotiation packets from a neighboring port; this port should never start the negotiation process or be configured with the "swi trunk enc dot" command.

On Cat-3:

```
Cat-3(config)#int f0/20
Cat-3(config-if)#switchport trunk encapsulation dot1q
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Cat-3(config-if)#swi mode dynamic desirable
Cat-3(config-if)#No shut
```

On Cat-4:

```
Cat-4(config)#int f0/20
Cat-4(config-if)#switchport mode dynamic auto
Cat-4(config-if)#No shut
```

To verify the configuration:

On Cat-3:

```
Cat-3#Sh int trunk | inc 802

Fa0/19      desirable    802.1q          trunking      1
Fa0/20      desirable    802.1q          trunking      1
```

On Cat-4:

```
Cat-4#Sh int trunk | inc 802

Fa0/19      desirable    802.1q          trunking      1
Fa0/20      auto         n-802.1q        trunking      1
Fa0/21      desirable    n-802.1q        trunking      1
Fa0/22      auto         n-802.1q        trunking      1
```

Task 7

Configure a Dot1q trunk between Cat-2 and Cat-3 using their F0/23 interfaces; these switches should be configured into permanent trunking mode and negotiate to convert the neighboring interface into a trunk.

On Both Switches:

```
Cat-x(config)#int f0/23
Cat-x(config-if)#switchport trunk encapsulation dot1q
Cat-x(config-if)#switchport mode trunk
Cat-x(config-if)#No shut
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-2:

```
Cat-2#Sh int trunk | inc 802

Fa0/21      on           802.1q         trunking      1
Fa0/22      on           802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1
```

On Cat-3:

```
Cat-3#Sh int trunk | inc 802

Fa0/19      desirable    802.1q         trunking      1
Fa0/20      desirable    802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1
```

Task 8

Configure a Dot1q trunk between Cat-2 and Cat-3 using interface F0/24; these ports should not use DTP to negotiate a Trunk.

Remember that the following commands are order specific, meaning that "swi mode trunk" cannot be entered before the "swi trunk enc dot" command, and/or the "swi nonegotiate" command cannot be entered before the "swi mode trunk" command.

On Both Switches:

```
Cat-x(config)#int f0/24
Cat-x(config-if)#switchport trunk encapsulation dot1q
Cat-x(config-if)#switchport mode trunk
Cat-x(config-if)#switchport nonegotiate
Cat-x(config-if)#No shut
```

To verify the configuration:

On Cat-2:

```
Cat-2#Sh int trunk | inc 802

Fa0/21      on           802.1q         trunking      1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/22      on            802.1q         trunking      1
Fa0/23      on            802.1q         trunking      1

Fa0/24      on            802.1q         trunking      1
```

On Cat-3:

```
Cat-3#Show int trunk | inc 802

Fa0/19      desirable     802.1q         trunking      1
Fa0/20      desirable     802.1q         trunking      1
Fa0/23      on            802.1q         trunking      1
Fa0/24      on            802.1q         trunking      1

Cat-3#Show int f0/24 swi | Inc Nego

Negotiation of Trunking: Off
```

Task 9

Configure trunking on F0/18 – F0/20 interfaces of Cat-1 and Cat-2, and F0/21-22 interfaces on Cat-1 and Cat-3 and the F0/23-24 interfaces of Cat-1 and Cat-4. These ports should be in permanent trunking mode.

On Cat-1:

```
Cat-1(config)#int range f0/18-24
Cat-1(config-if-range)#Swi trunk encap dot1q
Cat-1(config-if-range)#Swi mode trunk
Cat-1(config-if-range)#No shut
```

On Cat-2:

```
Cat-2(config)#int range f0/18-20
Cat-2(config-if-range)#Swi trunk encap dot
Cat-2(config-if-range)#Swi mode trunk
Cat-2(config-if-range)#No shut
```

On Cat-3:

```
Cat-3(config)#int range f0/21-22
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-3(config-if-range)#Swi trun encap dot1q
Cat-3(config-if-range)#Swi mode trunk

Cat-3(config-if-range)#No shut
```

On Cat-4:

```
Cat-4(config)#int range f0/23-24
Cat-4(config-if-range)#Swi trunk encap dot1q
Cat-4(config-if-range)#Swi mode trunk
Cat-4(config-if-range)#No shut
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show inter trunk | Inc 802.1q

Fa0/18      on              802.1q          trunking      1
Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
Fa0/21      on              802.1q          trunking      1
Fa0/22      on              802.1q          trunking      1
Fa0/23      on              802.1q          trunking      1
Fa0/24      on              802.1q          trunking      1
```

On Cat-2:

```
Cat-2#Show inter trunk | Inc 802.1q

Fa0/18      on              802.1q          trunking      1
Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
Fa0/21      on              802.1q          trunking      1
Fa0/22      on              802.1q          trunking      1
Fa0/23      on              802.1q          trunking      1
Fa0/24      on              802.1q          trunking      1
```

On Cat-3:

```
Cat-3#Show inter trunk | Inc 802.1q

Fa0/19      desirable       802.1q          trunking      1
Fa0/20      desirable       802.1q          trunking      1
Fa0/21      on              802.1q          trunking      1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/22      on          802.1q         trunking      1
Fa0/23      on          802.1q         trunking      1

Fa0/24      on          802.1q         trunking      1
```

Task 10

Configure the following VLANs on Cat-1 and ensure that they are propagated to the other switches.

VLANs 2 – 10, 100, 200, 300, 400, 120, 130, 140, 230, 240, and 340

On Cat-1

```
Cat-1(config)#vlan 2-10,100,200,300,400,120,130,140,230,240,340
```

You must "exit" for the VLANs to be propagated:

```
Cat-1(config-vlan)#exit
```

To verify the configuration:

On All Switches:

```
Cat-x#Sh vlan brie | inc VLAN

VLAN Name                                Status    Ports
2    VLAN0002                            active
3    VLAN0003                            active
4    VLAN0004                            active
5    VLAN0005                            active
6    VLAN0006                            active
7    VLAN0007                            active
8    VLAN0008                            active
9    VLAN0009                            active
10   VLAN0010                            active
100  VLAN0100                            active
120  VLAN0120                            active
130  VLAN0130                            active
140  VLAN0140                            active
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
200  VLAN0200                         active
230  VLAN0230                         active

240  VLAN0240                         active
300  VLAN0300                         active
340  VLAN0340                         active
400  VLAN0400                         active
```

Task 11

Configure the trunks based on the following policy:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/19	Cat-1 ←→ Cat-2	ONLY 120
2	F0/21	Cat-2 ←→ Cat-4	ONLY 240
3	F0/19	Cat-3 ←→ Cat-4	ONLY 340
4	F0/21	Cat-1 ←→ Cat-3	ONLY 130
5	F0/23	Cat-1 ←→ Cat-4	ONLY 140
6	F0/23	Cat-2 ←→ Cat-3	ONLY 230

Policy item 1: 

Let's verify the currect state of the trunk before the task is configured:

```
Cat-1#Show inter trunk | Inc Vlans|Fa0/19

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      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1-10,100,120,130,140,200,230,240,300,340,400
```

Note all the configured VLANs are allowed on the F0/19 trunk.

To configure the first item in the task:

On Both Switches:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Note: The following command only allows VLAN 120 on the trunk

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk allowed vlan 120
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show int trunk | Inc Vlans|Fa0/19

Fa0/19      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      120
Port        Vlans allowed and active in management domain
Fa0/19      120
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      120
```

Policy item 2: ◄

To configure the second item in the task:

On Cat-2 and Cat-4:

```
Cat-x(config)#int f0/21
Cat-x(config-if)#switchport trunk allowed vlan 240
```

To verify the configuration:

On Cat-2:

```
Cat-2#Show int trunk | Inc Vlan|Fa0/21

Fa0/21      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/21      240
Port        Vlans allowed and active in management domain
Fa0/21      240
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/21      240
```

On Cat-4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-4#Show int trunk | Inc Vlan|Fa0/21

Fa0/21      desirable    n-802.1q       trunking      1
Port        Vlans allowed on trunk
Fa0/21      240
Port        Vlans allowed and active in management domain
Fa0/21      240
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/21      240
```

Policy Item 3: 

To configure the third item in the task:

On Cat-3 and Cat-4:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk allowed vlan 340
```

To verify the configuration:

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/19

Fa0/19      desirable    802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      340
Port        Vlans allowed and active in management domain
Fa0/19      340
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      340
```

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/19

Fa0/19      desirable    802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      340
Port        Vlans allowed and active in management domain
Fa0/19      340
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      340
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Policy Item 4: 

To configure the fourth item in the task:

On Cat-1 and Cat-3:

```
Cat-x(config)#int f0/21
Cat-x(config-if)#switchport trunk allowed vlan 130
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/21

Fa0/21      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/21      130
Port        Vlans allowed and active in management domain
Fa0/21      130
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/21      none
```

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/21

Fa0/21      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/21      130
Port        Vlans allowed and active in management domain
Fa0/21      130
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/21      130
```

Policy Item 5: 

To configure the fifth item in the task:

On Cat-1 and Cat-4:

```
Cat-x(config)#int f0/23
Cat-x(config-if)#switchport trunk allowed vlan 140
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On Cat-1:

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      140
Port        Vlans allowed and active in management domain
Fa0/23      140
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      140
```

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      140
Port        Vlans allowed and active in management domain
Fa0/23      140
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      140
```

Policy Item 6: 

To configure the sixth item in the task:

On Cat-2 and Cat-3:

```
Cat-x(config)#int f0/23
Cat-x(config-if)#switchport trunk allowed vlan 230
```

To verify the configuration:

On Cat-2:

```
Cat-2#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/23      230
Port        Vlans allowed and active in management domain

Fa0/23      230
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      none
```

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      230
Port        Vlans allowed and active in management domain
Fa0/23      230
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      230
```

Task 11

Add VLANs to the allowed list of the trunks based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Add to the Allowed VLAN/s
1	F0/19	Cat-1 ←→ Cat-2	100
2	F0/21	Cat-2 ←→ Cat-4	200
3	F0/19	Cat-3 ←→ Cat-4	300
4	F0/23	Cat-1 ←→ Cat-4	400

Policy Item 1: 

To configure the first item in the task:

The following command adds VLAN 100 to the existing allowed VLAN list.

On Cat-1 and Cat-2:

```
Cat-x(config)#int f0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-x(config-if)#switchport trunk allowed vlan add 100
```

To verify the configuration:

As you can see VLAN 100 is added to the allowed VLAN list (120):

On Cat-1:

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/19

Fa0/19      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      100,120
Port        Vlans allowed and active in management domain
Fa0/19      100,120
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      100,120
```

On Cat-2:

```
Cat-2#Show inter trunk | Inc Vlan|Fa0/19

Fa0/19      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      100,120
Port        Vlans allowed and active in management domain
Fa0/19      100,120
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      none
```

Policy item 2: 

To configure the second item in the task:

On Cat-2 and Cat-4:

```
Cat-x(config)#int f0/21
Cat-x(config-if)#switchport trunk allowed vlan add 200
```

To verify the configuration:

On Cat-2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Cat-2#Show inter trunk | Inc Vlan|Fa0/21

Fa0/21      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/21      200,240
Port        Vlans allowed and active in management domain
Fa0/21      200,240
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/21      240
```

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/21

Fa0/21      desirable    n-802.1q       trunking      1
Port        Vlans allowed on trunk
Fa0/21      200,240
Port        Vlans allowed and active in management domain
Fa0/21      200,240
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/21      200,240
```

Policy item 3: 

To configure the third item in the task:

On Cat-3 and Cat-4:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk allowed vlan add 300
```

To verify the configuration:

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/19

Fa0/19      desirable    802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      300,340
Port        Vlans allowed and active in management domain
Fa0/19      300,340
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      300,340
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/19

Fa0/19      desirable    802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/19      300,340
Port        Vlans allowed and active in management domain
Fa0/19      300,340
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      300,340
```

Policy item 4: 

To configure the fourth item in the task:

On Cat-1 and Cat-4:

```
Cat-x(config)#int f0/23
Cat-x(config-if)#switchport trunk allowed vlan add 400
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      140,400
Port        Vlans allowed and active in management domain
Fa0/23      140,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      140
```

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      140,400
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Port        Vlans allowed and active in management domain

Fa0/23      140,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      140,400
```

Task 12

Remove VLANs from the allowed list of the trunks based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/22	Cat-1 ←→ Cat-3	Remove 1, 4 – 10 ONLY
2	F0/22	Cat-2 ←→ Cat-4	Remove 2, 4 – 10 ONLY

Policy item 1:

Let's see the existing VLANs that are allowed on the trunk:

```
Cat-1#Sh int trunk | Inc Vlans|Fa0/22
Fa0/22      on              802.1q          trunking      1
Port        Vlans allowed on trunk
Fa0/22      1-4094
Port        Vlans allowed and active in management domain
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      1-10,100,120,130,140,200,230,240,300,340,400
```

To configure the first item in the task:

The following command removes the specified VLANs from the trunk:

On Cat-1 and Cat-3:

```
Cat-x(config)#int f0/22
Cat-x(config-if)#switchport trunk allowed vlan remove 1,4-10
```

To verify the configuration:

On Cat-1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/22

Fa0/22      on          802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/22      2-3,11-4094
Port        Vlans allowed and active in management domain
Fa0/22      2-3,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      none
```

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/22

Fa0/22      on          802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/22      2-3,11-4094
Port        Vlans allowed and active in management domain
Fa0/22      2-3,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      2-3,100,120,130,140,200,230,240,300,340,400
```

Policy item 2: 

To configure the second item in the task:

On Cat-2 and Cat-4:

```
Cat-x(config)#int f0/22
Cat-x(config-if)#switchport trunk allowed vlan remove 2,4-10
```

To verify the configuration:

On Cat-2:

```
Cat-2#Show inter trunk | Inc Vlan|Fa0/22
Fa0/22      on          802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/22      1,3,11-4094
Port        Vlans allowed and active in management domain
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      1,3,100,120,130,140,230,300,340,400
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/22

Fa0/22      auto          n-802.1q       trunking      1
Port        Vlans allowed on trunk
Fa0/22      1,3,11-4094
Port        Vlans allowed and active in management domain
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      1,3,100,120,130,140,200,230,240,300,340,400
```

Task 13

Configure Cat-1, Cat-2 and Cat-4 based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/20	Cat-1 ←→ Cat-2	None
2	F0/24	Cat-1 ←→ Cat-4	None

Policy Item #1 ←

```
Cat-1#Sh int trunk | Inc Vlans|Fa0/20

Fa0/20      on              802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/20      1-4094
Port        Vlans allowed and active in management domain
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
```

To configure the first item in the task:

The following command removes ALL VLANs from the trunk

On Cat-1 and Cat-2:

```
Cat-x(config)#int f0/20
Cat-x(config-if)#swi trunk allow vlan none
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-1:

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/20

Fa0/20      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/20      none
Port        Vlans allowed and active in management domain
Fa0/20      none
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/20      none
```

On Cat-2:

```
Cat-2#Show inter trunk | Inc Vlan|Fa0/20

Fa0/20      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/20      none
Port        Vlans allowed and active in management domain
Fa0/20      none
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/20      none
```

Policy Item #2 

To configure the second item in the task:

On Cat-1 and Cat-4:

```
Cat-x(config)#int f0/24
Cat-x(config-if)#swi trunk allowed vlan none
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/24

Fa0/24      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/24      none
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Port        Vlans allowed and active in management domain

Fa0/24      none
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/24      none
```

On Cat-4:

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/24

Fa0/24      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/24      none
Port        Vlans allowed and active in management domain
Fa0/24      none
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/24      none
```

Task 14

Configure Cat-1, Cat-3 and Cat-4 based on the following chart:

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/20	Cat-3 ←→ Cat-4	All but 340
2	F0/22	Cat-1 ←→ Cat-3	All but 130

Policy Item #1 

```
Cat-3#Sh int trunk | Inc Vlans|Fa0/20
Fa0/20      desirable    802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/20      1-4094
Port        Vlans allowed and active in management domain
Fa0/20      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/20      1-10,100,120,130,140,200,230,240,400
```

To configure the first item in the task:

The following command allows all VLANs except 340

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-3 and Cat-4:

```
Cat-x(config)#int f0/20
Cat-x(config-if)#swi trunk allowed vlan except 340
```

To verify the configuration:

On Cat-3

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/20

Fa0/20      desirable    802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/20      1-339,341-4094
Port        Vlans allowed and active in management domain
Fa0/20      1-10,100,120,130,140,200,230,240,300,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/20      1-10,100,120,130,140,200,230,240,400
```

On Cat-4

```
Cat-4#Show inter trunk | Inc Vlan|Fa0/20

Fa0/20      auto         n-802.1q       trunking      1
Port        Vlans allowed on trunk
Fa0/20      1-339,341-4094
Port        Vlans allowed and active in management domain
Fa0/20      1-10,100,120,130,140,200,230,240,300,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/20      1-10,100,120,130,140,200,230,240,300,400
```

Policy Item #2 

To configure the second item in the task:

On Cat-1 and Cat-3:

```
Cat-x(config)#int f0/22
Cat-x(config-if)#swi trunk allowed vlan except 130
```

To verify the configuration:

On Cat-1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Cat-1#Show inter trunk | Inc Vlan|Fa0/22

Fa0/22      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/22      1-129,131-4094
Port        Vlans allowed and active in management domain
Fa0/22      1-10,100,120,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      1-10,100,120,200,230,240,300,340
```

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/22

Fa0/22      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/22      1-129,131-4094
Port        Vlans allowed and active in management domain
Fa0/22      1-10,100,120,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/22      1-10,100,120,140,200,230,240,300,340,400
```

Task 15

Configure Cat-2 and Cat-3 based on the following chart; you may override some of the previous task to accomplish this task.

Policy Item	Trunk Interface:	Between Switches	Allowed VLAN/s
1	F0/23	Cat-2 ←→ Cat-3	ALL
2	F0/24	Cat-2 ←→ Cat-3	ALL

Policy Item #1 & 2 ←

To configure the first item in the task:

The following command allows ALL VLANs on the trunk

On Cat-2 and Cat-3:

```
Cat-x(config)#int range f0/23-24
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-x(config-if)#swi trunk allow vlan all
```

To verify the configuration:

On Cat-3:

```
Cat-3#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      1-4094
Port        Vlans allowed and active in management domain
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      230

Cat-3#Show inter trunk | Inc Vlan|Fa0/24

Fa0/24      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/24      1-4094
Port        Vlans allowed and active in management domain
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
```

On Cat-2:

```
Cat-2#Show inter trunk | Inc Vlan|Fa0/23

Fa0/23      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/23      1-4094
Port        Vlans allowed and active in management domain
Fa0/23      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/23      2,4-10

Cat-2#Show inter trunk | Inc Vlan|Fa0/24

Fa0/24      on           802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/24      1-4094
Port        Vlans allowed and active in management domain
Fa0/24      1-10,100,120,130,140,200,230,240,300,340,400
Port        Vlans in spanning tree forwarding state and not pruned
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/24      none
```

Task 16

Erase the config.text and VLAN.dat on all four switches and reload them before proceeding to the next task.

On All Four Switches

```
Cat-x#delete vlan.dat
Cat-x#delete config.text

Cat-x#reload
```

Task 17

Configure Switch 1 and Switch 2 based on the following requirements:

- Configure a Dot1q trunk between Switch 1 and 2 using port F0/19 and shut down the remaining ports
- Configure the VTP domain name to be "TST"
- Configure the hostname of Switch 1 and Switch 2 to be Cat-1 and Cat-2 respectively

On the first switch:

```
Switch(config)#host Cat-1
```

On the second switch:

```
Switch(config)#host Cat-2
```

On both switches:

```
Cat-x(config)#int range f0/1-24
Cat-x(config-if-range)#shut
```

On Cat-1 and Cat-2:

```
Cat-x(config)#int f0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-x(config-if)#swi trunk encapsulation dot1q

Cat-x(config-if)#swi mode trunk
Cat-x(config-if)#No shut

Cat-x(config)#vtp domain TST
```

Task 18

Configure VLAN 100 on Cat-1 and assign its F0/1 interface to this VLAN.

On Cat-1:

```
Cat-1(config)#int f0/1
Cat-1(config-if)#swi mode acc

Cat-1(config-if)#swi acc v 100
Cat-1(config-if)#No shut
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show vlan brie | inc VLAN0100

100  VLAN0100                         active    Fa0/1
```

Task 19

Configure the switches such that they restrict flooded traffic to those trunk links that the traffic must use to access the appropriate network device(s).

This is an example of a vague question. This task is asking for VTP pruning to be enabled. To understand VTP pruning, it is helpful to know the VTP message types.

There are four types of VTP advertisements that are exchanged between the switches, and they are:

1. **Summary advertisements:** **This is an update sent by VTP servers or a client every 300**

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

seconds or when a VLAN database change occurs. This update includes the VTP version, domain name, configuration revision number, time stamp, and number of subset advertisements. If the advertisement results from a VLAN database change, one or more subset advertisements will follow.

2. **Subset advertisements:** An update that follows a summary advertisement resulting from a change in the VLAN database. A subset advertisement includes the specific changes that were made to a given VLAN(s).
3. **Advertisement requests from clients:** These are updates sent

   by a switch requesting more information so it can update its database. When a switch receives a VTP summary advertisement with a configuration revision number higher than its own, the local switch will send an advertisement request, requesting information about changes so it can update its VLAN database. A switch operating in VTP server mode then responds with one or more subset advertisements.
4. **VLAN membership announcement:** These messages are generated by the switches when VTP pruning is enabled and a port is associated to a given VLAN; these messages tell the neighboring switch that the local switch is interested in receiving traffic for that given VLAN. If the local switch does not send this message for a given VLAN, the neighboring switch will not send the traffic for that VLAN, and therefore the traffic for that VLAN will be pruned.

On Cat-1:

```
Cat-1#Show interface pruning

Pruning not currently enabled in this device's VTP administrative domain.
```

Note the above message states that the pruning feature is not enabled. The output of the following message reveals the same fact:

```
Cat-1#Show vtp status | inc Prun

VTP Pruning Mode                : Disabled
```

To enable VTP pruning:

```
Cat-1#vtp pruning

Pruning switched on
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show vtp status | inc Prun
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VTP Pruning Mode                    : Enabled
```

Note: This configuration will be propagated to all switches that have a trunk established with the local switch that are in the same VTP domain:

On Cat-2:

```
Cat-2#Show vtp status | inc Prun

VTP Pruning Mode                    : Enabled

Cat-2#Sh interface f0/19 pruning
```

Note: The following output has two sections. The first section lists VLANs that are pruned, because the local switch has not received a VLAN Membership Announcement message (VMA) from the neighboring switch. Therefore, the local switch is not pruning any VLANs.

The second section of the output identifies for what VLANs the local switch has sent VMAs, and therefore, not pruned, from the local switch's perspective:

On Cat-1:

```
Cat-1#Show interface f0/19 pruning
```

The local switch will NOT send traffic for VLAN out of this trunk interface (F0/19), because the local switch has NOT received VMAs for this VLAN. Meaning that Cat-1 has propagated the VLANs, in this case VLAN 100, but has not received VMAs from the neighboring switch through the F0/19 interface, for VLAN 100, therefore, Cat-1 will prune VLAN 100:

```
Port        Vlans pruned for lack of request by neighbor
Fa0/19      100
```

Note the local switch has sent VMAs for these two VLANs:

```
Port        Vlan traffic requested of neighbor
Fa0/19      1,100
```

Task 20

Configure VLANs 200, 300, 400, 500 and 600 on Cat-1 and ensure that these VLANs are propagated to Cat-2.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On Cat-1:

```
Cat-1(config)#vlan 200,300,400,500,600
Cat-1(config-vlan)#exit
```

On Cat-2:

```
Cat-2#Sh vlan br | Inc VLAN

VLAN Name                                Status    Ports
100  VLAN0100                            active
200  VLAN0200                            active
300  VLAN0300                            active
400  VLAN0400                            active
500  VLAN0500                            active
600  VLAN0600                            active
```

To verify the configuration:

On Cat-1:

Note the output of the following show command displays that VLANs 100, 200, 300, 400, 500 and 600 are pruned:

```
Cat-1#Show interface pruning
```

Note: The neighboring switch that is connected to the local switch's F0/19 interface has not sent VMAs for VLANs 100,200,300,400,500 or 600. Therefore, they are pruned. VMAs are messages that a switch uses to request traffic for a given VLAN(s):

```
Port        Vlans pruned for lack of request by neighbor
Fa0/19      100,200,300,400,500,600
```

Since the local switch has port membership in VLANs 1 and 100, it has sent VMAs for these VLANs. Therefore, the neighboring switch will not prune traffic for VLANs 1 or 100.

```
Port        Vlan traffic requested of neighbor
Fa0/19      1,100
```

On Cat-2:

```
Cat-2#Show interface pruning
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port          Vlans pruned for lack of request by neighbor
Fa0/19        200,300,400,500,600

Port          Vlan traffic requested of neighbor
Fa0/19        1
```

Task 21

Configure F0/2 interface of Cat-2 in VLAN 100.

On Cat-2:

```
Cat-2(config)#int f0/2
Cat-2(config-if)#swi mode acc
Cat-2(config-if)#swi acc v 100
Cat-2(config-if)#No shut
```

Note: You may have to wait for 30 seconds for STP convergence:

```
Cat-2#Show interface pruning

Port          Vlans pruned for lack of request by neighbor
Fa0/19        200,300,400,500,600

Port          Vlan traffic requested of neighbor
Fa0/19        1,100
```

Note: The output of the above show command reveals that the local switch has sent VMA messages for VLANs 1 and 100.

On Cat-1:

```
Cat-1#Sh interface pruning

Port          Vlans pruned for lack of request by neighbor
Fa0/19        200,300,400,500,600

Port          Vlan traffic requested of neighbor
Fa0/19        1,100
```


© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

Note: The first section of the above output reveals that VLAN 100 is no longer pruned.

Task 22

Configure the switches such that only VLAN 300 is pruned regardless of port membership.

On Cat-1:

```
Cat-1#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      200,300,400,500,600

Port        Vlan traffic requested of neighbor
Fa0/19      1,100
```

Note: The above output reveals that VLANs 200, 300, 400, 500 and 600 are pruned.

To configure the switches such that only VLAN 300 is pruned:

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk pruning vlan 300
```

Note: The above command instructs the trunk to prune VLAN 300 only, therefore, the rest of the VLANs in the VLAN Database will not be pruned.

On Cat-1:

```
Cat-1#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      300

Port        Vlan traffic requested of neighbor
Fa0/19      1,100,200,400,500,600
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-2:

```
Cat-2#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      300

Port        Vlan traffic requested of neighbor
Fa0/19      1,100,200,400,500,600
```

Task 23

Configure the switches such that VLAN 200 is also pruned. You should not use the command from the previous task to accomplish this task.

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk pruning vlan add 200
```

To verify the configuration:

On Cat-1:

Note: VLAN 200 is added to the list of pruned VLANs

```
Cat-1#Sh inter pruning

Port        Vlan pruned for lack of request by neighbor
Fa0/19      200,300

Port        Vlan traffic requested of neighbor
Fa0/19      1,100,400,500,600
```

On Cat-2:

```
Cat-2#Show interface pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      200,300
```

Note: VLAN 200 is added to the list of pruned VLANs

```
Port        Vlan traffic requested of neighbor
Fa0/19      1,100,400,500,600
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 24

Configure the switches such that none of the VLANs are pruned.

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk pruning vlan NONE
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      none
                          Note: None of the VLANs are pruned
Port        Vlan traffic requested of neighbor
Fa0/19      1,100,200,300,400,500,600
```

On Cat-2:

```
Cat-2#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      none
                          Note: None of the VLANs are pruned
Port        Vlan traffic requested of neighbor
Fa0/19      1,100,200,300,400,500,600
```

Task 25

Configure the switches such that all configured VLANs in the VLAN DB are pruned.

On Both Switches:

```
Cat-x(config)#int f0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-x(config-if)#switch trunk pruning vlan 1,100,200,300,400,500,600
```

You should receive the following errors:

```
Command rejected Bad VLAN pruning list.
```

The reason the error message was generated was that VLAN 1 cannot be pruned.

Let's re-enter the command without specifying VLAN 1:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switch trunk pruning vlan 100,200,300,400,500,600
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      200,300,400,500,600

Port        Vlan traffic requested of neighbor
Fa0/19      1,100
```

Note: VLAN 100 cannot be pruned because the local switch has port membership in that VLAN.

On Cat-2:

```
Cat-2#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      200,300,400,500,600

Port        Vlan traffic requested of neighbor
Fa0/19      1,100
```

To summarize the points in this task:

- **VLAN 1 cannot be pruned**
- **If the local switch has port membership in a given VLAN, that VLAN cannot be pruned**

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 26

Configure the switches such that VLAN 200 is no longer pruned; do not use a command that was used before to accomplish this task.

On Both Switches:

```
Cat-x(config)#int f0/19
Cat-x(config-if)#switchport trunk pruning vlan remove 200
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      300,400,500,600

Port        Vlan traffic requested of neighbor
Fa0/19      1,100,200
```

Note: VLAN 200 was removed from the list of VLANs being pruned.

On Cat-2:

```
Cat-2#Show interface f0/19 pruning

Port        Vlans pruned for lack of request by neighbor
Fa0/19      300,400,500,600

Port        Vlan traffic requested of neighbor
Fa0/19      1,100,200
```

With the "switchport trunk pruning vlan remove" command, you are asking the switch to remove the specified VLAN from the list of pruned VLANs, meaning "do not prune this VLAN".

Task 27

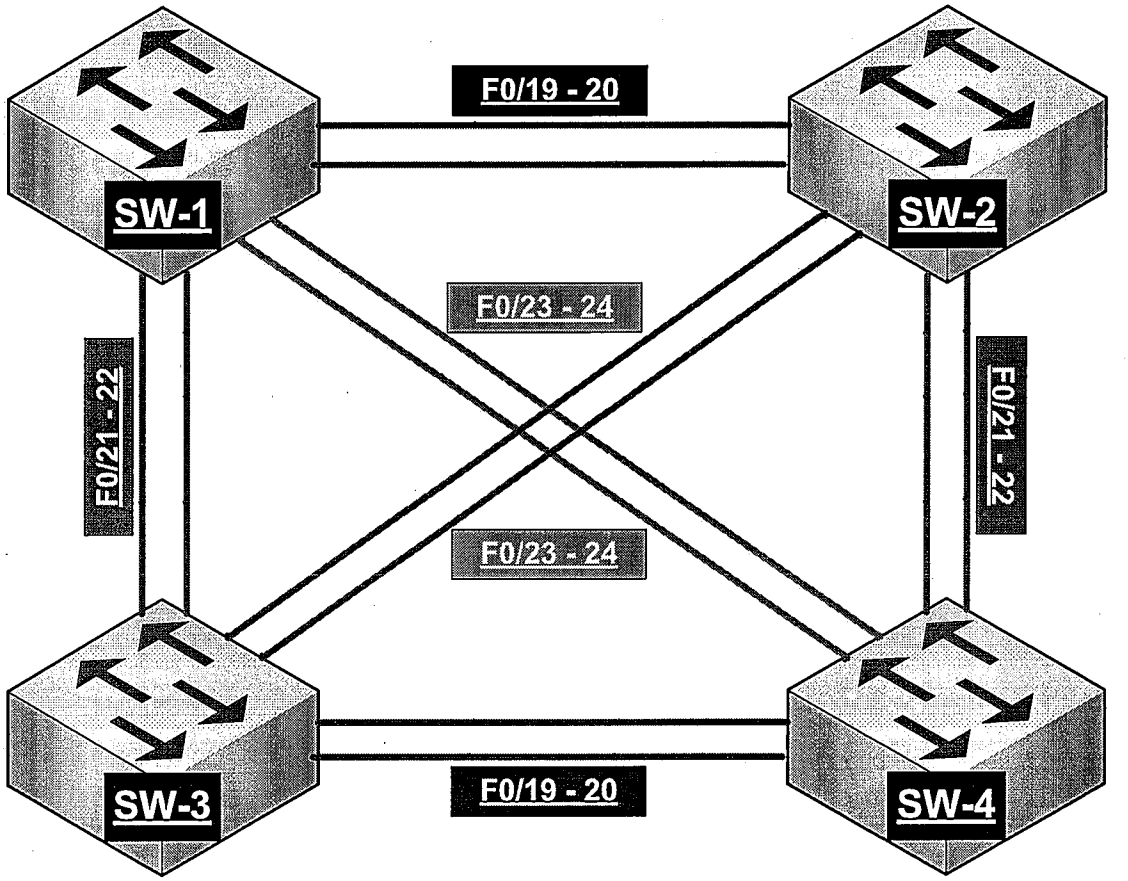
Erase the vlan.dat and config.text and reload the switches before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 2
Configuring EtherChannels



Task 1

Configure the hostname of the switches based on the above diagram. Ensure that the ports of these four switches are in shutdown mode. Configure these four switches in a VTP domain called "TST'.

On SW-1:

```
Switch(config)#hostname sw-1

SW-1(config)#int range f0/1-24
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW-1(config-if-range)#shut

SW-1(config)#vtp domain TST
```

On SW-2:

```
Switch(config)#hostname sw-2

SW-2(config)#int range f0/1-24
SW-2(config-if-range)#shut

SW-2(config)#vtp domain TST
```

On SW-3:

```
Switch(config)#hostname sw-3

SW-3(config)#int range f0/1-24
SW-3(config-if-range)#shut

SW-3(config)#vtp domain TST
```

On SW-4:

```
Switch(config)#hostname sw-4

SW-4(config)#int range f0/1-24
SW-4(config-if-range)#shut

SW-4(config)#vtp domain TST
```

Task 2

Configure ports F0/19 and F0/20 on SW-1 and SW-2 as trunk links using an industry standard protocol. These links should appear to spanning-tree protocol as a single link. If one of the links fails, the traffic should use the other link without any interruption. The ports on SW-1 should be configured such that they only respond to PAgP packets and never start the negotiation process.

EtherChannels provide the follows:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

- Fault-tolerant, high speed links between switches and routers
- Automatic recovery for the loss of a link by redistributing the traffic across the remaining links
- STP will not block one of the links in the bundle, because to STP the bundle looks like a single link
- Up to 8 links can be combined to provide more bandwidth
- The links within the bundle must have the same characteristics such as duplex, speed, etc
- EtherChannel can be configured as Layer 2 or Layer 3
- With Layer 3, a logical interface (Port-Channel) is statically configured and all Layer 3 configurations are performed under that interface
- With Layer 2, the logical interface is created automatically
- With both Layer 2 and Layer 3, physical interfaces must be manually assigned to the logical interface using the "channel-group" command
- EtherChannels can be configured automatically using Port Aggregation Protocol (PAgP) or Link Aggregation Control Protocol (LACP)
- PAgP is a Cisco proprietary protocol, whereas LACP is an industry standard IEEE 802.3ad protocol
- Switches can be configured to use PAgP by configuring them in AUTO or DESIRABLE mode
- Switches can be configured to use LACP by configuring them in ACTIVE or PASSIVE mode
- If the switches are configured in ON mode, they will not exchange LACP or PAgP packets

There are 5 modes:

- **ON** – Forces the interface into an EtherChannel without PAgP or LACP packets, both switches must be configured in ON mode for the EtherChannel to be established.

- **ACTIVE** – Used in LACP - the switches will actively negotiate an EtherChannel link.

- **PASSIVE** – Used in LACP - places the interface in a passive negotiation mode where it only responds to LACP packets that it receives. In this mode, the switch will not start the negotiation process; this setting minimizes the transmission of LACP packets.

- **AUTO** – Used in PAgP - places the interface in a passive negotiation mode; it only responds to PAgP packets that it receives. In this mode, the switch will not start the negotiation process; this setting minimizes the transmission of PAgP packets.

- **DESIRABLE** – Used in PAgP - the switches will actively negotiate an EtherChannel link.

The following table is very important to understand when configuring EtherChannels:

If SW-1 is configured in	If SW-2 is configured in	Will an EtherChannel be established?	The protocol used:
Desirable	Desirable	YES	PAgP

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Desirable	Auto	YES	PAgP
Auto	Auto	NO	--
Active	Active	YES	LACP
Active	Passive	YES	LACP
Passive	Passive	NO	--
ON	ON	YES	NONE
ON	Auto	NO	--
ON	Desirable	NO	--
ON	Passive	NO	--
ON	Active	NO	--

The configuration of EtherChannels should be done in a certain order, the following is my recommendation for creating EtherChannels:

1. Configure "default interface" for the interfaces involved.
2. Assign a channel-group and channel-group number to the physical interfaces, this step will create a port-channel interface automatically.
3. Configure the trunking encapsulation directly in port-channel interface configuration mode.
4. Reset the ports in the group by entering "shut" and then, "no shut".

Step One:

On SW-1:

```
SW-1(config)#default interface range f0/19-20
```

You should see the following console message:

```
%LINK-3-UPDOWN: Interface FastEthernet0/19, changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/20, changed state to up

SW-1(config)#interface range f0/19-20
SW-1(config-if-range)#No shut
```

Step Two:

```
SW-1(config)#int range f0/19-20
SW-1(config-if-range)#channel-group 12 mode Auto
```

You should see the following messages:

Creating a port-channel interface Port-channel 12

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Note: The interface Port-channel 12 is created automatically:

```
SW-1#Show run interface port-channel 12 |b inter

interface Port-channel12
end
```

Step Three:

```
SW-1(config)#int port-channel 12
SW-1(config-if)#switchport trunk encapsulation dot1q
SW-1(config-if)#switchport mode trunk
```

On SW-2:

```
SW-2(config)#default interface range f0/19-20
```

You should see the following console message:

```
%LINK-3-UPDOWN: Interface FastEthernet0/19, changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/20, changed state to up

SW-2(config)#int range f0/19-20
SW-2(config-if-range)#channel-group 21 mode Desirable

SW-2(config)#int port-channel 21
SW-2(config-if)#switchport trunk encapsulation dot1q
SW-2(config-if)#switchport mode trunk
```

Step Four:

On SW-1 and SW-2:

```
SW-x(config-if)#int range f0/19-20
SW-x(config-if-range)#shut
SW-x(config-if-range)#No shut
```

To verify the configuration:

On SW-1:

```
SW-1#Sh interface trunk
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Mode         Encapsulation  Status        Native vlan
Po12        on           802.1q         trunking      1

Port        Vlans allowed on trunk
Po12        1-4094

Port        Vlans allowed and active in management domain
Po12        1

Port        Vlans in spanning tree forwarding state and not pruned
Po12        1
```

On SW-2:

```
SW-2#Sh interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Po12        on           802.1q         trunking      1

Port        Vlans allowed on trunk
Po12        1-4094

Port        Vlans allowed and active in management domain
Po12        1

Port        Vlans in spanning tree forwarding state and not pruned
Po12        1
```

On SW-1:

```
SW-1#Show interface f0/19 switchport | inc Operational Mode

Operational Mode: trunk (member of bundle Po12)
```

On SW-2:

```
SW-2#Sh int f0/19 swi | inc Operational Mode

Operational Mode: trunk (member of bundle Po21)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 3

Configure ports F0/21 and F0/22 on SW-3 and SW-1 as trunk links using an industry standard protocol. These links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should **not** negotiate by exchanging **LACP** or **PAgP** protocols to accomplish this task.

On SW-1:

```
SW-1(config)#default interface range f0/21-22

SW-1(config)#int range f0/21 - 22
SW-1(config-if-range)#channel-group 13 mode on
SW-1(config-if-range)#NO shut

SW-1(config-if-range)#int port-channel 13
SW-1(config-if)#switchport trunk encapsulation dot1q
SW-1(config-if)#swi mode trunk

SW-1(config)#int range f0/21-22
SW-1(config-if-range)#shut
SW-1(config-if-range)#No shut
```

On SW-3:

```
SW-3(config)#default int range f0/21-22

SW-3(config)#int range f0/21 - 22
SW-3(config-if-range)#channel-group 31 mode on
SW-3(config-if-range)#NO shut

SW-3(config-if-range)#int port-channel 31
SW-3(config-if)#switchport trunk encapsulation dot1q
SW-3(config-if)#swi mode trunk

SW-3(config)#int range f0/21-22
SW-3(config-if-range)#shut
SW-3(config-if-range)#No shut
```

To verify the configuration:

On SW-1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

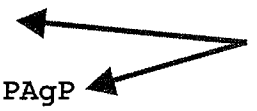
```
SW-1#Show interface trunk | inc 802

Po12        on           802.1q         trunking      1
Po13        on           802.1q         trunking      1

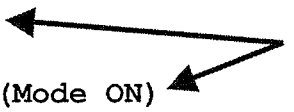
SW-1#Show etherchannel protocol

       Channel-group listing:
                -----------------------
Group: 12
----------
Protocol:  PAgP

Group: 13
----------
Protocol:   -   (Mode ON)
```



Note: PAgP is used for Etherchannel negotiation



Note: PAgP or LACP is not in use

On SW-3:

```
SW-3#Show interface trunk | inc 802

Po31        on           802.1q         trunking      1

SW-3#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
------+-------------+-----------+-----------------------
31     Po31(SU)         -        Fa0/21(P)   Fa0/22(P)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

Ensure that all the EtherChannels created on SW-1 are load-balanced based on destination MAC address.

Etherchannel load balancing can be done on 3550 or 3560 switches; however, 3560 switches are more flexible and provide more options. The following explains the load-balancing options available on each of these switch models:

On 3550 Switches:

Source MAC Address – Packets forwarded to an EtherChannel are distributed across the ports in the channel based on the **source MAC address** of the incoming packets. Therefore, different devices with different source MAC addresses use different interfaces in the bundle. When source MAC address load balancing is enabled, the load distribution based on the **source, destination IP address pair is also enabled,** and this is for routed IP traffic.

Destination MAC Address – If the EtherChannel is between a router and a switch and since the router has

a single MAC address, destination-based load balancing is the best way. In this load balancing method, packets forwarded to an EtherChannel are distributed across the ports in the channel based on the **destination MAC address** of the incoming packets.

Note there are only two choices on 3550 switches:

On SW3 (A 3550 switch):

```
SW-3(config)#port-channel load-balance ?
  dst-mac  Dst MAC Addr
  src-mac  Src MAC Addr
```

To verify the default settings on 3550 switches:

On SW-3:

```
SW-3#Show etherchannel load-balance

EtherChannel Load-Balancing Operational State (src-mac):
Non-IP: Source MAC address
    IPv4: Source MAC address
    IPv6: Source IP address
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

SW-1 is a 3560:

The following are the options available on 3560 switches:

Source MAC Address – When packets are forwarded to an Etherchannel, they are distributed across the ports in the channel based on the **source MAC address** of the incoming frame.

Destination MAC Address - When packets are forwarded to an Etherchannel, they are distributed across the ports in the channel based on the **destination MAC address** of the incoming frame.

Source and Destination MAC Address - When packets are forwarded to an Etherchannel, they are distributed across the ports in the channel based on the **source & destination MAC address pair** of the incoming frame.

Source IP Address - When packets are forwarded to an Etherchannel, they are distributed across the ports in the channel based on the **source IP address** of the incoming frame.

Destination IP Address - When packets are forwarded to an Etherchannel, they are distributed across the ports in the channel based on the **destination IP address** of the incoming frame.

Source & Destination IP Address - When packets are forwarded to an Etherchannel, they are distributed across the ports in the channel based on the **source & destination IP address pair** of the incoming frame.

To see the above options on 3560 switches:

On SW1:

```
SW-1(config)#port-channel load-balance ?
  dst-ip        Dst IP Addr
  dst-mac       Dst MAC Addr
  src-dst-ip    Src XOR Dst IP Addr
  src-dst-mac   Src XOR Dst MAC Addr
  src-ip        Src IP Addr
  src-mac       Src MAC Addr
```

To verify the default setting:

On SW1:

```
SW-1#Show etherchannel load-balance

EtherChannel Load-Balancing Operational State (src-mac):
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Non-IP: Source MAC address
    IPv4: Source MAC address
    IPv6: Source IP address
```

To configure the load balancing based on the destination MAC addresses:

On SW-1:

```
SW-1(config)#port-channel load-balance dst-mac
```

To verify the configuration:

On SW-1:

```
SW-1#Show etherchannel load

EtherChannel Load-Balancing Operational State (dst-mac):
Non-IP: Destination MAC address
     IPv4: Destination MAC address
     IPv6: Destination IP address
```

Note: Since the command is entered in the global configuration mode, it effects all EtherChannel ports created on the local switch.

Task 5

Ensure that all the EtherChannels created on SW-2 are load-balanced based on the following policy:

- For non-IP, source XOR and destination MAC address
- For IPv4, source XOR and destination IP address pair
- For IPv6, source XOR and destination IP address pair

On SW-2:

```
SW-2(config)#port-channel load-balance src-dst-ip
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW-2:

```
SW-2#Show etherchannel load-balance

EtherChannel Load-Balancing Operational State (src-dst-ip):
Non-IP: Source XOR Destination MAC address
  IPv4: Source XOR Destination IP address
  IPv6: Source XOR Destination IP address
```

The following reveals the behavior of a 3560 switch when the load balancing is changed:

If the load-balancing is changed to "src-mac":
Non-IP: Source MAC address
IPv4: Source MAC address
IPv6: Source IP address

If the load-balancing is changed to "dst-mac":
Non-IP: Destination MAC address
IPv4: Destination MAC address
IPv6: Destination IP address

If the load-balancing is changed to "src-ip":
Non-IP: Source MAC address
IPv4: Source IP address
IPv6: Source IP address

If the load-balancing is changed to "dst-ip":
Non-IP: Destination MAC address
IPv4: Destination IP address
IPv6: Destination IP address

If the load-balancing is changed to "src-dst-mac":
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination MAC address
IPv6: Source XOR Destination IP address

If the load-balancing is changed to "src-dst-ip":
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination IP address
IPv6: Source XOR Destination IP address

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 6

Configure ports F0/21 and F0/22 on SW-2 and SW-4 as trunk links using a Cisco proprietary trunking encapsulation. These links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should actively negotiate an EtherChannel using PAgP.

On SW-2:

```
SW-2(config)#default interface range f0/21-22

SW-2(config)#int range f0/21-22
SW-2(config-if-range)#channel-group 24 mode desirable
SW-2(config-if-range)#NO shut

SW-2(config)#int port-channel 24
SW-2(config-if)#switchport trunk encapsulation isl
SW-2(config-if)#switchport mode trunk

SW-2(config-if-range)#int range f0/21-22
SW-2(config-if-range)#shut
SW-2(config-if-range)#NO shut
```

On SW-4:

```
SW-4(config)#default interface range f0/21-22

SW-4(config)#int range f0/21-22
SW-4(config-if-range)#channel-group 42 mode desirable
SW-4(config-if-range)#No shut

SW-4(config)#int port-channel 42
SW-4(config-if)#switchport trunk encapsulation isl
SW-4(config-if)#switchport mode trunk

SW-4(config-if-range)#int range f0/21-22
SW-4(config-if-range)#shut
SW-4(config-if-range)#No shut
```

To verify the configuration:

On SW-4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW-4#Show interface trunk | inc isl

Po42        on          isl          trunking        1

SW-4#Show etherchannel protocol

                Channel-group listing:
                ----------------------

Group: 42
----------
Protocol:  PAgP
```



Note: PAgP is used for Etherchannel negotiation

On SW-2:

```
SW-2#Show interface trunk | Inc isl

Po24        on          isl          trunking        1

SW-2#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: 2
Number of aggregators            2

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------
21     Po21(SU)          PAgP         Fa0/19(P)   Fa0/20(P)
24     Po24(SU)          PAgP         Fa0/21(P)   Fa0/22(P)
```

Task 7

Configure ports F0/19 and F0/20 on SW-3 and SW-4 as trunk links using a Cisco proprietary trunking encapsulation. These links should appear to STP as a single link. If

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

one of the links fails, the traffic should use the other link without any interruption. These ports on SW-3 should be configured such that they only respond to LACP packets that are received from the appropriate ports on SW-4.

On SW-3:

```
SW-3(config)#default inter range f0/19-20

SW-3(config)#int range f0/19-20
SW-3(config-if-range)#channel-group 34 mode passive

SW-3(config-if-range)#No shut

SW-3(config)#int port-channel 34
SW-3(config-if)#switchport trunk encapsulation isl
SW-3(config-if)#switchport mode trunk

SW-3(config)#int range f0/19-20
SW-3(config-if-range)#shut
SW-3(config-if-range)#No shut
```

On SW-4:

```
SW-4(config)#default interface range f0/19-20

SW-4(config)#int range f0/19-20
SW-4(config-if-range)#channel-group 43 mode active
SW-4(config-if-range)#No shut

SW-4(config)#int port-channel 43
SW-4(config-if)#switchport trunk encapsulation isl
SW-4(config-if)#switchport mode trunk

SW-4(config)#int range f0/19-20
SW-4(config-if-range)#shut
SW-4(config-if-range)#No shut
```

To verify the configuration:

On SW-3:

```
SW-3#Show etherchannel protocol

                Channel-group listing:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                    --------------------------
Group: 31
--------------------------------------------------------------------------------

Protocol:    -   (Mode ON)

Group: 34
----------
Protocol:  LACP

SW-3#Show interface trunk | inc isl

Po34        on            isl             trunking        1

On SW-4:

SW-4#Show interface trunk | inc isl

Po42        on            isl             trunking        1
Po43        on            isl             trunking        1

SW-4#Sh ether summ | b Number

Number of channel-groups in use: 2
Number of aggregators:           2

Group   Port-channel  Protocol    Ports
--------+-------------+-----------+-----------------------
42        Po42(SU)        PAgP        Fa0/21(P)   Fa0/22(P)
43        Po43(SU)        LACP        Fa0/19(P)   Fa0/20(P)
```

Task 8

Configure ports F0/23 and F0/24 on SW-1 and SW-4 as trunk links using a Cisco proprietary trunking encapsulation.These links should appear to STP as a single link. If one of the links fails, the traffic should use the other link without any interruption. These ports should be configured such that they actively negotiate a LACP Etherchannel.

On SW-1:

```
SW-1(config)#default interface range f0/23-24
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW-1(config)#int range f0/23-24
SW-1(config-if-range)#channel-group 14 mode active
SW-1(config-if-range)#No shut

SW-1(config)#int port-channel 14
SW-1(config-if)#switchport trunk encapsulation isl
SW-1(config-if)#switchport mode trunk

SW-1(config-if)#int range f0/23-24
SW-1(config-if-range)#shut
SW-1(config-if-range)#No shut
```

On SW-4:

```
SW-4(config)#default interface range f0/23-24

SW-4(config)#int range f0/23-24
SW-4(config-if-range)#channel-group 41 mode active
SW-4(config-if-range)#No shut

SW-4(config)#int port-channel 41
SW-4(config-if)#switchport trunk encapsulation isl
SW-4(config-if)#switchport mode trunk

SW-4(config-if)#int range f0/23-24
SW-4(config-if-range)#shut
SW-4(config-if-range)#No shut
```

To verify the configuration:

On SW-1:

```
SW-1#Show inter trunk | inc isl

Po14        on          isl           trunking      1
```

On SW-4:

```
SW-4#Show inter trunk | inc isl

Po41        on          isl           trunking      1
Po42        on          isl           trunking      1
Po43        on          isl           trunking      1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW-4#Show etherchannel pro

                Channel-group listing:
                 ----------------------

Group: 41
----------
Protocol:  LACP

Group: 42
----------
Protocol:  PAgP

Group: 43
----------
Protocol:  LACP
```

Task 9

Configure ports F0/23 and F0/24 on SW-2 and SW-3 as a single layer three link; SW-2 should be configured with an IP address of 10.1.23.2 /24 and SW-3 should be configured with an IP address of 10.1.23.3 /24. These ports should not negotiate using LACP or PAgP.

Note: When configuring Layer 3 EtherChannels, I recommend the order of operation to be as follows:

1. **Default interface the physical interfaces**
2. **Configure the interface port-channel**
3. **Configure the port-channel interface with "no swi" and then configure the IP address**
4. **Configure the physical interfaces with "no swi"**
5. **Assign the port-channel ID to the interfaces using the channel-group interface configuration command**
6. **Reset the physical interfaces by using "shut" and "no shut"**

On SW-2:

```
SW-2(config)#default interface range f0/23-24

SW-2(config)#int port-channel 23
SW-2(config-if)#NO switchport
SW-2(config-if)#ip addr 10.1.23.2 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW-2(config)#int range f0/23-24
SW-2(config-if-range)#NO switchport
SW-2(config-if-range)#channel-group 23 mode on
SW-2(config-if)#NO shut

SW-2(config)#int range f0/23-24
SW-2(config-if-range)#shut
SW-2(config-if-range)#NO shut
```

On SW-3:

```
SW-3(config)#default interface range f0/23-24

SW-3(config)#int port-channel 32
SW-3(config-if)#No switchport
SW-3(config-if)#ip addr 10.1.23.3 255.255.255.0

SW-3(config)#int range f0/23-24
SW-3(config-if-range)#channel-group 32 mode on
```

Note: If the "no switchport" interface command is not configured, you should see the following error:

```
Command rejected (Port-channel32, Fa0/23): Either port is L2 and port-
channel is L3, or vice-versa % Range command terminated because it
failed on FastEthernet0/23

SW-3(config-if-range)#No swi
SW-3(config-if-range)#channel-group 32 mode on
SW-3(config-if-range)#No shut

SW-3(config)#int range f0/23-24
SW-3(config-if-range)#shut
SW-3(config-if-range)#No shut
```

To verify and test the configuration:

On SW-2:

```
SW-2#Show etherchannel summary | b Number

Number of channel-groups in use: 3
Number of aggregators:           3

Group  Port-channel  Protocol    Ports
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
-------+---------------+----------+-----------------------------
21         Po21(SU)         PAgP       Fa0/19(P)     Fa0/20(P)
23         Po23(RU)          -         Fa0/23(P)     Fa0/24(P)
24         Po24(SU)         PAgP       Fa0/21(P)     Fa0/22(P)
```

On SW-3:

```
SW-3#Ping 10.1.23.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.23.2, timeout is 2 seconds:

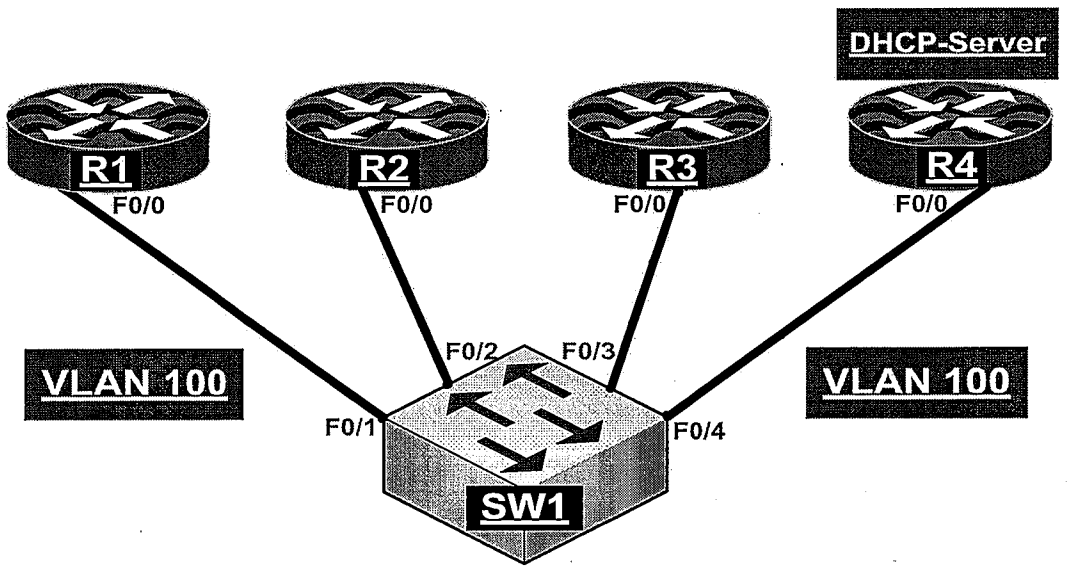
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Task 10

Erase the startup configuration and vlan.dat before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 3
DHCP Snooping



Task 1

Configure the F0/0 interface of R1 – 4 and SW1 based on the following policy:

- R4 should be configured as the DHCP server, offering IP addresses; the range should be from 10.1.1.0/24 network.
- R4 should be configured with an IP address of 10.1.1.254 /24.
- R1 – R3 should be configured to acquire an IP address from a DHCP server.
- R1 – R4 should be configured with 0000.1111.1111, 0000.2222.2222, 0000.3333.3333, and 0000.4444.4444 Mac-addresses respectively.
- Configure the F0/0 interfaces of R1 – R4 in VLAN 100.
- The F0/0 interfaces of R1 – R4 should be configured to bypass listeneing and learning state.

On SW1:

The following configuration places the F0/0 interfaces of R1 – R4 in VLAN 100, and enables portfast on

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

these interfaces:

```
SW1(config)#int range f0/1-4
SW1(config-if-range)#swi mode acc
SW1(config-if-range)#swi acc v 100
SW1(config-if-range)#Spanning-tree PortFast
```

On R4:

The following configuration, configures R4 as a DHCP server; the DHCP service is enabled, the DHCP pool called "TST" is created, the range of IP addresses in the scope are defined and the IP address of R4 is excluded:

```
R4(config)#service dhcp

R4(config)#ip dhcp pool TST
R4(dhcp-config)#Network 10.1.1.0 255.255.255.0

R4(config)#ip dhcp excluded-address 10.1.1.254

R4(config)#int f0/0
R4(config-if)#ip addr 10.1.1.254 255.255.255.0
R4(config-if)#No shut
```

The routers are configured based on the task's requirement:

On R1:

```
R1(config)#int f0/0
R1(config-if)#mac-address 0000.1111.1111
R1(config-if)#ip address dhcp
```

On R2:

```
R2(config)#int f0/0
R2(config-if)#mac-address 0000.2222.2222
R2(config-if)#IP address dhcp
```

On R3

```
R3(config)#int f0/0
R3(config-if)#mac-address 0000.3333.3333
R3(config-if)#ip address dhcp
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip int brie F0/0

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.1        YES DHCP   up                    up
```

On R2:

```
R1#Show ip int brie F0/0

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.2        YES DHCP   up                    up
```

On R3:

```
R1#Show ip int brie F0/0

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.3        YES DHCP   up                    up
```

Task 2

SW1 should be configured to provide security against rogue DHCP servers. You should disable option 82. R4 should be the ONLY DHCP server in this VLAN.

As you can see that any router can be configured as a DHCP server; even if there is a DHCP server, another router can be configured as the second DHCP server, but what if the second DHCP server is a rogue DHCP server? How do we protect ourselves against a rogue DHCP server?

DHCP Snooping is the answer, if the DHCP Snooping is configured on the switch, every port on that switch will transition into "Untrusted" state. The ports that are in "Untrusted" state will discard incoming DHCP Offers, ACKs, or NACKs which means that the existing DHCP server will NOT work either, unless the port that the real DHCP server is connected to is trusted.

Let's configure DHCP Snooping:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

DHCP Snooping must be enabled globally and then applied to a given Vlan or a range of Vlans; in this case VLAN 100:

On SW1:

```
SW1(config)#ip dhcp snooping
SW1(config)#ip dhcp snooping vlan 100
```

The following command disables Option 82:

```
SW1(config)#No ip dhcp snooping infor option
```

The following command ensures that R4 is the ONLY DHCP server. Since DHCP snooping is enabled, if this port is NOT trusted, DHCPOFFER, DHCPACK, DHCPNACK, DHCPLEASEQUERY packets will be dropped, and as a result, R1 – R3 won't be able to acquire an IP address from R4.

```
SW1(config)#int f0/4
SW1(config-if)#ip dhcp snoo trust
```

To verify the configuration:

On SW1:

```
SW1#Show ip dhcp snooping

Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
100
Insertion of option 82 is disabled
   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)
------------------------   -------     ----------------
FastEthernet0/4            yes         unlimited
```

NOTE: The output of the above command reveals that ONLY F0/4 interface of SW1 is in trusted state, therefore, the host connected to this port is the ONLY host that can act as a DHCP server.

Let's save the configuration on all routers and SW1, and reload them before verifying and testing this option.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On All Routers and SW1:

```
#wr
Building configuration...
[OK]

#Reload

Proceed with reload? [confirm]→ Press Enter to reload
```

To verify the configuration after the routers are up:

The routers may acquire a different IP address.

On R1:

```
R1#Show ip int brie F0/0

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.1        YES DHCP   up                    up
```

On R2:

```
R1#Show ip int brie F0/0

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.2        YES DHCP   up                    up
```

On R3:

```
R1#Show ip int brie F0/0

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.3        YES DHCP   up                    up
```

On SW1:

```
SW1#Show ip dhcp snooping binding

MacAddress          IpAddress       Lease(sec)  Type           VLAN  Interface
------------------  ----------      ----------  -------------  ----  ----------
00:00:22:22:22:22   10.1.1.2        86259       dhcp-snooping  100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1        86254       dhcp-snooping  100   FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
00:00:33:33:33:33   10.1.1.3    86277       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 3
```

Note the DHCP snooping feature eavesdrops on the conversation between the Untrusted ports and the DHCP Server (Which is connected to a trusted port), and it builds a Snooping database.

NOTE: ONLY the IP and MAC address of the devices that are connected to untructed ports are tracked.

Even though the database is auto-created, it is NOT used if the DHCP Snooping is the only option enabled. The database is used ONLY if IP source Guard or DAI is configured.

DO NOT RELOAD THE ROUTERS OR THE SWITCH

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 4
IP Source Guard

Task 1

Configure SW1 such that if the IP address of R1 is changed, the traffic from this router is dropped.

IP Source Guard is a security feature that can perform source IP address or source IP and Mac-address filtering on a non-routed layer 2 interface/s. IP Source Guard can be used to prevent IP address spoofing.

There are two prerequisites for using this feature:

1. IP Source Guard depends on DHCP Snooping for its operation, meaning that the DHCP Snooping MUST be enabled in order for us to use IP Source Guard.
2. The switch MUST run the EMI image.

IP Source Guard can be configured to do one of the followings:

Source IP Address Filtering:
With this feature, IP traffic is filtered based on the source IP address of ingress traffic. The traffic is allowed when the source IP address matches an entry in the DHCP snooping binding database, or a static IP source binding is configured.

Source IP and MAC Address Filtering:
With this feature, IP traffic is filtered based on the source IP and MAC address of ingress traffic. The traffic is allowed when the source IP and MAC address matches an entry in the DHCP snooping binding database, or a static IP source binding is confifgured.

In this task, Source IP Address Filtering is performed.

On SW1:

```
SW1(config)#Int F0/1
SW1(config-if)#ip verify source
```

To test and verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

To test this configuration, the IP Address of R1 is changed:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#IP addr 10.1.1.5 255.255.255.0

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 0 percent (0/5)
```

NOTE: The IP address is no longer in the DHCP Snooping database:

On SW1:

```
SW1#Sh ip dhcp snooping binding

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  ---------------
00:00:22:22:22:22   10.1.1.2         85473       dhcp-snooping  100   FastEthernet0/2
00:00:33:33:33:33   10.1.1.3         85799       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 2
```

To test the configuration further:

On R1:

```
R1(config)#Int f0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0

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:
.....
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

NOTE: Even though R1's IP address is changed back to what it was, it still can NOT communicate with the hosts in the same VLAN. Because R1 is no longer in the DHCP Snooping Database.

On SW1:

```
SW1#Sh ip dhcp bind

IP address       Client-ID/                Lease expiration        Type
                 Hardware address

SW1#Sh ip dhcp snoop bind

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
00:00:22:22:22:22   10.1.1.2         83820       dhcp-snooping  100   FastEthernet0/2
00:00:33:33:33:33   10.1.1.3         83868       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 2
```

Note: R1 is NO longer in the Snooping binding table. R1 will gain its connectivity ONLY if its added to the table; to test this:

On SW1:

```
SW1#ip dhcp snooping bind 0000.1111.1111 vlan 100 10.1.1.1 inter F0/1 exp 10000
```

To verify the configuration:

On SW1:

```
SW1#Sh ip dhcp snoop bind

MacAddress          IpAddress  Lease(sec)  Type           VLAN  Interface
------------------  ---------  ----------  -------------  ----  ----------
00:00:22:22:22:22   10.1.1.2   83516       dhcp-snooping   100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1   99736       dhcp-snooping   100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3   83564       dhcp-snooping   100   FastEthernet0/3
Total number of bindings: 3
```

To test the configuration:

On R1:

```
R1#Ping 10.1.1.2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

Note: The security is provided at layer 3, therefore, if the MAC address of R1 is changed, it should NOT effect the communication at all. To test and verify this fact:

On R1:

```
R1(config)#int f0/0
R1(config-if)#No mac-address

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

Let's change R1's Mac-address back to 0000.1111.1111:

On R1:

```
R1(config)#int f0/0
R1(config-if)#mac-address 0000.1111.1111
```

Since R1 is configured with a static IP address, you must exclude this address from the DHCP server:

On R4:

```
R4(config)#ip dhcp excluded-address 10.1.1.1
```

Task 2

Configure SW1 such that if the IP address or the MAC address of R2 is changed, the traffic from this router is dropped.

On SW1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1(config)#Int f0/2
SW1(config-if)#ip verify source port-security
SW1(config-if)#Switchport port-security
```

To test and verify the configuration:

On R2:

```
R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#No Mac-address

R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's change the Mac-address of R2's F0/0 back to 0000.2222.2222:

```
R2(config)#Int f0/0
R2(config-if)#mac-address 0000.2222.2222

R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

NOTE: When the mac address was changed (Removed), the entry was NOT removed from the DHCP Snooping DB, to see this:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show ip dhcp snooping binding

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------
00:00:22:22:22:22   10.1.1.2         84868       dhcp-snooping  100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1         99621       dhcp-snooping  100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3         84870       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 3
```

Let's test the IP address; when the static IP address is entered, you will see a slight hesitation, this is because the entry has to be removed from the Snooping DB.

```
R2(config)#Int f0/0
R2(config-if)#ip addr 10.1.1.5 255.255.255.0

R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note R2 is no longer in the DHCP snooping Database and once again it must be added statically:

On SW1:

```
SW1#Sh ip dhcp snoop bind

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
00:00:11:11:11:11   10.1.1.1         98585       dhcp-snooping   100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3         82413       dhcp-snooping   100   FastEthernet0/3
Total number of bindings: 2

SW1#ip dhcp snooping bind 0000.2222.2222 Vlan 100 10.1.1.2 inter F0/2 exp 10000

SW1#sh ip dhcp snoop bind

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------------
00:00:22:22:22:22   10.1.1.2         994         dhcp-snooping   100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1         98510       dhcp-snooping   100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3         82339       dhcp-snooping   100   FastEthernet0/3
Total number of bindings: 3
```

Let's change the IP address of R2 back to 10.1.1.2/24:

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config)#int F0/0
R2(config-if)#ip addr 10.1.1.2 255.255.255.0

R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Since R2 is configured with a static IP address, you must exclude this address from the DHCP server:

On R4:

```
R4(config)#ip dhcp excluded-address 10.1.1.2
```

Task 3

Assign 10.1.1.3 /24 to R3's F0/0 interface and add this address to the DHCP Snooping Database, and verify reachability to the other routers in the same VLAN.

On SW1:

```
SW1#Debug ip dhcp snooping event
DHCP Snooping Event debugging is on
```

On R3:

```
R3(config)#int f0/0
R3(config-if)#ip addr 10.1.1.3 255.255.255.0
```

Once the above IP address is statically configured you should see the following console message on SW1:

```
DHCP_SNOOPING: delete binding from port FastEthernet0/3.

DHCP_SNOOPING: dump binding entry: Mac=00:00:33:33:33:33 Ip=10.1.1.3
Lease=84495      ld Type=dhcp-snooping Vlan=100 If=FastEthernet0/3
```

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#IP dhcp snoop bind 0000.3333.3333 vlan 100 10.1.1.3 int f0/3 exp 1000
```

Once the static entry is configured you should get the following console message on SW1:

```
DHCP_SNOOPING: dump binding entry: Mac=00:00:33:33:33:33 Ip=10.1.1.3
Lease=100000    ld Type=dhcp-snooping Vlan=100 If=FastEthernet0/3
```

To test and verify reachability:

On SW1:

```
SW1#Sh ip dhcp snooping bin

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  ----------------
00:00:22:22:22:22   10.1.1.2         99325       dhcp-snooping  100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1         98674       dhcp-snooping  100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3         99668       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 3
```

On R3:

```
R3#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R3#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
```

Task 4

Configure SW1 such that if the IP address or the MAC address of R3 is changed, the traffic from this router is dropped. Configure a static binding to bind IP address 10.1.1.3 /24 to a MAC-Address of 0000.0123.0123 for Vlan 100 through F0/3 interface. This configuration should be part of the running config.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW1:

```
SW1(config)#int f0/3
SW1(config-if)#ip verify source port-security
SW1(config-if)#swi port-security
```

Sicne the task stated that this configuration should be part of the startup configuration, the snooping DB can not be manipulated, and a static IP Source Guard should be configured. The static IP source Guard will be part of the running configuration if it's saved.

```
SW1(config)#ip source binding 0000.1234.1234 vlan 100 10.1.1.3 inter f0/3
```

On R3:

```
R3#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Note even though a static entry is configured for the F0/3 interface of SW1 which is connected to R3's F0/0 and a different MAC-Address is bound to this IP address, R3 is still able to communicate with the other hosts in VLAN 100 using its current MAC-Address of 0000.3333.3333. This is because, the static entry is checked first, and if the correct binding is NOT found, the DHCP Snooping Database is consulted, and obviously the correct binding is found there. To see the correct binding:

On SW1:

```
SW1#Sh ip dhcp snooping binding

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  --------------
00:00:22:22:22:22   10.1.1.2         99279       dhcp-snooping  100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1         98832       dhcp-snooping  100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3         99403       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 3
```

To test this fact, let's remove the configured MAC-Address of R3:

On R3:

```
R3(config)#int f0/0
R3(config-if)#No mac-address

R3#Ping 10.1.1.2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note: The ping was not sucessful, because R3 is using its real MAC-Address which is NOT configured statically nor in the snooping DB.

In the following configuration, R3's F0/0 is configured with the MAC-Address used in the static binding, which should be consulted before going to the DHCP Snooping Database. Let's test and verify this configuration:

On R3:

```
R3(config)#int f0/0
R3(config-if)#mac-address 0000.1234.1234
```

You should see the following console message:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed
state to down
```

The reason the line protocol is down, is because the port-security configuration on SW1; to verify:

```
SW1#Sh inter f0/3 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/3                        err-disabled 100        auto   auto 10/100BaseTX
```

To reset the F0/3 interface of SW1:

On SW1:

```
SW1(config)#int f0/3
SW1(config-if)#Shut
SW1(config-if)#No Shut

SW1#Sh inter f0/3 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/3                        connected    100        a-full  a-100 10/100BaseTX
```

Let's test the configuration on R3:

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3#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
```

Note: The ONLY reason R3 can communicate with the other hosts in VLAN 100, is because of the static binding configured earlier, and if the static binding is removed, R3 will NOT be able to communicate with any host in VLAN 100. To test his configuration:

On SW1:

```
SW1(config)#No ip source bind 0000.0123.0123 vlan 100 10.1.1.3 inter f0/3
```

On R3:

```
R3#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

To resolve this problem, the MAC-Address of R3's F0/0 interface is changed back to 0000.3333.3333, remember that this MAC address is still in the DHCP Snooping Binding. The reason it was NOT removed is because it was configured statically.

On R3:

```
R3(config)#int f0/0
R3(config-if)#mac-address 0000.3333.3333

R3#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

To see the entry in the DHCP Snooping Database:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show ip DHCP Snooping Binding

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  ----------------
00:00:22:22:22:22   10.1.1.2         97885       dhcp-snooping  100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1         97234       dhcp-snooping  100   FastEthernet0/1
00:00:33:33:33:33   10.1.1.3         98228       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 3
```

DO NOT RELOAD THE ROUTERS OR THE SWITCH

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 5
Dynamic ARP Inspection

Task 1

Remove the following configuration from SW1's F0/1, F0/2, F0/3 interfaces and configure R4 and SW1 according to the following policy:

1. SW1's F0/1 : IP Verify source
2. SW1's F0/2 : Switchport Port-security and IP verify source Port-Security
3. SW1's F0/3 : Switchport Port-security and IP verify source Port-Security
4. SW1's F0/4 : ip dhcp snooping trust
5. The DHCP configuration.
6. Add R4's IP and MAC address to the Snooping database

To configure Item #1:

On SW1:

```
SW1(config)#Int F0/1
SW1(config-if)#No ip verify source
```

To configure Item #2:

On SW1:

```
SW1(config)#Int F0/2
SW1(config-if)#No switchport port-security
SW1(config-if)#No ip verify source port-security
```

To configure Item #3:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1(config)#Int F0/3

SW1(config-if)#No switchport port-security
SW1(config-if)#No ip verify source port-security
```

To configure Item #4:

On SW1:

```
SW1(config-if)#int f0/4
SW1(config-if)#No ip dhcp snooping trust
```

To configure Item #5:

On R4:

```
R4(config)#No ip dhcp pool TST

R4(config)#No ip dhcp excluded-address 10.1.1.254
R4(config)#No ip dhcp excluded-address 10.1.1.1
R4(config)#No ip dhcp excluded-address 10.1.1.2
```

To configure Item #6:

On SW1:

```
SW1#IP dhcp snoo bind 0000.4444.4444 vlan 100 10.1.1.4 inter f0/4 exp 1000
```

To verify the configuration:

On SW1:

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

interface FastEthernet0/1
 switchport access vlan 100
 switchport mode access
 spanning-tree portfast

SW1#Sh run int f0/2 | B inter

interface FastEthernet0/2
 switchport access vlan 100
 switchport mode access
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
 spanning-tree portfast

SW1#Sh run int f0/3 | B inter

interface FastEthernet0/3
 switchport access vlan 100
 switchport mode access
 spanning-tree portfast

SW1#Sh run int f0/4 | B inter

interface FastEthernet0/4
 switchport access vlan 100
 switchport mode access
 spanning-tree portfast
```

To test the configuration:

On R4:

```
R4#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R4#Ping 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 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

R4#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 2

Configure SW1 based on the following policy:

- SW1 MUST inspect all the ARP messages on untrusted interfaces in VLAN 100.
- SW1 MUST intercept and verify that the IP to MAC binding matches an entry in the prebuilt or statically built table.
- SW1 MUST drop the ARP packets, if the IP-to-MAC binding does NOT match the entry in the dynamic or statically built table.

This task calls for Dynamic ARP inspection or "DAI". DAI inspects all ARP packets and compares the IP to MAC binding in the incoming packets against DHCP Snooping Database, if it matches, the packets are forwarded, if it does NOT match, the packets are dropped.

On SW1:

```
SW1(config)#IP arp inspection VLAN 100

SW1#Show ip dhcp snooping binding

MacAddress          IpAddress        Lease(sec)  Type           VLAN  Interface
------------------  ---------------  ----------  -------------  ----  ------------------
00:00:22:22:22:22   10.1.1.2         93277       dhcp-snooping  100   FastEthernet0/2
00:00:11:11:11:11   10.1.1.1         93257       dhcp-snooping  100   FastEthernet0/1
00:00:44:44:44:44   10.1.1.4         99534       dhcp-snooping  100   FastEthernet0/4
00:00:33:33:33:33   10.1.1.3         93271       dhcp-snooping  100   FastEthernet0/3
Total number of bindings: 4
```

On R1 – R4:

```
Rx(config)#int f0/0
Rx(config-if)#shu
Rx(config-if)#No shut

RX#Sh arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.1.X                -   0000.XXXX.XXXX  ARPA   FastEthernet0/0
```

Note xxxx.xxxx in the above show command is the router number, for example R1 should have 0000.1111.1111, R2 should have 0000.2222.2222 and R3 should have 0000.3333.3333.

To test the DAI configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

NOTE: The MAC-Address of R1 is removed and a ping is generated:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#No Mac

R1#Ping 10.1.1.2 rep 2

Type escape sequence to abort.
Sending 2, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
..
Success rate is 0 percent (0/2)
```

You should see the following console messages on SW1:

```
%SW_DAI-4-DHCP_SNOOPING_DENY: 1 Invalid ARPs (Req) on Fa0/1, vlan
100.([0015.f9ec.3110/10.1.1.1/0000.0000.0000/10.1.1.2/01:19:29 UTC Mon Mar 1
1993])
```

The output states that DAI denied one Invalid ARP request received on F0/1 interface of the local switch for VLAN 100, with a MAC-Address of 0015.f9ec.3110 (R1's real MAC address) and an IP address of 10.1.1.1 (R1's IP Address) for host 10.1.1.2 (R2).

NOTE: The MAC address of R1 in your pod may be different.

Task 3

R1's NIC was changed, and the new Mac-address is 0000.1234.1234; Configure SW1 to allow the traffic from R1 with its new Mac address. However, if it uses the old MAC address, the traffic should still be allowed. DO NOT change the DHCP snooping table to accomplish this task.

On SW1:

```
SW1(config)#IP arp inspection Filter TST Vlan 100

SW1(config)#Arp access-list TST
SW1(config-arp-nacl)#Permit ip host 10.1.1.1 mac host 0000.1234.1234
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To test the configuration:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#Mac-address 0000.1234.1234

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
```

Note: The access-list called "TST" is looked at before the DHCP Snooping Database, since the correct MAC to IP binding was found in the Access-list, the traffic was allowed.

Let's change the MAC address back to 0000.1111.1111 and test again:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#mac-address 0000.1111.1111

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

Let's change R1's Mac-address back to the new address:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#mac-address 0000.1234.1234
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

Configure SW1 such that it does NOT use the DHCP Snooping database to provide DAI. DO NOT modify or remove the entries from the DHCP Snooping Database to accomplish this task.

On SW1:

```
SW1(config)#Arp access-list TST
SW1(config-arp-nacl)#Permit ip host 10.1.1.4 mac host 0000.4444.4444
SW1(config-arp-nacl)#Permit ip host 10.1.1.3 mac host 0000.3333.3333
SW1(config-arp-nacl)#Permit ip host 10.1.1.2 mac host 0000.2222.2222

SW1(config)#IP arp inspection filter TST Vlan 100 STATIC
```

To verify the configuration:

On SW1:

```
SW1#Show arp access-list

ARP access list TST
     permit ip host 10.1.1.1 mac host 0000.1234.1234
     permit ip host 10.1.1.4 mac host 0000.4444.4444
     permit ip host 10.1.1.3 mac host 0000.3333.3333
     permit ip host 10.1.1.2 mac host 0000.2222.2222
```

To test the configuration:

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

Let's change R1's MAC-Address to 0000.1111.1111 and test the configuration:

```
R1(config)#int f0/0
R1(config-if)#Mac-address 0000.1111.1111
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Ping 10.1.1.2 rep 2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
..
Success rate is 0 percent (0/2)
```

You should see the following console messages on SW1

On SW1:

```
%SW_DAI-4-ACL_DENY: 1 Invalid ARPs (Req) on Fa0/1, vlan
100.([0000.1111.1111/10.1.1.1/0000.0000.0000/10.1.1.2/01:34:04 UTC Mon Mar
1 1993])
```

**The "Static" keyword instructs the switch NOT to check the DHCP Snooping database.
Therefore, the ONLY table checked are the entries in the "ARP access-list" called "TST".**

Since the DHCP Snooping table is NOT used at all; DHCP snooping does NOT need to be enabled. If the "Dhcp Snooping" is configured in the CCIE lab where you configured the "Static" option with an ARP access-list, and I was the proctor, I would subtract points, and I think they should as well, you should be very meticulous and very specific with your configuration, if a command or a keyword is NOT required, IT SHOULD NOT BE CONFIGURED.

To verify:

On SW1:

```
SW1(config)#No ip dhcp Snooping
SW1(config)#No IP dhcp snooping Vlan 100
```

On R1:

```
R1#Ping 10.1.1.2 rep 1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
.
Success rate is 0 percent (0/1)
```

You should still see the following console messages on SW1, because DAI is still enabled:

On SW1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
%SW_DAI-4-ACL_DENY: 1 Invalid ARPs (Req) on Fa0/1, vlan
100.([0000.1111.1111/10.1.1.1/0000.0000.0000/10.1.1.2/01:37:23 UTC Mon Mar
1 1993])
```

To test this further:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#mac-address 0000.1234.1234
```

To test and verify:

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
```

Task 5

Erase the startup configuration of all routers and the VLAN.dat, and config.text of SW1, and reload them before proceeding to the next task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 6
MAC Access-list

Task 1

Configure the F0/1 interface of R5, R6, and R7 using the following chart. You should configure a MAC ACL to filter ARP protocol such that a static ARP must be configured in order for these routers to have the ability to ping each other.

Router	Interface	MAC address	IP address	VLAN
R5	F0/1	000A.AA55.5555	20.1.1.5 /24	200
R6	F0/1	000A.AA66.6666	20.1.1.6 /24	200
R7	F0/1	000A.AA11.1111	20.1.1.10 /24	200

********** DO NOT PING TO TEST THE CONFIGURATION **********

On R7:

```
R7(config)#Int F0/1
R7(config-if)#ip addr 20.1.1.10 255.255.255.0
R7(config-if)#mac-address 000A.AA11.1111
R7(config-if)#No shut
```

On R6:

```
R6(config)#Int f0/1
R6(config-if)#ip addr 20.1.1.6 255.255.255.0
R6(config-if)#mac-address 000A.AA66.6666
R6(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/1
R5(config-if)#ip addr 20.1.1.5 255.255.255.0
R5(config-if)#Mac-address 000A.AA55.5555
R5(config-if)#No shut
```

On SW2:

```
SW2(config)#Int range f0/5-7
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2(config-if-range)#Swi mode acc
SW2(config-if-range)#Swi acc v 200
SW2(config-if-range)#Spanning portf

*************** DO NOT PING TO TEST THE CONFIGURATION ***************
```

To configure the policy:

```
SW2(config)#Mac access-list extended TST
SW2(config-ext-macl)#deny 000A.AA00.0000 0000.00FF.FFFF any 0x806 0x0
SW2(config-ext-macl)#permit any any

Switch(config)#Int range f0/5-6 , f0/11
Switch(config-if-range)#Mac access-group TST in
```

To verify the configuration:

On R6:

```
R6#Ping 20.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 20.1.1.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.10, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Sh arp

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  20.1.1.5                0  Incomplete      ARPA
Internet  20.1.1.6                -  000a.aa66.6666  ARPA   FastEthernet0/1
Internet  20.1.1.10               0  Incomplete      ARPA
```

To test the configuration:

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6(config)#arp 20.1.1.5 000a.aa55.5555 arpa
R6(config)#arp 20.1.1.10 000a.aa11.1111 arpa

R6#Ping 20.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note: R6 is sending the ICMP packets to R5, but since R5 does NOT have an entry for R6 in its ARP cache, it will get encapsulation failure messages, unless an APR entry is statically configured for R6. To verify:

On R5:

```
R5#Debug ip packet
```

On R6:

```
R6#Ping 20.1.1.5 rep 1

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
.
Success rate is 0 percent (0/1)
```

You should see the following console message on R5:

On R5:

```
IP: tableid=0, s=20.1.1.6 (FastEthernet0/1), d=20.1.1.5 (FastEthernet0/1),
routed via RIB
IP: s=20.1.1.6 (FastEthernet0/1), d=20.1.1.5 (FastEthernet0/1), len 100,
rcvd 3
IP: tableid=0, s=20.1.1.5 (local), d=20.1.1.6 (FastEthernet0/1), routed
via RIB
IP: s=20.1.1.5 (local), d=20.1.1.6 (FastEthernet0/1), len 100, sending
IP: s=20.1.1.5 (local), d=20.1.1.6 (FastEthernet0/1), len 100,
encapsulation failed
```

To fix this problem, let's add R6's IP-to-MAC binding into R5's ARP caches:

On R5:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5(config)#arp 20.1.1.6 000a.aa66.6666 arpa
```

On R6:

```
R6#Ping 20.1.1.5

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

To have full reachability, you must repeat the same process for the other routers.

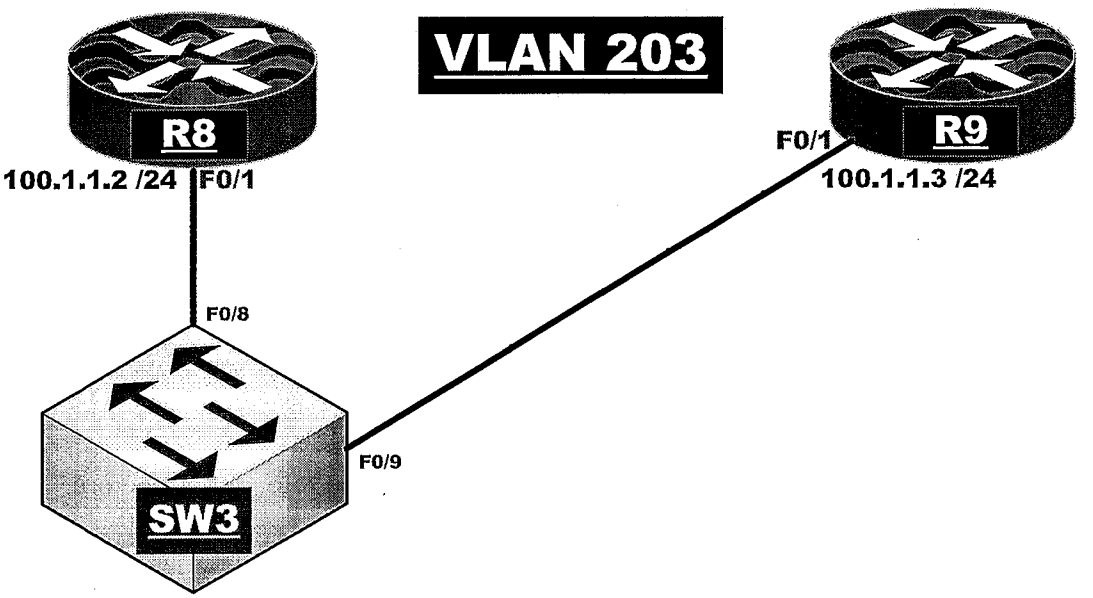
Task 2

Erase the startup configuration of these routers and SW2 and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 7
VLAN Access maps -I



Task 1

Confiure R8and R9's F0/1 interface in VLAN 203 and assign 100.1.1.2 /24 and 100.1.1.3 /24 to R8and R9's F0/1 interface respectively.

On R9:

```
R9(config)#int f0/1
R9(config-if)#ip addr 100.1.1.3 255.255.255.0
R9(config-if)#No shut
```

On R8:

```
R8(config)#int f0/1
R8(config-if)#ip addr 100.1.1.2 255.255.255.0
R8(config-if)#No shut
```

On SW3:

```
SW3(config)#Int range f0/8-9
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW3(config-if-range)#Swi mode acc
SW3(config-if-range)#Swi acc v 203
```

To test the configuration:

On R8:

```
R8#Ping 100.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
```

Task 2

Implement the following intra-VLAN policy on SW3 for VLAN 203:

- ONLY ICMP traffic is allowed within this VLAN.
- All other traffic should be dropped.
- DO NOT use any interface configuration command to accomplish this task.

This task calls for Vlan access-maps, the following describes the configuration of VLAN access-maps:

- **The ONLY way to conrol filtering within a VLAN.**
- **VACLs have NO direction, and if a direction must be specified, it should be configured in the access-list that the VACL references.**
- **The default action of VACLs is to drop, and if a Vlan access-map is configured with NO match or action statement, then, the traffic is forwarded.**
- **VACLs are order specific, meaning that the IOS reads them from the top to the bottom, and the first match wins philosophy applies, meaning that the most specific items should be configured at the top and the general entries should be configured at the bottom.**
- **If a large number of entries are configured in a given VACL, it will slow down the boot process.**
- **VACLs have no logging option.**
- **The port ACL will take precedence over the configured VACL.**

Configure VACLs should be done in three steps:

The first step is to identify the traffic, in this step the traffic MUST be permitted in an access-list or a MAC

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

access-list in order for it to be identified:

On SW3:

```
SW3(config)#IP access-list extended ICMP
SW3(config-ext-nacl)#Permit icmp any any
```

The second step is to reference the access-list in the VLAN access-map and apply a policy to it, the policy can either be to forward the traffic, or to drop the traffic:

```
SW3(config)#VLAN access-map TST 10
SW3(config-access-map)#match ip addr ICMP
SW3(config-access-map)#Action forward
```

The third and last step is to apply the policy to a given VLAN or a range of VLANs:

```
SW3(config)#Vlan filter TST vlan-list 203
```

To verify the configuration:

On SW3:

```
SW3#Show vlan access-map TST

Vlan access-map "TST" 10
  Match clauses:
    ip address: ICMP
  Action:
    Forward
```

To test the configuration:

On R8:

```
R8#Ping 100.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R8#Telnet 100.1.1.3
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Trying 100.1.1.3 ...
% Connection timed out; remote host not responding
```

Note: Since the configured policy in the VLAN access-map does NOT have a catch-all (Other traffic besides the one specified in the ACL) statement, all other traffic will fall under the default action and it will be dropped.

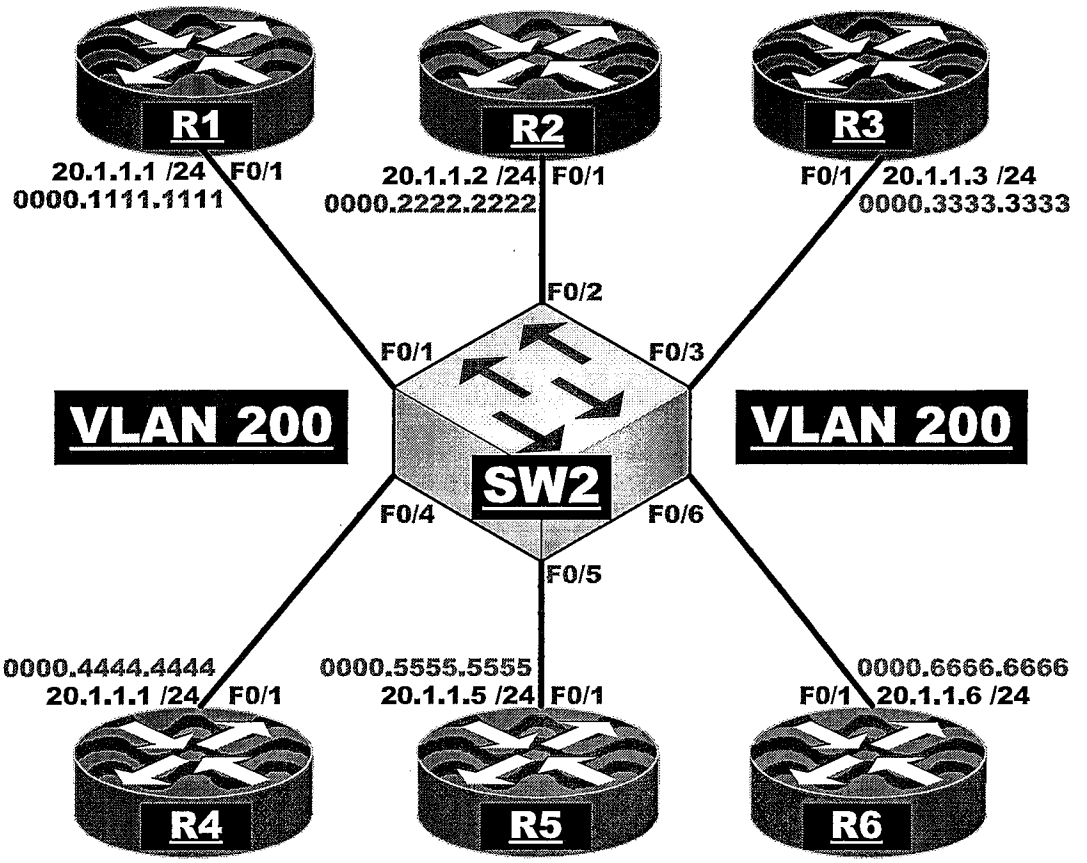
Task 3

Erase the startup configuration of these routers and SW2 and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 8
VLAN Access maps -II



IP Addressing:

Router	Interface	IP Address	MAC address	VLAN
R1	F0/1	20.1.1.1 /24	0000.1111.1111	200
R2	F0/1	20.1.1.2 /24	0000.2222.2222	200
R3	F0/1	20.1.1.3 /24	0000.3333.3333	200
R4	F0/1	20.1.1.4 /24	0000.4444.4444	200
R5	F0/1	20.1.1.5 /24	0000.5555.5555	200
R6	F0/1	20.1.1.6 /24	0000.6666.6666	200

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 1

Configure VLANs and IP address assignments based on the above chart/diagram.

On SW2:

```
SW2(config)#Int range F0/1-6
SW2(config-if-range)#Swi mode acc
SW2(config-if-range)#Swi acc v 200
SW2(config-if-range)#No shut
```

On R1:

```
R1(config)#Int F0/1
R1(config-if)#ip addr 20.1.1.1 255.255.255.0
R1(config-if)#mac-address 0000.1111.1111
R1(config-if)#No shut
```

On R2:

```
R2(config-if)#Int F0/1
R2(config-if)#ip addr 20.1.1.2 255.255.255.0
R2(config-if)#mac-address 0000.2222.2222
R2(config-if)#No shut
```

On R3:

```
R3(config-if)#Int F0/1
R3(config-if)#ip addr 20.1.1.3 255.255.255.0
R3(config-if)#mac-address 0000.3333.3333
R3(config-if)#No shut
```

On R4:

```
R4(config-if)#Int F0/1
R4(config-if)#Ip addr 20.1.1.4 255.255.255.0
R4(config-if)#mac-address 0000.4444.4444
R4(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/1
R5(config-if)#IP address 20.1.1.5 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config-if)#mac-address 0000.5555.5555
R5(config-if)#No shut
```

On R6:

```
R6(config)#Int F0/1
R6(config-if)#ip addr 20.1.1.6 255.255.255.0
R6(config-if)#Mac-address 0000.6666.6666
R6(config-if)#No shut
```

To test the configuration:

On R1:

```
R1#Ping 20.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 20.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 20.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 20.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 20.1.1.6
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 2

Configure the following Intra-VLAN policy for VLAN 200:

- SW2 should be configured to drop all traffic between R1 and R2
- SW2 should be configured to drop all traffic between R3 and R4
- SW2 should be configured to drop all traffic between R5 and R6
- All other traffic should be allowed.
- R6 should NOT be able to perform an ARP request/response.

On SW2:

Step 1: Identify the traffic:

```
SW2(config)#Access-list 100 permit ip host 20.1.1.1 host 20.1.1.2
SW2(config)#Access-list 100 permit ip host 20.1.1.2 host 20.1.1.1
SW2(config)#Access-list 100 permit ip host 20.1.1.3 host 20.1.1.4
SW2(config)#Access-list 100 permit ip host 20.1.1.4 host 20.1.1.3
SW2(config)#Access-list 100 permit ip host 20.1.1.5 host 20.1.1.6
SW2(config)#Access-list 100 permit ip host 20.1.1.6 host 20.1.1.5

SW2(config)#Mac access-list extended tst
SW2(config-ext-macl)#permit host 0000.6666.6666 any 0x806 0x0
```

Reference the identified traffic in the VLAN Access-map and assign a policy to it:

```
SW2(config)#Vlan access-map TST 10
SW2(config-access-map)#match ip addr 100
SW2(config-access-map)#action drop

SW2(config)#vlan access-map TST 20
SW2(config-access-map)#match mac address R6
SW2(config-access-map)#action drop

SW2(config)#vlan access-map TST 90
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Note in this case we need the catch-all Vlan access-map:

```
SW2(config)#Vlan access-map TST 90
```

Apply the policy to a VLAN:

```
SW2(config)#Vlan filter TST vlan-list 200
```

To test and verify the configuration:

On R1:

```
R1#Ping 20.1.1.2 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.2, timeout is 2 seconds:
...
Success rate is 0 percent (0/3)

R1#Ping 20.1.1.3 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.3, timeout is 2 seconds:
.!!
Success rate is 66 percent (2/3), round-trip min/avg/max = 1/1/1 ms

R1#Ping 20.1.1.4 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
.!!
Success rate is 66 percent (2/3), round-trip min/avg/max = 1/1/1 ms

R1#Ping 20.1.1.5 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
.!!
Success rate is 66 percent (2/3), round-trip min/avg/max = 1/1/1 ms

R1#Ping 20.1.1.6 rep 3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

On R3:

```
R3#Ping 20.1.1.4 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
...
Success rate is 0 percent (0/3)

R3#Ping 20.1.1.5 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.5, timeout is 2 seconds:
.!!
Success rate is 66 percent (2/3), round-trip min/avg/max = 1/1/1 ms

R3#Ping 20.1.1.6 rep 3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R5:

```
R5#Ping 20.1.1.6 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
...
Success rate is 0 percent (0/3)

R5#Ping 20.1.1.4 rep 3

Type escape sequence to abort.
Sending 3, 100-byte ICMP Echos to 20.1.1.4, timeout is 2 seconds:
.!!
Success rate is 66 percent (2/3), round-trip min/avg/max = 1/1/1 ms

R5#Ping 20.1.1.1 rep 3

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 3, 100-byte ICMP Echos to 20.1.1.1, timeout is 2 seconds:
!!!
Success rate is 100 percent (3/3), round-trip min/avg/max = 1/1/1 ms
```

Since R6 can NOT send any ARP replies or respond to one, you must configure static entry/ies in the ARP table of these routers:

On R6:

```
R6(config)#arp 20.1.1.1 0000.1111.1111 arpa
R6(config)#arp 20.1.1.2 0000.2222.2222 arpa
R6(config)#arp 20.1.1.3 0000.3333.3333 arpa
R6(config)#arp 20.1.1.4 0000.4444.4444 arpa
```

Since the other routers don't know the MAC address of R6, the other routers must also be configured with a static route for the MAC address of R6.

On R1:

```
R1(config)#arp 20.1.1.6 0000.6666.6666 arpa
```

To test the configuration:

```
R1#Ping 20.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R2:

```
R2(config)#arp 20.1.1.6 0000.6666.6666 arpa
```

To test the configuration:

```
R2#Ping 20.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config)#arp 20.1.1.6 0000.6666.6666 arpa
```

To test the configuration:

```
R3#Ping 20.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R4:

```
R4(config)#arp 20.1.1.6 0000.6666.6666 arpa
```

To test the configuration:

```
R4#Ping 20.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R5:

Remember that R5 and R6 are NOT allowed to communicate.

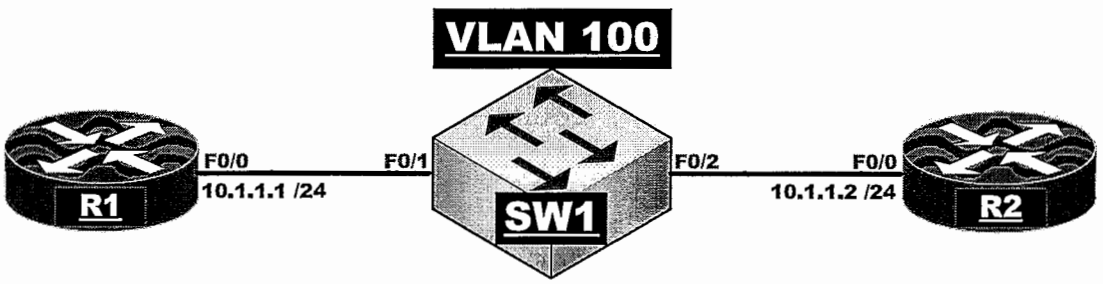
Task 3

Erase the startup configuration of these routers and SW2 and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 9
Storm Control - I



Task 1

Configure the routers and SW1 according to the above diagram.

On R1:

```
R1(config)#Int F0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0
R1(config-if)#No shut
```

On R2:

```
R2(config)#Int F0/0
R2(config-if)#ip addr 10.1.1.2 255.255.255.0
R2(config-if)#No shut
```

On SW1:

```
SW1(config)#Int range f0/1-2
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc v 100
SW1(config-if-range)#Spanning portf
SW1(config-if-range)#No shut
```

To verify and test the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
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
```

Task 2

Configure the F0/1 interface of SW1 such that it NEVER allows incoming Unicast traffic to exceed .5 percent of its interface's bandwidth. If traffic exceeds this policy, the F0/1 interface of SW1 should shutdown in "**err-disable**" mode and stay in this mode until the administrator of the switch re-enables this port.

On SW1:

```
SW1(config)#Int F0/1
SW1(config-if)#Storm-control Unicast level 0.5
SW1(config-if)#Storm-control action Shutdown
```

To verify the configuration:

On SW1:

```
SW1#Show storm-control unicast

Interface  Filter State   Upper         Lower         Current
---------  -------------  ------------  ------------  -----------
Fa0/1      Forwarding          0.50%         0.50%         0.00%
```

To test the configuration:

R1 will generate traffic using a Ping command, and if the amount of generated traffic exceeds half percent of the interfaces bandwidth, the interface should transition into "err-disable" mode.

The following ping will generate the desired amount of traffic:

On R1:

```
R1#Ping 10.1.1.10 repeat 1000 size 1500
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 1000, 1500-byte ICMP Echos to 10.1.1.10, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!..
```

After the timeouts (.) you should see the following console message:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed
state to down
```

This means that the F0/1 interface of SW1 is in "Err-disable" mode, this interface is in this mode because the traffic exceeded the configured threshold. To verify:

On SW1:

You should see the following console messages on SW1:

```
%PM-4-ERR_DISABLE: storm-control error detected on Fa0/1, putting Fa0/1 in
err-disable state

%STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/1. The
interface has been disabled.

%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down

%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down

SW1#Show interface status err-disabled

Port      Name               Status       Reason
Fa0/1                        err-disabled storm-control
```

Note the output of the above show command reveals that F0/1 interface is in err-disable mode because of Storm-control violation.

On SW1:

```
SW1#Show storm-control unicast

Interface  Filter State   Upper         Lower         Current
---------  -------------  ------------  ------------  ----------
Fa0/1      Link Down          0.50%         0.50%         0.00%
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 3

Re-configure SW1 such that it automatically recovers from err-disable mode every 30 seconds, if the configured storm-control policy from the previous task is violated, it should transition back into err-disable mode; however, if the policy is NOT violated, the state of this interface should remain in UP/UP state.

On SW1:

```
SW1(config)#errdisable recovery cause storm-control
SW1(config)#errdisable recovery interval 30
```

You should see the following console messages. The switch recovers from the storm control violation that placed the F0/1 interface into err-disable mode:

```
00:40:16: %PM-4-ERR_RECOVER: Attempting to recover from storm-control err-disable state on Fa0/1

00:40:20: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
00:40:21: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
```

Since the existing traffic exceeds the set threshold, the interface transitions into err-disable mode again:

```
%PM-4-ERR_DISABLE: storm-control error detected on Fa0/1, putting Fa0/1 in err-disable state

%STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/1. The interface has been disabled.

%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down

%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
```

And the cycle repeats:

```
%PM-4-ERR_RECOVER: Attempting to recover from storm-control err-disable state on Fa0/1

%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%PM-4-ERR_DISABLE: storm-control error detected on Fa0/1, putting Fa0/1 in
err-disable state

%STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/1. The
interface has been disabled.

%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down

%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
```

Let's stop the Pings being generated by R1. This is done by pressing and holding the "Shift" and the "control" key on the keyboard while pressing the "6" key twice.

You should see the following console message on SW1, and you can see that the interface is recovered once again and it remains in UP/UP state.

```
%PM-4-ERR_RECOVER: Attempting to recover from storm-control err-disable state on Fa0/1

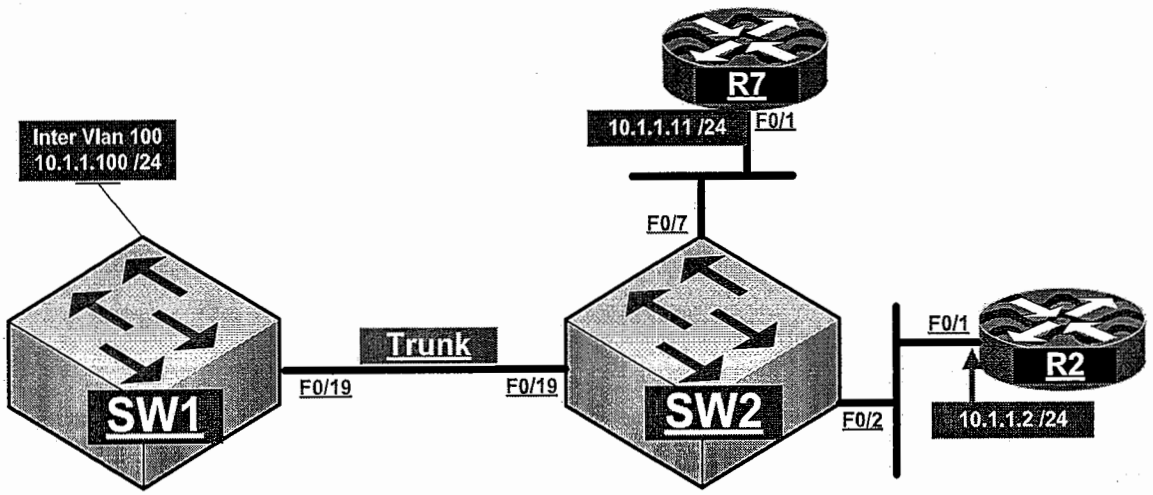
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
```

Task 4

Erase the startup and VLAN.DAT of these devices and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 10
Storm Control - II



Task 1

Configure the routers and switches based on the following diagram, the ports NOT referenced in the diagram should be disable.

On R2:

```
R2(config)#Int F0/1
R2(config-if)#ip addr 10.1.1.2 255.255.255.0
R2(config-if)#No Shut
```

On R7:

```
R7(config)#Int F0/1
R7(config-if)#ip addr 10.1.1.11 255.255.255.0
R7(config-if)#No Shut
```

On SW1:

```
SW1(config)#Int Range F0/1-24
SW1(config-if-range)#Shut

SW1(config)#Int F0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1(config-if)#Swi trunk enca Dot

SW1(config-if)#Swi mode trunk
SW1(config-if)#No Shut

SW1(config)#Vlan 100
SW1(config-vlan)#exit

SW1(config)#Int Vlan 100
SW1(config-if)#IP address 10.1.1.100 255.255.255.0
```

On SW2:

```
SW2(config)#Int Range f0/1-24
SW2(config-if-range)#Shut

SW2(config)#Int F0/19
SW2(config-if)#Swi tru enca Dot
SW2(config-if)#Swit mode trun
SW2(config-if-range)#No shut

SW2(config)#Int Range F0/2 , F0/11
SW2(config-if-range)#Swi mode acc
SW2(config-if-range)#Swi acc V 100
SW2(config-if-range)#No shut
```

To test he configuration:

On R2:

```
R2#Ping 10.1.1.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.100, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms

R2#Ping 10.1.1.11

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.11, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R7:

```
R7#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 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

R7#Ping 10.1.1.100

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.100, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

Task 2

Configure the F0/19 interface of SW1 such that it NEVER allows the incoming Unicast traffic to exceed 830 packets per second. This port should be configured such that if the number of packets per second matches or exceeds the set threshold, the port transitions into err-disable mode.

On SW1:

```
SW1(config)#Int F0/19
SW1(config-if)#Storm-control Unicast level pps 830
SW1(config-if)#Storm-control action Shutdown
```

To test the configuration:

On SW1:

```
SW1#Show storm-control Unicast

Interface  Filter State   Upper         Lower         Current
---------  -------------  ------------  ------------  ----------
Fa0/19     Forwarding        830 pps       830 pps         0 pps
```

On R7:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7#Ping 10.1.1.100 repeat 1000
```

On SW1:

Note: If the number of packets per second (pps) does NOT increment, it is the IOS and it has nothing to do with your configuration. If you continue with this lab and the F0/19 interface shuts down, it proves that the feature actually worked, but the show command did not. You MUST do the following show command quickly after the ping command is entered on the R7 router to see the currect counter.

```
SW1#Show storm-control Unicast

Interface  Filter State   Upper         Lower         Current
---------  -------------  ------------  ------------  ----------
Fa0/19     Forwarding        1.6k pps      1.6k pps      420 pps
```

R7 is generating around 400 to 750 pps. Therefore, let's ping with a higher repeat count so the number of packets per second exceeds the set threshold, once this happens the port should transition into err-disable state.

On R7:

```
R7#Ping 10.1.1.100 repeat 100000
```

On SW1:

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

%STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/19. The
interface has been disabled.

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

%LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to
down

%LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan100, changed state to
down

%LINK-3-UPDOWN: Interface FastEthernet0/19, changed state to down
```

To verify the state of the F0/19 interface on SW1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW1:

```
SW1#Show interface F0/19 status err-disabled

Port      Name               Status       Reason              Err-disabled Vlans
Fa0/19                       err-disabled storm-control
```

Task 3

Remove the Storm-control configuration from the F0/19 interface of SW1 and ensure that the interface is in "**Connected**" state.

On SW1:

```
SW1(config)#Int F0/19
SW1(config-if)#No storm-control Unicast level pps 830
SW1(config-if)#No storm-control Action Shutdown
```

On SW1:

```
SW1(config)#Int F0/19
SW1(config-if)#Shut
SW1(config-if)#No Shut
```

To verify the configuration:

On SW1:

```
SW1#Show run int F0/19 | B inter

interface FastEthernet0/19
 switchport trunk encapsulation dot1q
 switchport mode trunk

SW1#Sh Inter F0/19 Status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/19                       connected    trunk      a-full  a-100 10/100BaseTX
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 4

Configure F0/19 interface of SW1 such that it blocks Unicast traffic all together.

On SW1:

```
SW1(config)#int f0/19
SW1(config-if)#storm-control Unicast level 0 0
```

You should receive the following console message:

```
%STORM_CONTROL-3-FILTERED: A Unicast storm detected on Fa0/19. A packet
filter action has been applied on the interface.
```

To verify the configuration:

On SW1:

```
SW1#Show storm-control Unicast

Interface  Filter State    Upper          Lower          Current
---------  -------------   -----------    -----------    ----------
Fa0/19     Blocking          0.00%          0.00%          0.00%
```

To test the configuration:

Configure R1's F0/0 interface in Vlan 100 and assign an IP address of 10.1.1.1 /24 to this interface:

On SW1:

```
SW1(config)#Int F0/1
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc v 100
SW1(config-if)#Spanning portf
SW1(config-if)#No Shut
```

On R1:

```
R1(config)#Int F0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0
R1(config-if)#No shut
```

Configure a Lo0 and assign an IP address of 1.1.1.1 /8 to this interface:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config)#Int Lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```

Configure a Lo0 on R2 and assign an IP address of 2.2.2.2 /8 to this interface:

On R2:

```
R2(config)#Int Lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0
```

Since the "Storm-control" applies the configured suppression level for Unicast traffic inbound F0/19 on SW1, it will block all Unicast traffic while allowing Broadcast and Multicast through.

In this test, RIPv2 will be configured on R1 and R2 to test the Unicast, and Multicast traffic, as follows:

On R2:

```
R2(config)#router rip
R2(config-router)#ver 2
R2(config-router)#passive-interface F0/1
R2(config-router)#Neighbor 10.1.1.1
R2(config-router)#Netw 10.0.0.0
R2(config-router)#Netw 2.0.0.0
```

On R1:

```
R1(config)#Router rip
R1(config-router)#ver 2
R1(config-router)#passive-interface F0/0
R1(config-router)#Neighbor 10.1.1.2
R1(config-router)#Netw 10.0.0.0
R1(config-router)#Netw 1.0.0.0
```

Note R2 is configured to send Unicast updates to R1 which is blocked by the F0/19 interface of SW1, but since storm-control is NOT configured on the interface facing R1, R2 will receive R1's Network, but R1 will NOT receive R2's Network/s.

To verify:

On R2:

```
R2#Show ip route rip
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R    1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:06, FastEthernet0/1
```

On R1:

```
R1#Show ip route rip
R1#
```

The above test verifies that all ingress Unicast traffic is blocked by SW1's F0/19 interface.

Task 5

Remove the Storm-control configuration from the F0/19 interface of SW1 and ensure that the interface is in "Connected" state. You should also remove the RIP configuration from R1 and R2.

On SW1:

```
SW1(config)#int f0/19
SW1(config-if)#No storm-control Unicast level 0 0
```

On R1 and R2:

```
Rx(config)#No router rip
```

Task 6

Configure the F0/19 interface of SW1 such that it blocks Unicast traffic that exceeds 800Kbps. If the traffic exceeds the set threshold, the port should transition into "**Err-disable**" mode.

On SW1:

```
SW1(config)#Int F0/19
SW1(config-if)#storm-control unicast level bps 800k
SW1(config-if)#storm-control action shutdown
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On SW1:

```
SW1#Show storm-control unicast

Interface  Filter State   Upper         Lower         Current
---------  -------------  ------------  ------------  ----------
Fa0/19     Forwarding        800k bps      800k bps       0 bps
```

To test the configuration:

On R2:

```
R2#Ping 10.1.1.1 rep 100000
```

On SW1:

Note, the traffic exceeded the configured threshold and as a result of that the F0/19 interface transitioned into "Err-Disable" state.

You should see the following message stating that the F0/19 is in "Err-disable" state because the Unicast traffic exceeded the set threshould.

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

%STORM_CONTROL-3-SHUTDOWN: A packet storm was detected on Fa0/19. The
interface has been disabled.

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

%LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to
down

%LINK-3-UPDOWN: Interface FastEthernet0/19, changed state to down
```

To verify the test:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show storm-control unicast

Interface  Filter State    Upper         Lower         Current
---------  -------------  -----------   -----------   ----------
Fa0/19     Link Down          800k bps      800k bps        0 bps

SW1#Show inter F0/19 status err-disabled

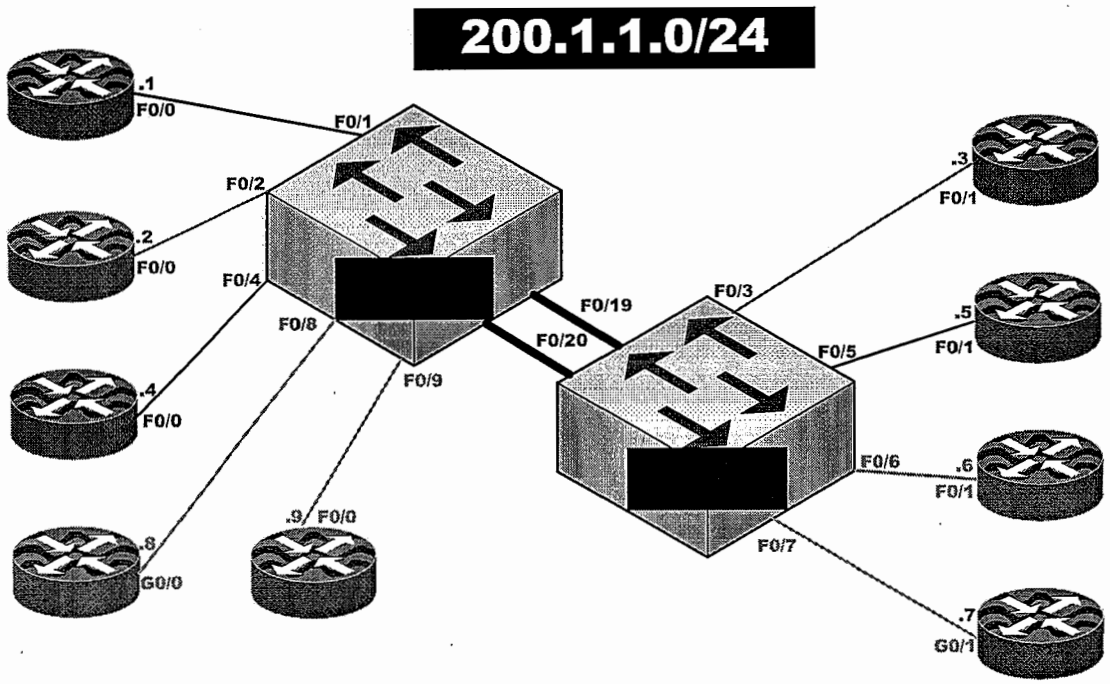
Port      Name               Status       Reason              Err-disabled Vlans
Fa0/19                       err-disabled storm-control
```

Task 7

Erase the startup configuration and vlan.dat and reload the switches and the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 11
Configuring Private VLANs



Task 1

The first switch should be configured with a hostname of SW1 and the second switch should be configured with a hostname of SW2.

On the First Switch

```
Switch(config)#Hostname SW1
```

On the Second Switch

```
Switch(config)#Hostname SW2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 2

Shutdown ports F0/21-24 on SW1 and SW2.

On Both Switches:

```
SWx(config)#int range f0/21-24
SWx(config-if-range)#Shut
```

Task 3

Configure trunking between SW1 and SW2 using ports F0/19 and F0/20. Use an industry standard trunking protocol for this purpose. Assign a brief meaningful description to these interfaces.

On SW1:

```
SW1(config)#Interface range f0/19-20
SW1(config-if-range)#Switch trunk encap dot1q

SW1(config-if-range)#Switch mode trunk
SW1(config-if-range)#Description Trunk to SW2
```

On SW2:

```
SW2(config)#Interface range f0/19-20
SW2(config-if-range)#Switch trunk encap dot1q
SW2(config-if-range)#Switch mode trunk
SW2(config-if-range)#Description Trunk to SW1
```

To verify the configuration:

On SW1:

```
SW1#Show int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

Port        Vlans allowed and active in management domain
Fa0/19      1
Fa0/20      1

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
Fa0/20      none
```

On SW2:

```
SW2#Show int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/20      on           802.1q         trunking      1

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

Port        Vlans allowed and active in management domain

Fa0/19      1
Fa0/20      1
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
Fa0/20      1
```

Task 4

Assign IP addressing to the interface of the routers using the following chart and ensure that these routers can ping each other: You should assign a brief meaningful interface description on the switchports.

Router	Interface	IP address and Subnet mask
R1	F0/0	200.1.1.1 /24
R2	F0/0	200.1.1.2 /24
R3	**F0/1**	**200.1.1.3 /24**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R4	F0/0	200.1.1.4 /24
R5	**F0/1**	**200.1.1.5 /24**
R6	**F0/1**	**200.1.1.6 /24**
R7	**F0/1**	**200.1.1.7 /24**
R8	F0/0	200.1.1.8 /24
R9	F0/0	200.1.1.9 /24

On R1:

```
R1(config)#Int F0/0
R1(config-if)#Ip address 200.1.1.1 255.255.255.0
R1(config-if)#No shut
```

On R2:

```
R2(config)#Int F0/0
R2(config-if)#Ip address 200.1.1.2 255.255.255.0
R2(config-if)#No shut
```

On R3:

```
R3(config)#Int F0/1
R3(config-if)#Ip address 200.1.1.3 255.255.255.0
R3(config-if)#No shut
```

On R4:

```
R4(config)#Int F0/0
R4(config-if)#Ip address 200.1.1.4 255.255.255.0
R4(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/1
R5(config-if)#Ip address 200.1.1.5 255.255.255.0
R5(config-if)#No shut
```

On R6:

```
R6(config)#Int F0/1
R6(config-if)# Ip address 200.1.1.6 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R6(config-if)#No shut
```

On R7:

```
R7(config)#Int G0/1
R7(config-if)# Ip address 200.1.1.7 255.255.255.0
R7(config-if)#No shut
```

On R8:

```
R8(config)#int G0/0
R8(config-if)#ip address 200.1.1.8 255.255.255.0
R8(config-if)#No shut
```

On R9:

```
R9(config)#int F0/0
R9(config-if)#ip address 200.1.1.9 255.255.255.0
R9(config-if)#No shut
```

On SW1:

```
SW1(config)#Int F0/1
SW1(config-if)#Description R1's F0/0

SW1(config)#Int F0/2
SW1(config-if)#Description R2's F0/0

SW1(config)#Int range F0/3 , F0/5-9 , F0/12-18 , F0/21-24
SW1(config-if-range)#Description --

SW1(config)#Int F0/4
SW1(config-if)#Description R4's F0/0

SW1(config)#Int F0/12
SW1(config-if)#Description R8's F0/0

SW1(config)#Int F0/13
SW1(config-if)#Description R9's F0/0
```

On SW2:

```
SW2(config)#Int range F0/1-2 , F0/4 , F0/10-18 , F0/21-24
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2(config-if-range)#Description --

SW2(config)#Int F0/3
SW2(config-if)#Description R3's F0/1

SW2(config)#Int F0/5
SW2(config-if)#Description R5's F0/1

SW2(config)#Int F0/6
SW2(config-if)#Description R6's F0/1

SW2(config)#Int F0/7
SW2(config-if)#Description R7's F0/1
```

To test and verify the configuration:

On R1:

```
R1#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
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

```
R1#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.9

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 5

Configure the switches such that the ports that are not used are in Administratively down state. Use minimum number of commands for this task.

On SW1:

```
SW1(config)#int range F0/3 , F0/5 , F0/10, F0/14-18 , F0/21-24
SW1(config-if-range)#Shut
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1#Sh int status | Inc Port|connected
```

```
Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/1     R1's F0/0          connected    1          a-full  a-100 10/100BaseTX
Fa0/2     R2's F0/0          connected    1          a-full  a-100 10/100BaseTX
Fa0/4     R4's F0/0          connected    1          a-full  a-100 10/100BaseTX
Fa0/8     R8's F0/0          connected    1          a-full  a-100 10/100BaseTX
Fa0/9     R9's F0/0          connected    1          a-full  a-100 10/100BaseTX
Fa0/19    Trunk to SW2       connected    trunk      a-full  a-100 10/100BaseTX
Fa0/20    Trunk to SW2       connected    trunk      a-full  a-100 10/100BaseTX
```

On SW2:

```
SW2(config)#int range F0/1-2 , F0/4 , F0/8-10, F0/12-18 , F0/21-24
SW2(config-if)#Shut
```

To verify the configuration:

On SW2:

```
SW2#Sh int status | Inc Port|connected
```

```
Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/3     R3's F0/1          connected    1          a-full  a-100 10/100BaseTX
Fa0/5     R5's F0/1          connected    1          a-full  a-100 10/100BaseTX
Fa0/6     R6's F0/1          connected    1          a-full  a-100 10/100BaseTX
Fa0/7     R7's F0/1          connected    1          a-full  a-100 10/100BaseTX
Fa0/19    Trunk to SW1       connected    trunk      a-full  a-100 10/100BaseTX
Fa0/20    Trunk to SW1       connected    trunk      a-full  a-100 10/100BaseTX
```

The interface description can be extremely helpful specially if the switches are configured in transparent mode, and/or the task asks for the configuration of allowed VLANs on the trunks.

Task 6

Configure Private VLANs based on the following policy:

Router	Interface	VLAN-Type	VLAN-ID
R1	F0/0	Primary	10
R2	F0/0	Community	20

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

R3	F0/1	Community	20
R4	F0/0	Community	30
R5	F0/1	Community	30
R6	F0/1	Isolated	40
R7	F0/1	Isolated	40
R8	F0/0	Isolated	40
R9	F0/0	Isolated	40

Private-VLANs are typically seen in service provider networks, this feature addresses two major problems that the providers used to face:

1. Number of Clients: If every client was in a VLAN of their own, the provider will be restricted to 4094 clients, which is the maximum number of VLANs on a given switch.

2. Routing between VLANs & IP addressing: Routing between VLANs will be a nightmare, and the number of wasted IP addresses that result from Subnetting will be enormous.

Private-VLANs solves these two issues, with Private-VLANs a VLAN is sub-divided into sub-VLANs or sub-domains.

Private-VLANs consist of one primary, and one or more secondary VLANs, the secondary VLANs can be either Community VLANs or Isolated VLANs.

A Primary VLAN can have many Community VLANs, but it can ONLY have a Single Isolated VLAN.

Ports in a Private-VLAN:

There are three types of ports in Private-VLAN and they are as follows:

1. Promiscuous: A promiscuous port belongs to the primary VLAN; this port can communicate with all ports that are member of a secondary VLAN/s (Community and/or Isolated) that are associated with the primary VLAN that it belongs.

2. Isolated: An isolated port is a host port that belongs to an isolated secondary VLAN. The host ports that are member of a given Isolated VLAN can NOT Communicate with each other. These ports can ONLY communicate with the Port configured as Promiscuous port.

3. Community: A community port is a host port that belongs to a community Secondary VLAN. Community ports can communicate with ports in the same Community VLAN and with the port

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

that is configured as promiscuous ports. These ports can't Communicate with other ports in other Community VLANs.

On Both Switches:

In order to configure private-vlans, the switches must be configured in Transparent mode as follows:

```
SWx(config)#vtp mode transparent
```

The following commands configures the primary VLAN

```
SWx(config)#vlan 10
SWx(config-vlan)#private-vlan primary
SWx(config-vlan)#Exit
```

The following two VLANs are defined as the community secondary VLANs, there could be many community VLANs:

```
SWx(config)#vlan 20
SWx(config-vlan)#private-vlan community

SWx(config)#vlan 30
SWx(config-vlan)#private-vlan community
```

There can ONLY be one isolated secondary VLAN:

```
SWx(config)#vlan 40
SWx(config-vlan)#private-vlan isolated
```

The following command associates the secondary VLANs to the primary:

```
SWx(config)#vlan 10
SWx(config-vlan)#private-vlan association add 20,30,40
```

To verify the configuration:

On Both Switches:

```
SWx#Show vlan private-vlan

Primary  Secondary  Type              Ports
-------  ---------  ----------------- ----------------
10       20         community
10       30         community
```


© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

```
10       40          isolated
```

The output of the above show command displays the secondary VLANs that are created so far and the primary VLAN to which they are associated.

On SW1:

The following command sets F0/1 interface in promiscuous mode, assigns the port to primary VLAN 10 and maps VLANs 20, 30 and 40 to this interface:

```
SW1(config)#Int F0/1
SW1(config-if)#Switchport mode private-vlan promiscuous
SW1(config-if)#Switchport private-vlan mapping 10 add 20,30,40
```

The ports that belong to a given secondary VLAN must be configured in host mode. The following command sets F0/2 interface in a host mode, associates this port to VLAN 10 (The primary VLAN) and assigns this port to VLAN 20 which was configured as a community secondary VLAN earlier:

```
SW1(config-if)#Int F0/2
SW1(config-if)#Switchport mode private-vlan host
SW1(config-if)#Switchport private-vlan host-association 10 20
```

The following command sets F0/4 interface in a host mode, associates this port to VLAN 10 (The primary VLAN) and assigns this port to VLAN 30 which was configured as a community secondary VLAN earlier:

```
SW1(config-if)#Int F0/4
SW1(config-if)#Switchport mode private-vlan host
SW1(config-if)#switchport private-vlan host-association 10 30
```

The following command sets F0/8 and F0/9 interfaces in a host mode, associates these ports to VLAN 10 (The primary VLAN) and assigns these ports to VLAN 40 which was configured as an isolated secondary VLAN earlier:

```
SW1(config)#Int range F0/8-9
SW1(config-if)#Switchport mode private-vlan host
SW1(config-if)#Switchport private-vlan host-association 10 40
```

To verify the configuration:

On SW1:

```
SW1#Sh vlan pri
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Primary Secondary Type                Ports
------- --------- ------------------ --------------------------
10      20        community          Fa0/1, Fa0/2
10      30        community          Fa0/1, Fa0/4
10      40        isolated           Fa0/1, Fa0/8, Fa0/9
```

On SW2:

```
SW2(config)#Int F0/3
SW2(config-if)#Switchport mode private-vlan host
SW2(config-if)#Switchport private-vlan host-association 10 20

SW2(config)#Int F0/5
SW2(config-if)#Switchport mode private-vlan host
SW2(config-if)#Switchport private-vlan host-association 10 30

SW2(config)#Int range F0/6 , F0/7
SW2(config-if)#Switchport mode private-vlan host
SW2(config-if)#switchport private-vlan host-association 10 40
```

To verify the configuration:

On SW2:

```
SW2#Show vlan private-vlan

Primary Secondary Type                Ports
------- --------- ------------------ ----------------
10      20        community          Fa0/3
10      30        community          Fa0/5
10      40        isolated           Fa0/6, Fa0/7
```

To test the configuration:

On R1:

```
R1#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1#Ping 200.1.1.3
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, 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 = 1/2/4 ms
```

Note R1 is able to ping all routers because it is configured to be in promiscuous mode, this interface can be thought of as the default gateway.

On R2:

```
R2#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R2#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note R2 is able to ping R1 which is the port in the primary VLAN and R3 which is in the same community VLAN. R2 can NOT communicate with the hosts in the other secondary VLANs. The following verifies this information:

```
R2#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

R2#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R3:

```
R3#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note R3 is able to ping R1 which is the port in primary VLAN and the router in its own community secondary VLAN, which is R2.

```
R3#Ping 200.1.1.4

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.10, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.10, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note R3 can NOT ping the other routers because they are in another secondary VLAN.

On R4:

```
R4#Ping 200.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R4#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note R4 is able to ping R1 which is the port in primary VLAN and the router in its own community secondary VLAN, which is R5.

```
R4#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.8
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R4#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note R4 can NOT ping the other routers because they are in another secondary VLAN.

On R5:

```
R5#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R5#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Note R5 is able to ping R1 which is the port in primary VLAN and the router in its own community secondary VLAN (R2).

```
R5#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

R5#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R5#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Note R5 can NOT ping the other routers because they are in another secondary VLAN.

On R6:

```
R6#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Note R6 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R7, R8 and R9 are in the same VLAN, but remember that the VLAN is defined as isolated; the hosts in isolated VLAN do NOT have reachability to each other.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R6#Ping 200.1.1.9

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R7:

```
R7#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

R7 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R6, R8and R9 are in the same VLAN, but remember that the VLAN is defined as an isolated secondary VLAN; the hosts in isolated VLAN do NOT have reachability to each other.

```
R7#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R7#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R7#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R7#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R7#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R7#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R8:

```
R8#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

R8 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R6, R7 and R9 are in the same VLAN, but remember that the VLAN is defined as an isolated secondary VLAN; the hosts in isolated VLAN do NOT have reachability to each other.

```
R8#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

R8#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R8#Ping 200.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.9, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R9:

```
R9#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, 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
```

R9 is able to ping R1 which is the port in primary VLAN but it can NOT ping any other router, even though R6, R7 and R8 are in the same VLAN, but remember that the VLAN is defined as an isolated secondary VLAN; the hosts in isolated VLAN do NOT have reachability to each other.

```
R9#Ping 200.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R9#Ping 200.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R9#Ping 200.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R9#Ping 200.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R9#Ping 200.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.6, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R9#Ping 200.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.7, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

R9#Ping 200.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.8, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Task 7

Reconfigure the IP addressing of the hosts that belong to the two community secondary VLANs based on the following chart and provide InterVlan routing between them: The hosts in the other secondary VLANs should still be able to reach the host in the primary VLAN. You can use static routes and any IP addressing to accomplish this task.

Routers / Interface	IP address	VLAN-ID
R2 – F0/0	202.1.1.2 /24	20
R3 – F0/1	202.1.1.3 /24	20
R4 – F0/0	**203.1.1.4 /24**	**30**
R5 – F0/1	**203.1.1.5 /24**	**30**

On R2:

```
R2(config)#int f0/0
R2(config-if)#ip addr 202.1.1.2 255.255.255.0

R2(config)#ip route 0.0.0.0 0.0.0.0 202.1.1.100
```

On R3:

```
R3(config)#int f0/1
R3(config-if)#ip addr 202.1.1.3 255.255.255.0

R3(config)#ip route 0.0.0.0 0.0.0.0 202.1.1.100
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#int f0/0
R4(config-if)#ip addr 203.1.1.4 255.255.255.0

R4(config)#ip route 0.0.0.0 0.0.0.0 203.1.1.100
```

On R5:

```
R5(config)#int f0/1
R5(config-if)#ip addr 203.1.1.5 255.255.255.0

R5(config)#ip route 0.0.0.0 0.0.0.0 203.1.1.100
```

On SW1:

```
SW1(config)#IP routing
```

Note two IP addresses are configured under interface VLAN 10, a primary and a secondary, the primary IP address is used by the hosts in VLAN 20 and the secondary is used by the hosts in VLAN 30. The "Private-vlan mapping" command maps the secondary VLANs to their layer 3 VLAN interface, in this case VLAN 10 which is the layer 3 interface of the primary VLAN.

```
SW1(config)#int vlan 10
SW1(config-if)#ip address 202.1.1.100 255.255.255.0
SW1(config-if)#ip address 203.1.1.100 255.255.255.0 sec
SW1(config-if)#private-vlan mapping 20,30
```

With the "Private-vlan mapping" interface configuration command, secondary VLANs can be added or removed using the "Private-vlan mapping add, or Private-vlan mapping remove" interface configuration command. After this command is entered, you should get the following messages:

```
%PV-6-PV_MSG: Created a private vlan mapping, Primary 10, Secondary 20
%PV-6-PV_MSG: Created a private vlan mapping, Primary 10, Secondary 30
```

To verify the configuration:

On SW1:

```
SW1#Show interfaces private-vlan mapping

Interface Secondary VLAN Type
--------- -------------- -----------------

vlan10    20             community
vlan10    30             community
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To test the configuration:

On R2:

```
R2#Ping 203.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 203.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

On R7:

```
R7#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

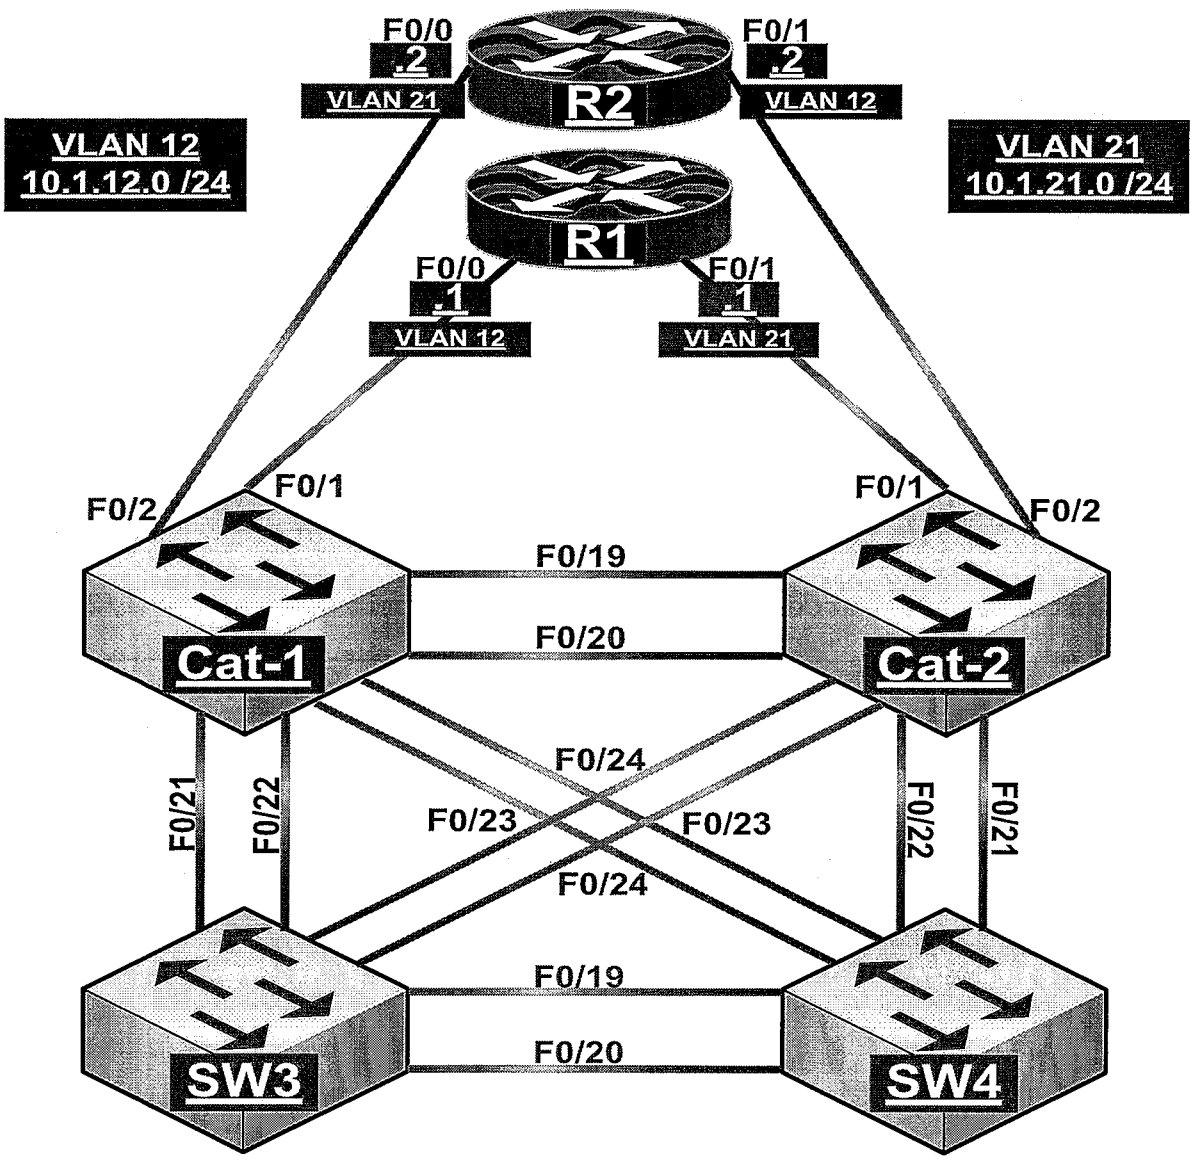
Task 8

Erase the startup config and reload the routers before proceeding to the next task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 12
Basic Spanning-tree Configuration



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

- Name the first switch CAT-1 and shut down interfaces F0/3 through F0/18
- Name the second switch CAT-2 and shut down interfaces F0/3 through F0/18
- Name the third switch SW-3 and shut down interfaces F0/1 through F0/18
- Name the fourth switch SW-4 and shut down interfaces F0/1 through F0/18
- Create VLANs 12 and 21 on all four switches
- Place R1 F0/0 and R2 F0/1 interfaces in VLAN 12
- Place R1 F0/1 and R2 F0/0 interfaces in VLAN 21
- All switch-to-switch links should be 802.1Q trunks.
- Permit only VLANs 12 and 21 on the switch-to-switch links
- Assign IP and MAC addresses as shown in the table below
- Verify that R1 can ping R2 on both subnets

You can also use the initial configuration file from the "Boot-Camp-Init" directory:

IP & MAC Addressing:

Router	Interface	IP address	Mac-address
R1	F0/0	10.1.12.1 /24	0000.**1111.0000** → Identifies **R1**'s **F0/0**
	F0/1	10.1.21.1 /24	0000.**1111.1111** → Identifies **R1**'s **F0/1**
R2	F0/0	10.1.21.2 /24	0000.**2222.0000** → Identifies **R2**'s **F0/0**
	F0/1	10.1.12.2 /24	0000.**2222.1111** → Identifies **R2**'s **F0/1**

Task 1

Make sure that traffic on VLAN 12 beetween R1 and R2 uses only the following switch-to-switch links.

- F0/21 between CAT-1 and SW-3
- F0/23 between SW-3 and CAT-2

Use only the command "**spanning-tree vlan** *vlan-id* **priority** *priority*," use it only on one switch, and make the minimum possible change from the default value.

The spanning-tree algorithm creates loop-free paths connecting all active switched ports. For each broadcast domain it chooses one root bridge, each non-root bridge then chooses one root port, and each link chooses one designated port. Ports not chosen as either root or designated ports go into the blocking

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

state. Task 1 requires you to affect the path that traffic will take by changing only the bridge priority on one switch, by setting which switch will be the root bridge.

Determine the existing topology

On CAT-1:

```
CAT-1#Show spanning-tree vlan 12

VLAN0012
  Spanning tree enabled protocol ieee
  Root ID     Priority    32780
              Address     000c.ce44.7d80
              Cost        19
              Port        25 (FastEthernet0/23)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32780  (priority 32768 sys-id-ext 12)
              Address     001a.a174.1f00
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
Fa0/21           Altn BLK 19        128.23   P2p
Fa0/22           Altn BLK 19        128.24   P2p
Fa0/23           Root FWD 19        128.25   P2p
Fa0/24           Altn BLK 19        128.26   P2p
```

The required path uses port F0/21 on CAT-1, but you can see that in this topology that port is in blocking state. F0/23 is the root port, the one closest to the root bridge, and it is forwarding. The root bridge has a cost of 19 out F0/23 from CAT-1, so the root bridge in this topology is SW-4.

SW-4 is chosen as the root bridge because it has the lowest bridge ID of all four switches in VLAN 12. The original 802.1D protocol specification defined the bridge ID as a 16 bits priority field and a 48 bits MAC address. The default priority value is the midpoint of the range, 32768 decimal, or 8000 hexadecimal. So by default, the root bridge will be the one with the lowest built-in MAC address. You can see this MAC address by doing a "Show version | Inc Base" or by noting the output of "Show spanning-tree".

In the above output you see the MAC address for the root bridge and it's 000C.CE44.7D80 and the MAC address of CAT-1 is 001A.A174.1F00. Note the root bridge MAC address is lower in the third hexadecimal place from the left. These results will be different on your pod, and you may have a different

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

root bridge at this point in the lab.

To achieve the required path between R1 and R2, you need to make SW-3 the root bridge for VLAN 12. This will cause CAT-1 to choose F0/21 as its root port and move it to forwarding state (It has a lower port-ID). Similarly, CAT-2 will change its root port from F0/21 to the required F0/23. You cannot change the MAC address portion of the bridge ID,

Therefore, to change the root bridge you change the bridge priority portion of the bridge ID for that VLAN. Let's change the priority to 100:

To Lower the priority for VLAN 12 on SW-3:

On SW-3

```
SW-3(config)#spanning-tree vlan 12 priority 100

% Bridge Priority must be in increments of 4096.
% Allowed values are:
  0     4096  8192  12288 16384 20480 24576 28672
  32768 36864 40960 45056 49152 53248 57344 61440
```

Catalyst 3550 and 3560 switches implement the extended system ID feature, so you can only set bridge priority to one of 16 different values, shown above. This feature divides the original bridge priority field into a 4-bit priority and 12-bit system ID. The system ID codes the VLAN ID. It allows for 4096 values (the number of VLANs permitted in the 802.1Q trunking protocol). In the output below you see that the bridge priority field for SW-3 has a value of 32780, which is the sum of the default priority and the extended system ID fields.

On SW-3:

```
SW-3#Show spanning-tree vlan 12 bridge

                                                    Hello  Max  Fwd
Vlan                      Bridge ID                  Time  Age  Dly
---------------- ---------------------------------- -----  ---  ---
VLAN0012         32780 (32768,  12) 000d.2870.fe00    2    20   15
```

Lower SW-3 VLAN 12 bridge priority to 28672

On SW-3:

```
SW-3(config)#spanning-tree vlan 12 priority 28672
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

```
CAT-1#Show spanning-tree vlan 12

VLAN0012
  Spanning tree enabled protocol ieee
  Root ID     Priority    28684
              Address     000d.2870.fe00
              Cost        19
              Port        23 (FastEthernet0/21)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32780  (priority 32768 sys-id-ext 12)
              Address     001a.a174.1f00

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

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
Fa0/21           Root FWD 19        128.23   P2p
Fa0/22           Altn BLK 19        128.24   P2p
Fa0/23           Altn BLK 19        128.25   P2p
Fa0/24           Altn BLK 19        128.26   P2p
```

CAT-1 shows the root bridge with priority 28684, which is the sum of the configured bridge priority on SW-3, 28672, and the VLAN ID, 12. The root bridge address is the base MAC address on SW-3. Only interfaces F0/1 and F0/21 are forwarding for VLAN 12, so traffic arriving from R1 on F0/1 will be forwarded out F0/21, as required.

On CAT-2:

```
CAT-2#Show spanning-tree vlan 12


VLAN0012
  Spanning tree enabled protocol ieee
  Root ID     Priority    28684
              Address     000d.2870.fe00
              Cost        19
              Port        25 (FastEthernet0/23)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32780  (priority 32768 sys-id-ext 12)
              Address     0019.e780.6d80
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/2            Desg FWD 19        128.4    P2p
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Altn BLK 19        128.23   P2p
Fa0/22           Altn BLK 19        128.24   P2p
Fa0/23           Root FWD 19        128.25   P2p
Fa0/24           Altn BLK 19        128.26   P2p
```

Similarly on CAT-2, the root bridge is shown with the priority and address from SW-3. Traffic arriving on root port F0/23 can be forwarded toward R2 on F0/2, as required by the task. For these switches, ports F0/19 and F0/20 are also forwarding, but recall that these links are blocked on the CAT-1 side, so traffic cannot flow between the routers on these links. You will come back to this issue in a later task.

On SW-3:

```
SW-3#Show spanning-tree vlan 12

VLAN0012
  Spanning tree enabled protocol ieee
  Root ID     Priority    28684
              Address     000d.2870.fe00
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    28684  (priority 28672 sys-id-ext 12)
              Address     000d.2870.fe00
              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        128.19   P2p
Fa0/20           Desg FWD 19        128.20   P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/22           Desg FWD 19        128.22   P2p
Fa0/23           Desg FWD 19        128.23   P2p
Fa0/24           Desg FWD 19        128.24   P2p
```

Because it is the root bridge, all of the ports on SW-3 are in forwarding state. You will find a blocked port on a root bridge only when it is looped to itself.

On SW-3:

If you do not get any output, you should generate a ping from the hosts within VLAN 12, example: On R1, ping 10.1.12.2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW-3#Show mac-address-table dynamic vlan 12

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

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  12    0000.1111.0000    DYNAMIC     Fa0/21 → R1's F0/0
  12    0000.2222.1111    DYNAMIC     Fa0/23 → R2's F0/1
Total Mac Addresses for this criterion: 2
```

Using "show interface" on the routers and "show mac-address-table" on the switches you can further verify that traffic is taking the required path. In the output of the above "show" command you see traffic to MAC address of 0000.1111.0000 is sent out of the F0/21 interface, whereas, the traffic to MAC address 0000.2222.1111 is sent out SW-3's F0/23 interface toward R2's F0/1.

Task 2

Raise the spanning-tree cost on a port in VLAN 12 so that CAT-1 chooses F0/22 as its root port instead of F0/21.

A non-root bridge chooses one root port based on the following criteria:

- **Choose the port with the lowest root path cost**
- **If root path costs are equal, choose the port connected to the upstream neighbor with the lowest bridge ID**
- **If the bridge IDs of the upstream neighbors are the same (parallel links between two switches) choose the port that receives the lowest port ID.**

The parallel links connecting CAT-1 and SW-3 have an equal cost. You are asked to raise the cost on a port for VLAN 12 so that CAT-1 chooses F0/22 as its root port.

You will need to raise the cost of the F0/21 link, but which port on this link needs to be adjusted, the root port on CAT-1 or the designated port on SW-3?

Root path cost is incremented as BPDUs arrive on designated ports. Changing the spanning-tree cost on root the bridge has no effect.

On CAT-1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On CAT-1, enable debugging and raise the cost on F0/21:

```
CAT-1#Debug spanning-tree events
Spanning Tree event debugging is on

CAT-1(config)#interface f0/21
CAT-1(config-if)#spanning-tree vlan 12 cost 20

23:57:18: STP: VLAN0012 new root port Fa0/22, cost 19
23:57:18: STP: VLAN0012 Fa0/22 -> listening
23:57:18: STP: VLAN0012 sent Topology Change Notice on Fa0/22
23:57:18: STP: VLAN0012 Fa0/21 -> blocking
23:57:33: STP: VLAN0012 Fa0/22 -> learning
23:57:48: STP: VLAN0012 sent Topology Change Notice on Fa0/22
23:57:48: STP: VLAN0012 Fa0/22 -> forwarding
```

As soon as the cost of F0/21 is raised above the default value of 19 for a 100 Mbps link, CAT-1 chooses F0/22 as its root port and the port transitions from blocking state to listening state. The old root port, F0/21 immediately goes into blocking state.

F0/22 spends 15 seconds in listening state (the default value of the forward delay timer) and then transitions to learning state. While in learning state F0/22 spends another 15 seconds building its MAC address table, then transitions to forwarding state. Traffic between R1 and R2 would have been disrupted for a total of 30 seconds, twice the forward delay timer.

To verify the configuration:

On CAT-1:

```
CAT-1#Show spanning-tree vlan 12 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
Fa0/21           Altn BLK 20        128.23   P2p
Fa0/22           Root FWD 19        128.24   P2p
Fa0/23           Altn BLK 19        128.25   P2p
Fa0/24           Altn BLK 19        128.26   P2p
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 3

Links F0/19 and F0/20 connecting CAT-1 and CAT-2 are in blocking state on one side for VLAN 12. Use the command "**spanning-tree** *vlan-id* **priority** *priority*" to move the VLAN 12 designated ports on these links from one switch to the other. Do not change the root bridge for VLAN 12.

For VLAN 12 there is a physical loop connecting CAT-1, SW-3 and CAT-2. The spanning-tree algorithm broke this loop at OSI Layer 2 by placing ports in STP blocking state. Because all of the forwarding link costs are equal (19), the blocked ports will be on the parallel links connecting CAT-1 and CAT-2. But which side will block and which side will have the designated port? With equal path costs to the root, the algorithm chooses the port on the link that is connected to the switch that has the lowest bridge ID. As you saw earlier, this will be the switch with the lowest MAC address by default, and this is purely a matter of chance.

Determine which switch has the designated ports

On CAT-1:

```
CAT-1#Show spanning-tree vlan 12 bridge

                                                    Hello  Max  Fwd
Vlan                      Bridge ID                  Time  Age  Dly
---------------- ---------------------------------- -----  ---  ---
VLAN0012         32780 (32768,  12) 001a.a174.1f00    2    20   15

CAT-1#Show spanning-tree vlan 12 interface f0/19 detail

 Port 21 (FastEthernet0/19) of VLAN0012 is blocking
   Port path cost 19, Port priority 128, Port Identifier 128.21.
   Designated root has priority 28684, address 000d.2870.fe00
   Designated bridge has priority 32780, address 0019.e780.6d80
   Designated port id is 128.21, designated path cost 19
   Timers: message age 3, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 27, received 41329
```

On CAT-2:

```
CAT-2#Show spanning-tree vlan 12 bridge
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Vlan                Bridge ID                 Hello Time  Max Age  Fwd Dly
---------------- -------------------------------- -----  ---  ---
VLAN0012         32780 (32768,  12) 0019.e780.6d80    2    20   15

CAT-2#Show spanning-tree vlan 12 interface f0/19 detail

 Port 21 (FastEthernet0/19) of VLAN0012 is forwarding
   Port path cost 19, Port priority 128, Port Identifier 128.21.
   Designated root has priority 28684, address 000d.2870.fe00
   Designated bridge has priority 32780, address 0019.e780.6d80
   Designated port id is 128.21, designated path cost 19
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 41380, received 7
```

On our pod, CAT-2 happens to have the lowest base MAC address, and so the lowest default bridge ID. Link F0/19 is therefore forwarding on the CAT-2 side and blocking on the CAT-1 side.

Note that this result may be reversed on your pod, depending on your switch base MAC addresses.

The command "Show spanning-tree vlan 12 interface f0/19 detail" indicates the bridge ID of the switch that has the link's designated port. Note the BPDU counts: designated ports source BPDUs, root ports and blocking ports receive them. In our case, we need to move the blocked ports from CAT-1 to CAT-2.

Raise the bridge ID for VLAN 12 on CAT-2:

```
CAT-2#Debug spanning-tree events
Spanning Tree event debugging is on

CAT-2(config)#spanning-tree vlan 12 priority 100

% Bridge Priority must be in increments of 4096.
% Allowed values are:
  0     4096  8192  12288 16384 20480 24576 28672
  32768 36864 40960 45056 49152 53248 57344 61440

CAT-2(config)#spanning-tree vlan 12 priority 36864

1d01h: setting bridge id (which=1) prio 36876 prio cfg 36864 sysid 12
(on) id 900C.0019.e780.6d80
1d01h: STP: VLAN0012 sent Topology Change Notice on Fa0/23
1d01h: STP: VLAN0012 Fa0/19 -> blocking
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
1d01h: STP: VLAN0012 Fa0/20 -> blocking
```

Instead of raising the bridge priority of CAT-2, we might have been able to lower the bridge ID of CAT-1, but this might have made CAT-1 the root bridge for VLAN 12, rather than SW-3, as required. It would depend on the switch base MAC addresses.

To verify the configuration:

On CAT-1:

```
CAT-1#Show spanning-tree vlan 12
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/1	Desg	FWD	19	128.3	P2p
Fa0/19	Desg	FWD	19	128.21	P2p
Fa0/20	Desg	FWD	19	128.22	P2p
Fa0/21	Altn	BLK	20	128.23	P2p
Fa0/22	Root	FWD	19	128.24	P2p
Fa0/23	Altn	BLK	19	128.25	P2p
Fa0/24	Altn	BLK	19	128.26	P2p

Note the output of the above "Show" command reveals that F0/19 and F0/20 are in forwarding state on CAT-1, whereas, the output of the following "Show" command reveals that ports F0/19 and F0/20 are in Blocking state on CAT-2:

On CAT-2:

```
CAT-2#Show spanning-tree vlan 12 | B Inter
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/2	Desg	FWD	19	128.4	P2p
Fa0/19	Altn	BLK	19	128.21	P2p
Fa0/20	Altn	BLK	19	128.22	P2p
Fa0/21	Altn	BLK	19	128.23	P2p
Fa0/22	Altn	BLK	19	128.24	P2p
Fa0/23	Root	FWD	19	128.25	P2p
Fa0/24	Altn	BLK	19	128.26	P2p

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

Make sure that traffic on VLAN 21 between R1 and R2 uses only the following switch-to-switch links.

- F0/21 between CAT-2 and SW-4
- F0/23 between SW-4 and CAT-1

Use only the command "**spanning-tree vlan** *vlan-id* **root**," and use it only on one switch. Set the appropriate diameter value.

Similar to Task 1, you must configure a root bridge so that traffic takes the required path. In this case you must set the bridge priority indirectly, by using the "root primary" macro. This macro determines the bridge priority of the current root bridge and sets the local bridge priority to a lower value. To achieve the required path you need to make SW-4 the root bridge for VLAN 21.

But what is diameter, and how should it be set? The root macro will calculate the optimum maximum delay and forwarding delay timer values, based on the size of the switched network: the smaller the network, the lower the timer values, and the faster the the topology will converge. For this calculation the diameter of the switched network is the maximum number of switches that could possibly be in the path between two end stations. In our case, the correct diameter is 4. The default and maximum value is 7.

To configure the root macro on SW-4 for VLAN 21:

```
SW-4(config)#spanning-tree vlan 21 root primary diameter 4
```

To verify the configuration:

On SW-4:

```
SW-4#Show spanning-tree vlan 21

VLAN0021
  Spanning tree enabled protocol ieee

  Root ID     Priority    24597
              Address     000c.ce44.7d80
              This bridge is the root
              Hello Time   2 sec  Max Age 14 sec  Forward Delay 10 sec

  Bridge ID   Priority    24597  (priority 24576 sys-id-ext 21)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
              Address     000c.ce44.7d80
              Hello Time   2 sec  Max Age 14 sec  Forward Delay 10 sec
              Aging Time 300
[output is omitted for brevity]
```

Note that bridge priority on SW-4 was lowered to 24576, making SW-4 the root bridge, the Max Age timer has been reduced from the default 20 seconds to 14 seconds, and the Forward Delay timer has been reduced from the default 15 seconds to 10 seconds. However, you will not see the "root" command in the running configuration of the switch. Instead, you will see the actual command it entered:

On SW-4:

```
SW-4#Show run | include vlan 21

spanning-tree vlan 21 priority 24576
spanning-tree vlan 21 forward-time 10
spanning-tree vlan 21 max-age 14
```

Here is the output on CAT-1:

On CAT-1:

```
CAT-1#Show spanning-tree vlan 21

VLAN0021
  Spanning tree enabled protocol ieee


  Root ID     Priority    24597
              Address     000c.ce44.7d80
              Cost        19
              Port        25 (FastEthernet0/23)
              Hello Time   2 sec  Max Age 14 sec  Forward Delay 10 sec

  Bridge ID   Priority    32789  (priority 32768 sys-id-ext 21)
              Address     001a.a174.1f00
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300
```

Note: CAT-1 has learned the Max Age and Forward Delay timers configured on SW-4. These values are advertised in configuration BPDUs sourced by the root bridge and are forwarded from switch to switch. It is these timers that all switches on VLAN 21 will actually use. The locally configured timers are only used if that switch happens to be the root bridge.

To verify the required path between R1 and R2 for VLAN 21, you could issue the commands "Show

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

interface" and "Show mac-address-table dynamic VLAN 21" as you did in Task 1 for VLAN 12.

Task 5

Raise the spanning-tree port ID on a port in VLAN 21 so that CAT-2 chooses F0/22 as its root port instead of F0/21.

Here is the spanning tree topology for VLAN 21 on CAT-2 for our pod. Note that CAT-2 chose F0/21 as its root port, rather than F0/22. Both ports have a local cost of 19, and BPDUs received on each port report the same value for sender bridge ID. The ports differ in their local port IDs; the port ID for F0/21 is 128.23, and the local port ID for F0/22 is 128.24.

```
CAT-2#show span vlan 21 | B Inter
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/1	Desg	FWD	19	128.3	P2p
Fa0/19	Desg	FWD	19	128.21	P2p
Fa0/20	Desg	FWD	19	128.22	P2p
Fa0/21	**Root**	**FWD**	**19**	**128.23**	**P2p**
Fa0/22	**Altn**	**BLK**	**19**	**128.24**	**P2p**
Fa0/23	Altn	BLK	19	128.25	P2p
Fa0/24	Altn	BLK	19	128.26	P2p

The port ID is sent in BPDUs sourced by that interface. It is a 16-bit field; originally the first 8 bits were priority and the last 8 bits were an interface number. This limited the number of ports on a switch to 255. To increase the number of potential ports on a switch this field is now divided so that only the first 4 bits can be used for priority, and the last 12 bits can be used for interface number. The priority value is still evaluated as if it was a full 8-bit number with the four lowest bits set to zero. Here you see the resulting configuration options:

```
CAT-2(config-if)#spanning-tree vlan 21 port-priority ?

  <0-240>  port priority in increments of 16
```

To change CAT-2's root port from F0/21 to f0/22 you will need to raise the port priority on the F0/21 link for VLAN 21, but should you do this on CAT-2 or SW-4?

Raise the port priority for f0/21 on SW-4

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW-4(config-if)#spanning-tree vlan 21 port-priority 144
```

All else being equal, non-root bridges choose their root port based on RECEIVED port IDs, not locally configured port IDs. Change the port ID on the upstream switch, that is, on the designated port for the link. Note this is the opposite of the cost configuration, where local, downsteam cost is effective, as you saw in Task 2.

To verify the configuration:

On SW-4:

```
SW-4#Show spanning-tree vlan 21 int f0/21

Vlan                Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
VLAN0021            Desg FWD 19        144.21   P2p
```

On Cat-2:

```
CAT-2#Show spanning-tree vlan 21 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Altn BLK 19        128.23   P2p
Fa0/22           Root FWD 19        128.24   P2p
Fa0/23           Altn BLK 19        128.25   P2p
Fa0/24           Altn BLK 19        128.26   P2p
```

By the way, notice that CAT-2, a 3560, shows F0/21 as link number 23, but SW-4, a 3550, shows F0/21 as link number 21. This difference is irrelevant to the spanning-tree algorithm, because port IDs are only compared when the sending bridge IDs are equal. In other words, spanning-tree never compares a port ID on one switch to a port ID on another switch.

When dealing with parallel links, you change local cost to change the local root port, and you change the local port ID to influence a neighbor's choice of root port.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

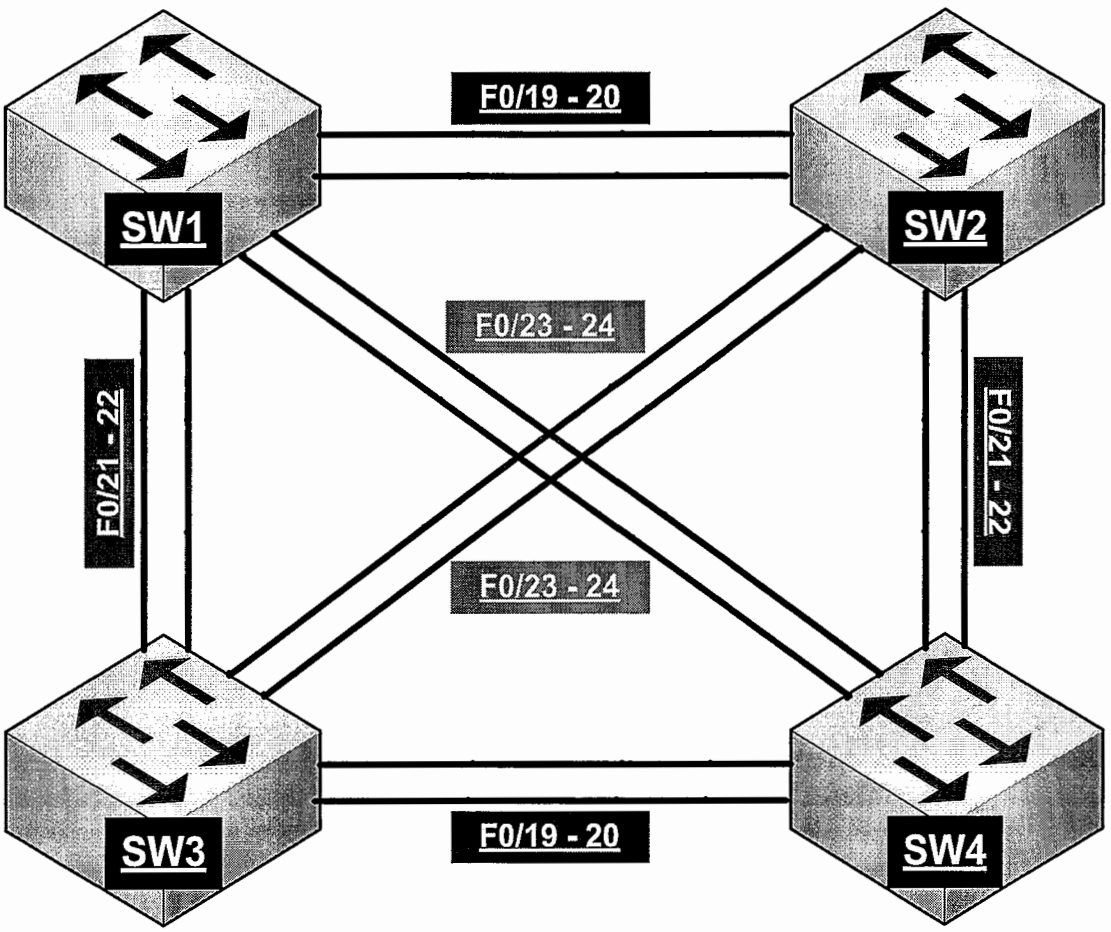
Task 6

Erase the "config.text" and "vlan.dat" on all four switches and reload before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 13
Advanced Spanning-tree Configuration



Task 1

Shut down all ports on all switches. And configure the hostnames of these switches based on the above diagram.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On All Switches:

```
SWx(config)#Interface range F0/1 - 24
SWx(config-if-range)#Shutdown
```

To verify the configuration:

On All Switch:

```
SWx#Show interface Status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Fa0/1		disabled	1	auto	auto	10/100BaseTX
Fa0/2		disabled	1	auto	auto	10/100BaseTX
Fa0/3		disabled	1	auto	auto	10/100BaseTX
Fa0/4		disabled	1	auto	auto	10/100BaseTX
Fa0/5		disabled	1	auto	auto	10/100BaseTX
Fa0/6		disabled	1	auto	auto	10/100BaseTX
Fa0/7		disabled	1	auto	auto	10/100BaseTX
Fa0/8		disabled	1	auto	auto	10/100BaseTX
Fa0/9		disabled	1	auto	auto	10/100BaseTX
Fa0/10		disabled	1	auto	auto	10/100BaseTX
Fa0/11		disabled	1	auto	auto	10/100BaseTX
Fa0/12		disabled	1	auto	auto	10/100BaseTX
Fa0/13		disabled	1	auto	auto	10/100BaseTX
Fa0/14		disabled	1	auto	auto	10/100BaseTX
Fa0/15		disabled	1	auto	auto	10/100BaseTX
Fa0/16		disabled	1	auto	auto	10/100BaseTX
Fa0/17		disabled	1	auto	auto	10/100BaseTX
Fa0/18		disabled	1	auto	auto	10/100BaseTX
Fa0/19		disabled	1	auto	auto	10/100BaseTX
Fa0/20		disabled	1	auto	auto	10/100BaseTX
Fa0/21		disabled	1	auto	auto	10/100BaseTX
Fa0/22		disabled	1	auto	auto	10/100BaseTX
Fa0/23		disabled	1	auto	auto	10/100BaseTX
Fa0/24		disabled	1	auto	auto	10/100BaseTX

On SW1:

```
SW1(config)#Hostname SW1
```

On SW2:

```
SW2(config)#Hostname SW2
```

On SW3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3(config)#Hostname SW3
```

On SW4:

```
SW4(config)#Hostname SW4
```

Task 2

Configure ports F0/19-20 connecting SW1 to SW2, and SW3 to SW4 as trunk ports; you should use an industry standard protocol to accomplish this task. These ports should never become an access port through negotiation.

On SW1, SW2, SW3 and SW4:

```
SWx(config)#int range f0/19-20
SWx(config-if-range)#Switch trunk encap dot1q
SWx(config-if-range)#Switch mode trunk
SWx(config-if-range)#No Shut
```

To verify the configuration:

On SW1:

```
SW1#Show inter trunk | Inc trunking

Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
```

On SW2:

```
SW2#Show inter trunk | Inc trunking

Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
```

On SW3:

```
SW3#Show inter trunk | Inc trunking

Fa0/19      on              802.1q          trunking      1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/20      on              802.1q          trunking      1
```

On SW4:

```
SW4#Show inter trunk | Inc trunking

Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
```

Task 3

Configure ports F0/21-22 connecting SW2 to SW4, and SW1 to SW3 as trunk ports; you should use an industry standard protocol to accomplish this task. These ports should never become an access port through negotiation.

On SW1, SW2, SW3 and SW4:

```
SWx(config)#int range f0/21-22
SWx(config-if-range)#Switch trunk encap dot1q
SWx(config-if-range)#Switch mode trunk
SWx(config-if-range)#No Shut
```

To verify the configuration:

On SW1:

```
SW1#Show int trunk | Inc trunking

Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
Fa0/21      on              802.1q          trunking      1
Fa0/22      on              802.1q          trunking      1
```

On SW2:

```
SW2#Show int trunk | Inc trunking

Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/21      on                802.1q          trunking      1
Fa0/22      on                802.1q          trunking      1
```

On SW3

```
SW3#Show int trunk | Inc trunking
```

```
Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
Fa0/21      on              802.1q          trunking      1
Fa0/22      on              802.1q          trunking      1
```

On SW4

```
SW4#Show int trunk | Inc trunking
```

```
Fa0/19      on              802.1q          trunking      1
Fa0/20      on              802.1q          trunking      1
Fa0/21      on              802.1q          trunking      1
Fa0/22      on              802.1q          trunking      1
```

Task 4

These switches should be configured in a VTP domain called "CCIE".

On SW1:

```
SW1(config)#Vtp domain CCIE
```

This configuration will be propagated via VTP to the other switches.

To verify the configuration:

On All Switches:

```
SWx#Show vtp status | inc VTP Domain Name
```

```
VTP Domain Name                 : CCIE
```

NOTE: the domain name is propagated by VTP messages when the trunk is established between the

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

switches.

Task 5

Create the following VLANs and ensure that they are propagated to all four switches: **100, 200, 300, and 400.**

On SW1:

```
SW1(config)#vlan 100,200,300,400
SW1(config-vlan)#exit
```

You must "exit" for the VLANs to be propagated.

To verify the configuration:

On SW1:

```
SW1#Show vlan Brief | Inc VLAN

VLAN Name                                Status    Ports
100  VLAN0100                            active
200  VLAN0200                            active
300  VLAN0300                            active
400  VLAN0400                            active
```

On SW2:

```
SW2#Show vlan Brief | Inc VLAN

VLAN Name                                Status    Ports
100  VLAN0100                            active
200  VLAN0200                            active
300  VLAN0300                            active
400  VLAN0400                            active
```

On SW3:

```
SW3#Show vlan Brief | Inc VLAN

VLAN Name                                Status    Ports
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
100  VLAN0100                         active
200  VLAN0200                         active
300  VLAN0300                         active
400  VLAN0400                         active
```

On SW4:

```
SW4#Show vlan Brief | Inc VLAN

VLAN Name                             Status    Ports
100  VLAN0100                         active
200  VLAN0200                         active
300  VLAN0300                         active
400  VLAN0400                         active
```

Task 6

Ensure that SW1 is the root bridge for VLAN 100, SW2 is the root bridge for VLAN 200, SW3 is the root bridge for VLAN 300 and SW4 is the root bridge for VLAN 400. You should use a macro to accomplish this task.

On SW1:

```
SW1(config)#spanning-tree vlan 100 root primary
```

On SW2:

```
SW2(config)#Spanning-tree vlan 200 root primary
```

On SW3:

```
SW3(config)#Spanning-tree vlan 300 root primary
```

On SW4:

```
SW4(config)#Spanning-tree vlan 400 root primary
```

To verify the configuration:

On SW1:

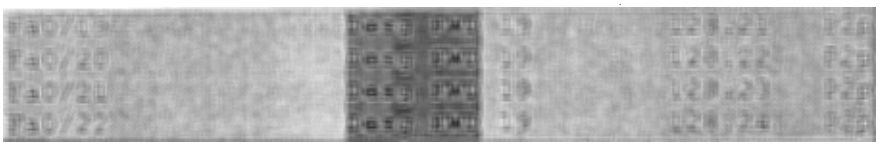
© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1#Sh spanning-tree VLAN 100

VLAN0100
  Spanning tree enabled protocol ieee
  Root ID     Priority    24676
              Address     0012.7f40.9380
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24676  (priority 24576 sys-id-ext 100)
              Address     0012.7f40.9380
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 15

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------
```



Note: Since SW1 is the root bridge for VLAN 100, all the trunk interfaces are Designated ports.

```
SW1#Sh ver | Inc Base

Base ethernet MAC Address       : 00:12:7F:40:93:80
```

On SW2:

```
SW2#Sh spanning-tree VLAN 200

VLAN0200
  Spanning tree enabled protocol ieee
  Root ID     Priority    24776
              Address     001d.e5d6.0000
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24776  (priority 24576 sys-id-ext 200)
              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
Fa0/20              Desg FWD 19        128.22   P2p
Fa0/21              Desg FWD 19        128.23   P2p
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/22              Desg FWD 19        128.24   P2p

SW2#Show ver | Inc Base

Base ethernet MAC Address       : 00:1D:E5:D6:00:00
```

On SW3:

```
SW3#Sh spanning-tree VLAN 300

VLAN0300
  Spanning tree enabled protocol ieee
  Root ID     Priority    24876
              Address     000c.858b.7a00
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24876  (priority 24576 sys-id-ext 300)
              Address     000c.858b.7a00
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 15

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/19           Desg FWD 19        128.19   P2p
Fa0/20           Desg FWD 19        128.20   P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/22           Desg FWD 19        128.22   P2p

SW3#Show ver | Inc Base

Base ethernet MAC Address: 00:0C:85:8B:7A:00
```

On SW4:

```
SW4#Sh spanning-tree VLAN 400

VLAN0400
  Spanning tree enabled protocol ieee
  Root ID     Priority    24976
              Address     000c.302d.9980
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24976  (priority 24576 sys-id-ext 400)
              Address     000c.302d.9980
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
            Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/19           Desg FWD 19        128.19   P2p
Fa0/20           Desg FWD 19        128.20   P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/22           Desg FWD 19        128.22   P2p

SW4#Sh ver | Inc Base

Base ethernet MAC Address: 00:0C:30:2D:99:80
```

Task 7

Implement the following policy:

1. VLAN 100 should never traverse SW4, this switch should NOT receive/send traffic for this VLAN.
2. VLAN 200 should never traverse SW3, this switch should NOT receive/send traffic for this VLAN.
3. VLAN 300 should never traverse SW2, this switch should NOT receive/send traffic for this VLAN.
4. VLAN 400 should never traverse SW1, this switch should NOT receive/send traffic for this VLAN.

To configure the first Policy:

On SW2:

```
SW2(config)#int range F0/21-22
SW2(config-if-range)#switchport trunk allowed vlan except 100
```

On SW3:

```
SW3(config)#int range F0/19-20
SW3(config-if-range)#switchport trunk allowed vlan except 100
```

On SW4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW4(config)#int range F0/19-20 , F0/21-22
SW4(config-if-range)#switchport trunk allowed vlan except 100
```

To verify the configuration:

On SW2:

```
SW2#Show inter trunk | Exc trunking|200|300|domain|pruned|Enc

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/20      1-4094
Fa0/21      1-99,101-4094
Fa0/22      1-99,101-4094
```

← Note VLAN 100 is NOT allowed on the Trunk links to SW4

On SW3:

```
SW3#Show inte trunk | exc trunking|200|300|domain|pruned|Enc

Port        Vlans allowed on trunk
Fa0/19      1-99,101-4094
Fa0/20      1-99,101-4094
Fa0/21      1-4094
Fa0/22      1-4094
```

← Note VLAN 100 is NOT allowed on the Trunk links to SW4

On SW4:

```
SW4#Show int trunk | exc trunking|200|300|pruned|domain|Enc

Port        Vlans allowed on trunk
Fa0/19      1-99,101-4094
Fa0/20      1-99,101-4094
Fa0/21      1-99,101-4094
Fa0/22      1-99,101-4094
```

← Note Traffic for VLAN 100 will NOT be sent out of the trunk links

To configure the second policy:

Ensure that the first policy in this task is NOT overridden by the next policy:

On SW1:

```
SW1(config)#Int range f0/21-22
SW1(config-if-range)#Switch trunk allowed vlan except 200
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW4:

```
SW4(config)#Int range f0/19-20
SW4(config-if-range)#Switch trunk allowed vlan except 100,200
```

On SW3:

```
SW3(config)#Int range f0/19-20
SW3(config-if-range)#Switchport trunk allowed vlan except 100,200

SW3(config)#Int range f0/21-22
SW3(config-if-range)#Switchport trunk allowed vlan except 200
```

To verify the configuration:

On SW1:

```
SW1#Show int trunk | exc trunking|200|300|100|pruned|domain|Enc

Port        Vlans allowed on trunk
Fa0/19      1-4094
Fa0/20      1-4094
Fa0/21      1-199,201-4094
Fa0/22      1-199,201-4094
```

Note VLAN 200 is NOT allowed on the Trunk links to SW3

On SW4:

```
SW4#Show int trunk | exc trunking|200|300|400|pruned|domain|Enc

Port        Vlans allowed on trunk
Fa0/19      1-99,101-199,201-4094
Fa0/20      1-99,101-199,201-4094
Fa0/21      1-99,101-4094
Fa0/22      1-99,101-4094
```

Note the traffic for VLAN 100 or 200 is NOT allowed on the trunk links to SW3

The traffic for VLAN 100 is NOT allowed on these Trunk links

On SW3:

```
SW3#Show inter trunk | exc trunking|200|300|pruned|domain|Enc

Port        Vlans allowed on trunk
Fa0/19      1-99,101-199,201-4094
Fa0/20      1-99,101-199,201-4094
Fa0/21      1-199,201-4094
Fa0/22      1-199,201-4094
```

The traffic from VLAN 100 and 200 is NOT allowed on the trunk links

The traffic from VLAN 200 is NOT allowed on the trunk links

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To configure the third Policy:

On SW2:

```
SW2(config-if-range)#int range f0/21-22
SW2(config-if-range)#swi trunk allowed vlan except 100,300

SW2(config-if-range)#int range f0/19-20
SW2(config-if-range)#swi trunk allowed vlan except 300
```

On SW4:

```
SW4(config)#int range f0/21-22
SW4(config-if-range)#swi trunk allowed vlan except 100,300
```

On SW1:

```
SW1(config)#int range f0/19-20
SW1(config-if-range)#Switch trunk allowed vlan except 300
```

To verify the configuration:

On SW2:

```
SW2#Show inter trunk | exc trunking|200|400|Enc|pruned|domain
```

Port	Vlans allowed on trunk	
Fa0/19	1-299,301-4094	← The traffic for VLAN 300 is NOT allowed on these trunk links
Fa0/20	1-299,301-4094	
Fa0/21	1-99,101-299,301-4094	← The traffic for VLAN 100 or 300 is NOT allowed on these trunk links
Fa0/22	1-99,101-299,301-4094	

On SW4:

```
SW4#Show int trunk | exc trunking|200|300|400|pruned|domain|Enc
```

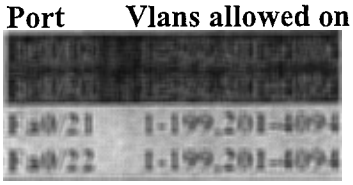
Port	Vlans allowed on trunk	
[redacted]		← The traffic for VLAN 100 or 200 are NOT allowed on these trunk links
[redacted]		← The traffic for VLANs 100 or 300 are NOT allwed on these trunk links

On SW1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1#Show int trunk | exc trunking|200|300|100|pruned|domain|Enc
```

Port Vlans allowed on trunk



← The traffic for VLAN 300 is NOT allowed on this trunk link

← The traffic for VLAN 200 is NOT allowed on this trunk link

To configure the forth Policy:

On SW1:

```
SW1(config)#int range F0/19-20
SW1(config-if-range)#swi trunk allowed vlan except 300,400

SW1(config)#int range f0/21-22
SW1(config-if-range)#swi trunk allowed vlan except 200,400
```

On SW2:

```
SW2(config)#int range f0/19-20
SW2(config-if-range)#swi trunk allowed vlan except 300,400
```

On SW3:

```
SW3(config)#int range f0/21-22
SW3(config-if-range)#swi trunk allowed vlan except 200,400
```

To verify the configuration:

On SW1:

```
SW1#Show int trunk | exc trunking|200|300|100|pruned|domain|Enc
```

Port Vlans allowed on trunk



← The traffic for VLANs 300 or 400 are NOT allowed on these trunk links.

← The traffic for VLANs 200 or 400 are NOT allowed on these trunk lonks.

On SW2:

```
SW2#Show inter trunk | exc trunking|100|200|300|Enc|pruned|domain
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Vlans allowed on trunk
Fa0/19      1-299,301-399,401-4094
Fa0/20      1-299,301-399,401-4094
Fa0/21      1-99,101-299,301-4094
Fa0/22      1-99,101-299,301-4094
```

The traffic for VLANs 300 or 400 are NOT allowed on these trunk links.

The traffic for VLANs 100 or 300 are NOT allowed on these trunk links.

On SW3:

```
SW3#Show inter trunk | exc trunking|200|300|pruned|domain|Enc

Port        Vlans allowed on trunk
Fa0/19      1-99,101-199,201-4094
Fa0/20      1-99,101-199,201-4094
Fa0/21      1-199,201-399,401-4094
Fa0/22      1-199,201-399,401-4094
```

The traffic for VLANs 100 or 200 are NOT allowed on these trunk links.

The traffic for VLANs 200 or 400 are NOT allowed on these trunk links.

Task 8

Configure SW1 such that it's the root bridge for VLAN 500, if this switch goes down, SW2 should become the root bridge for this VLAN. DO NOT override any of the previous tasks to accomplish this one.

On SW1:

```
SW1(config)#Spanning-tree vlan 500 root primary

SW1#Debug spanning-tree even
Spanning Tree event debugging is on

SW1(config)#VLAN 500
SW1(config-vlan)#
```

Note when creating a VLAN in global config mode, the VLAN is NOT created & propagated unless we exit out of VLAN configuration mode.

```
SW1(config-vlan)#Exit
```

Once exited, you should see the following console messages:

```
setting bridge id (which=3) prio 25076 prio cfg 24576 sysid 500 (on) id
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
 61F4.0012.7f40.9380
set portid: VLAN0500 Fa0/19: new port id 8015
STP: VLAN0500 Fa0/19 -> listening
set portid: VLAN0500 Fa0/20: new port id 8016
STP: VLAN0500 Fa0/20 -> listening
set portid: VLAN0500 Fa0/21: new port id 8017
STP: VLAN0500 Fa0/21 -> listening

STP: VLAN0500 heard root 33268-000c.858b.7a00 on Fa0/21
STP: VLAN0500 heard root 33268-000c.858b.7a00 on Fa0/22
STP: VLAN0500 heard root 33268-001d.e5d6.0000 on Fa0/19
STP: VLAN0500 heard root 33268-001d.e5d6.0000 on Fa0/20

STP: VLAN0500 Fa0/19 -> learning

STP: VLAN0500 Fa0/20 -> learning
STP: VLAN0500 Fa0/21 -> learning
STP: VLAN0500 Fa0/22 -> learning

STP: VLAN0500 Fa0/19 -> forwarding
STP: VLAN0500 Fa0/20 -> forwarding
STP: VLAN0500 Fa0/21 -> forwarding
STP: VLAN0500 Fa0/22 -> forwarding
```

To disable the Debug:

On SW1:

```
SW1#u all
```

On SW2:

```
SW2(config)#Spanning-tree vlan 500 root secondary
```

To verify the configuration:

On SW2:

```
SW2#Sh spanning-tree vlan 500 | exc Interface|Fa0

VLAN0500
  Spanning tree enabled protocol ieee
  Root ID     Priority    25076
              Address     0012.7f40.9380
              Cost        19
```


© 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    29172  (priority 28672 sys-id-ext 500)
             Address     001d.e5d6.0000
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

------------------ ---- --- --------- -------- --------------------------------
```

Note: SW1 is the root

On SW1:

```
SW1#Show ver | Inc Base

Base ethernet MAC Address       : 00:12:7F:40:93:80
```

Task 9

Create VLAN 500 and ensure that the traffic from SW2 for VLAN 500 takes the following path:

SW2 → SW4 → SW3 → SW1

SW2 should use ports F0/19 or F0/20 ONLY if the path through SW4 → SW3 → SW1 is NOT possible due to a link failure.

On SW2:

```
SW2#Sh spanning-tree vlan 500

VLAN0500
  Spanning tree enabled protocol ieee
  Root ID     Priority    25076
              Address     0012.7f40.9380

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

  Bridge ID   Priority    29172  (priority 28672 sys-id-ext 500)
              Address     001d.e5d6.0000
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Fa0/19              Root FWD 19        128.21   P2p
Fa0/20              Altn BLK 19        128.22   P2p
Fa0/21              Desg FWD 19        128.23   P2p
Fa0/22              Desg FWD 19        128.24   P2p
```

Note: SW2 is taking port F0/19 to get to the root bridge for VLAN 500, and its root cost is 19 which is the cost of a 100 Mbps link. If these ports (F0/19 & F0/20) are shut down, you should see the cost of the local switch (SW2) to the root bridge; it should be 57, the following reveals the cumulative cost of the local switch to the root bridge for VLAN 500:

To Add the total cost:

The cost of the link from SW2 → SW4 which is 19 +
The cost of the link from SW4 → SW3 which is 19
The cost of the link from SW3 → SW1 which is also 19

The total cost should is 57

To reveal this information:

On SW2:

```
SW2(config)#int range f0/19-20
SW2(config-if-range)#Shut
```

To verify the new cost:

```
SW2#Sh spanning-tree vlan 500

VLAN0500
  Spanning tree enabled protocol ieee
  Root ID     Priority    25076
              Address     0012.7f40.9380
              Cost        57
              Port        23 (FastEthernet0/21)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    29172  (priority 28672 sys-id-ext 500)

              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
------------------- ---- --- --------- -------- --------------------------------
Fa0/21              Root FWD 19        128.23   P2p
Fa0/22              Altn BLK 19        128.24   P2p
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Note: it is taking port F0/21 toward the root bridge through SW4. If this is traced all the way back to SW1 you will see that the path from SW2's perspective is as follows:

SW2 → SW4 → SW3 → SW1.

To configure this task, you should "NO Shutdown" ports F0/19-20 first.

On SW2:

```
SW2(config)#int range f0/19-20
SW2(config-if-range)#No Shut

SW2(config-if-range)#Spanning-tree vlan 500 cost 58
```

Note: the above command sets the cost through ports F0/19-20 higher than 57; therefore, traffic from SW2 for VLAN 500 will have two paths:
Primary: SW2 → SW4 → SW3 → SW1
Backup : SW2 → SW1

To verify the configuration:

On SW2:

```
SW2#Show spanning-tree vlan 500

VLAN0500

  Spanning tree enabled protocol ieee
  Root ID     Priority    25076
              Address     0012.7f40.9300
              Cost        57
              Port        23 (FastEthernet0/21)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    29172  (priority 28672 sys-id-ext 500)
              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              Altn BLK 58        128.21   P2p
Fa0/20              Altn BLK 58        128.22   P2p
Fa0/21              Root FWD 19        128.23   P2p
Fa0/22              Altn BLK 19        128.24   P2p
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW2:

```
SW2(config)#int range f0/21-22
SW2(config-if-range)#Shut
```

You should wait for 30 seconds before performing the following show command:

```
SW2#Show spanning-tree vlan 500

VLAN0500
  Spanning tree enabled protocol ieee
  Root ID     Priority    25076
              Address     0012.7f40.9380
              Cost        58
              Port        21 (FastEthernet0/19)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    29172  (priority 28672 sys-id-ext 500)
              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              Root FWD 58        128.21   P2p
Fa0/20              Altn BLK 58        128.22   P2p
```

On SW2

```
SW2(config)#int range f0/21-22
SW2(config-if-range)#No Shut
```

Task 10

Create VLAN 600 and configure SW3 as the root bridge for this VLAN; this switch should be configured such that traffic for VLAN 600 uses the following ports:

- F0/22 from SW1
- F0/20 from SW4

On SW-3:

```
SW3(config)#VLAN 600
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3(config-vlan)#Exit

SW3(config)#Spanning-tree vlan 600 root primary
```

To see the existing path from SW1 and SW4:

On SW1:

```
SW1#Show Spanning-tree vlan 600

VLAN0600
  Spanning tree enabled protocol ieee
  Root ID     Priority    25176
              Address     000c.858b.7a00
              Cost        19
              Port        23 (FastEthernet0/21)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    33368  (priority 32768 sys-id-ext 600)
              Address     0012.7f40.9380
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 15

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Fa0/19              Desg FWD 19        128.21   P2p
Fa0/20              Desg FWD 19        128.22   P2p
Fa0/21              Root FWD 19        128.23   P2p
Fa0/22              Altn BLK 19        128.24   P2p
```

On SW4:

```
SW4#Show Spanning-tree vlan 600

VLAN0600
  Spanning tree enabled protocol ieee
  Root ID     Priority    25176
              Address     000c.858b.7a00
              Cost        19
              Port        19 (FastEthernet0/19)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    33368  (priority 32768 sys-id-ext 600)
              Address     000c.302d.9980
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/19           Root FWD 19        128.19   P2p
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/20           Altn BLK 19        128.20   P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/22           Desg FWD 19        128.22   P2p
```

To configure the task:

On SW3:

```
SW-3(config)#int range F0/20 , F0/22
SW-3(config-if)#Spanning-tree vlan 600 port-priority 0
```

To verify the configuration:

On SW1:

By default, the port-priority of all interfaces for all VLANs are set to 128, remember that the lower cost has more preference, therefore, if the port-priority of the ports connecting this switch to SW1 and SW4 (F0/20 and F0/22 for SW4 and SW1 respectively) is reduced, these ports will be the preferred ports for connecting to the root bridge, in this case SW4.

To verify the configuration:

On SW3:

```
SW3#Show spanning-tree vlan 600 | B Fa0

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------

Fa0/19           Desg FWD 19        128.19   P2p
Fa0/20           Desg FWD 19        0.20     P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/22           Desg FWD 19        0.22     P2p
```

On SW1

```
SW1#Show spanning-tree vlan 600 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ----------------

Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
Fa0/21           Altn BLK 19        128.23   P2p
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/22               Root FWD 19        128.24   P2p
```

On SW4:

```
SW4#Show spanning-tree vlan 600 | B Interface

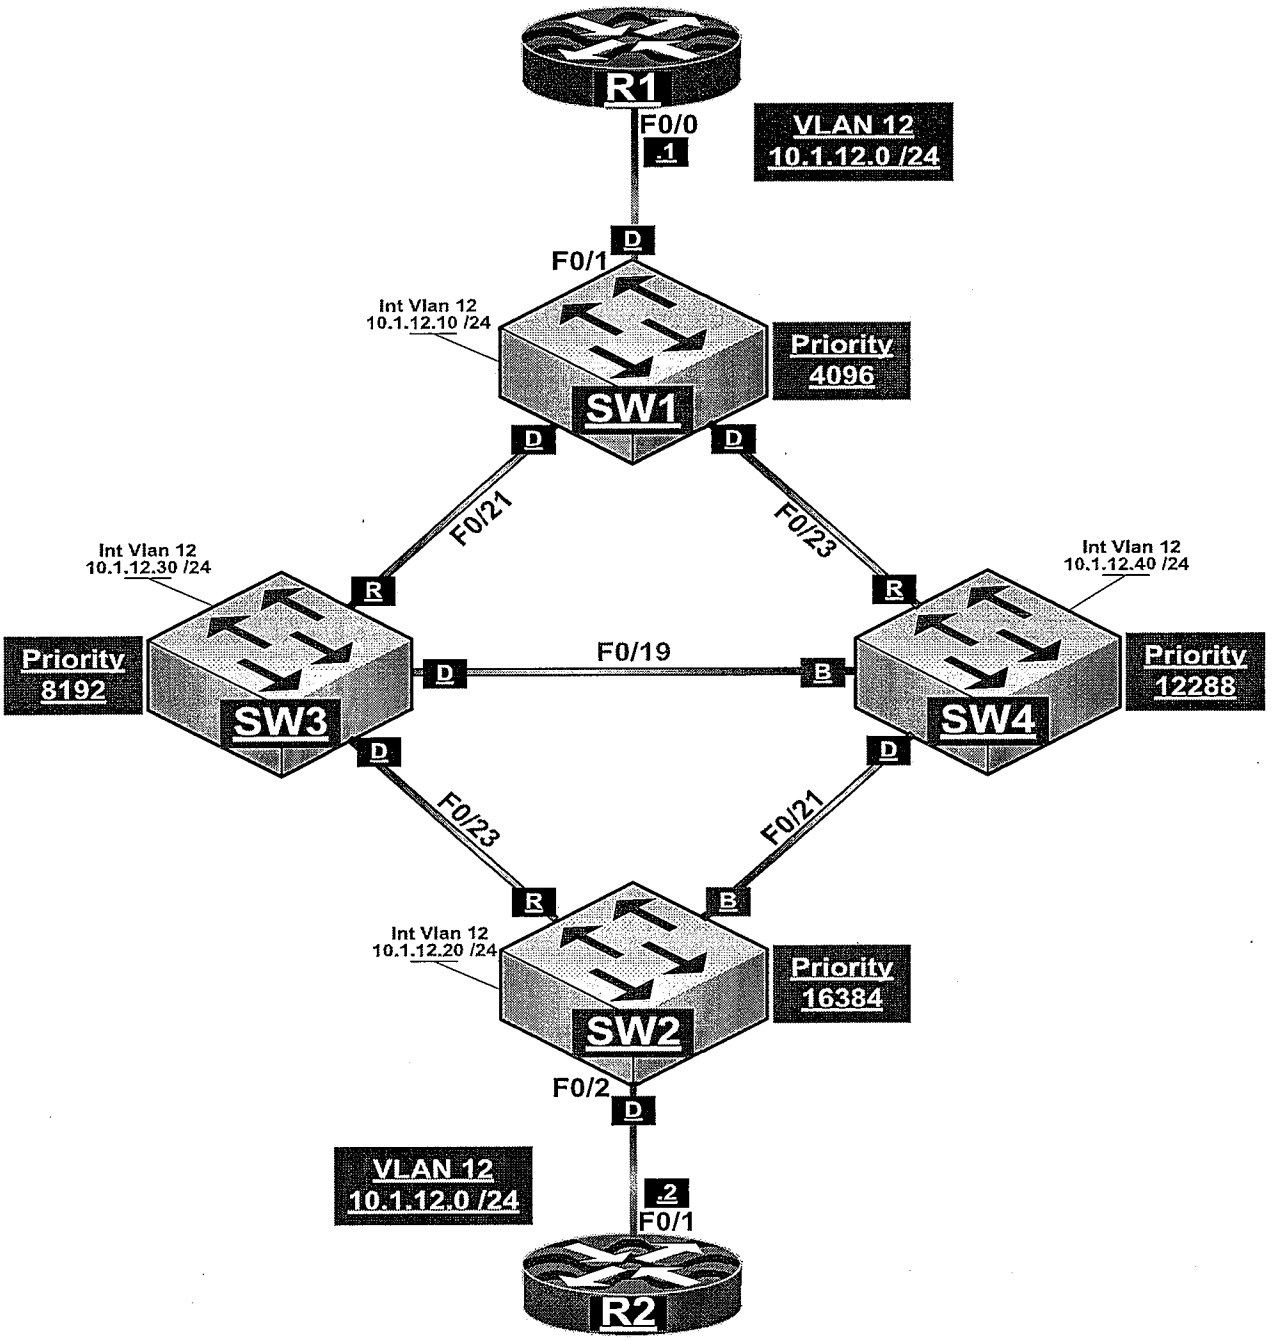
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------
Fa0/19           Altn BLK 19        128.19   P2p
Fa0/20           Root FWD 19        128.20   P2p
Fa0/21           Desg FWD 19        128.21   P2p
Fa0/22           Desg FWD 19        128.22   P2p
```

Task 11

Erase the "config.text" and "vlan.dat" and reload the switches before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 14
Rapid Spanning-Tree Protocol (802.1w)



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Setup:

- Shut down all interfaces that are not shown in the lab diagram, above
- Configure all of interfaces shown as access links in VLAN 12
- On each switch configure the bridge priority for VLAN 12 as shown in the diagram
- Switch ports should be root (R), designated (D) or blocking (B) as shown in the diagram
- Assign IP addresses as shown in the table below
- Enter "ntp server 10.1.12.1" on all devices, except R1
- On R1 enter "ntp master 4"
- Enter the following commands on each switch:
  - service timestamps debug datetime msec
  - service timestamps log datetime msec
  - logging buffered 100000 debugging
- Verify all IP addresses are reachable and all switches have learned the time from R1.

IP Addressing:

Router	Interface	IP address
R1	F0/0	10.1.12.1 /24
R2	F0/1	10.1.12.2 /24
SW1	Interface VLAN 12	10.1.12.10 /24
SW2	Interface VLAN 12	10.1.12.20 /24
SW3	Interface VLAN 12	10.1.12.30 /24
SW4	Interface VLAN 12	10.1.12.40 /24

Task 1

Configure ONLY the following switches to operate in rapid-pvst mode:

- SW1
- SW3
- SW4

Examine the output of the **"Show spanning-tree vlan 12"** for any changes.

The Catalyst 3550 and 3560 can operate in one of three spanning-tree modes:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

1. PVST, per-VLAN Spanning-tree, the default 802.1D mode
2. Rapid-pvst, rapid per-VLAN spanning-tree, based on IEEE standard 802.1w
3. MST, multiple spanning-tree, based on IEEE standard 802.1s

Note that the switch is configured for just one mode at a time.

Enable Rapid Spanning-tree on SW1, SW3 and SW4

```
SWx(config)#Spanning-tree mode rapid-pvst
```

To Verify the configuration:

On SW1:

```
SW1#Show spanning-tree summary | Inc mode

Switch is in rapid-pvst mode
```

On SW3:

```
SW3#Show spanning-tree summary | Inc mode

Switch is in rapid-pvst mode
```

On SW4:

```
SW4#Show spanning-tree summary | Inc mode

Switch is in rapid-pvst mode
```

On SW2:

```
SW2#Show spanning-tree summary | Inc mode

Switch is in pvst mode
```

When you examine the output of the "Show spanning-tree vlan 12," you will see NO difference for SW1 and SW2, and only a minor difference on SW3 and SW4. The bridge priorities, timer values, port costs and port IDs are identical. RSTP and 802.1D use the same algorithm and the same criteria to determine the active topology. On SW3 and SW4 you should notice the following differences:

```
SW3#Show spanning-tree vlan 12 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
---------------- ---- --- --------- -------- -------------
Fa0/19           Desg FWD 19        128.19   P2p
Fa0/21           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.23   P2p Peer(STP)

SW4#Show spanning-tree vlan 12 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -------------
Fa0/19           Altn BLK 19        128.19   P2p
Fa0/21           Desg FWD 19        128.21   P2p Peer(STP)
Fa0/23           Root FWD 19        128.23   P2p
```

SW3 and SW4 are both in the RSTP mode, but they detected a PVST-mode switch on the indicated ports. SW3 and SW4 will send standard 802.1D BPDUs on these border ports.

Task 2

Use debug output to explain how RSTP achieves rapid convergence when SW3 loses its root port.

- Shut SW3's F0/23 and SW4's F0/21 to simplify the problem
- Verify that SW3's F0/21 and SW4's F0/23 are root ports
- Verify that SW3's F0/19 is a designated port, and SW4's F0/19 is in discarding (blocking) state
- On SW3 and SW4 enter the following commands
  - Clear logging
  - Debug spanning-tree events
  - Debug spanning-tree synchronization
- Shut SW3's F0/21 and observe the logs on these three switches

Shut the ports connected to the PVST switch:

On SW3:

```
SW3(config)#int f0/23
SW3(config-if)#shut
```

On SW4:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW4(config)#int f0/21
SW4(config-if)#shut
```

Verify the starting STP topology:

On SW3:

```
SW3#Show spanning-tree vlan 12

[output omitted for brevity]

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.19   P2p
Fa0/21           Root FWD 19        128.21   P2p

SW4#Show spanning-tree vlan 12

[output omitted for brevity]

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.19   P2p
Fa0/23           Root FWD 19        128.23   P2p
```

When SW3 loses its root port, F0/19 on SW4 will need to transition to forwarding state. Standard 802.1D would require up to 50 seconds (20 seconds Max Age plus twice the Forward Delay timer, 15 seconds).

Clear the logs and start debugs on SW1 and SW3:

```
SWx#Clear logging
Clear logging buffer [confirm]

SWx#Debug spanning-tree events
Spanning Tree event debugging is on

SWx#Debug spanning-tree synchronization
Spanning Tree state sync support debugging is on
```

Shutdown SW3 F0/21 and Analyze the results:

```
SW3(config)#int f0/21
SW3(config-if)#shut

RSTP(12): updt roles, root port Fa0/21 going down
RSTP(12): we become the root bridge
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
RSTP(12): updt roles, received superior bpdu on Fa0/19
RSTP(12): Fa0/19 is now root port
RSTP(12): synced Fa0/19
RSTP(12): transmitting an agreement on Fa0/19 as a response to a
proposal
```

According to the Debug output above, RSTP was able to determine the proper role for F0/21 and move F0/19 to forwarding state in less than 20 milliseconds. Port F0/19 became active and immediately sent a BPDU claiming to be the root bridge. 8 milliseconds later it received a BPDU with the proposal bit set from SW4. SW3 synchronized all of its ports (only F0/19 is active) and replied to the proposal with an agreement, completing its transition. Here you see the final state on SW3.

```
SW3#Show spanning-tree vlan 12

VLAN0012
[output removed for brevity]

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Root FWD 19        128.19   P2p Peer(STP)
```

Below you see the debug output on SW4. (NTP has synchronized clocks to the second, but not to the millisecond, so the sequence may not always be precisely comparable to SW3).

From this output you can see that RSTP took less than 10 milliseconds to determine its role and move to forwarding state.

```
SW4#Sh log

[output omitted for brevity]

18:45:31.076: RSTP(12): updt roles, received superior bpdu on Fa0/19
18:45:31.076: RSTP(12): Fa0/19 is now designated
18:45:31.080: RSTP(12): transmitting a proposal on Fa0/19
18:45:31.084: RSTP(12): received an agreement on Fa0/19
```

As soon as F0/19 received the first BPDU from SW3 it changed its role from discarding to designated. SW4 sent a proposal to SW3 and went to forwarding state as soon as it received the agreement. Below you see the final state for SW4:

```
SW4#Show spanning-tree vlan 12 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.19   P2p Peer(STP)
Fa0/23           Root FWD 19        128.23   P2p
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

In 802.1D this would be a scenario where the Cisco backbonefast feature could be used to timeout Max Age and speed convergence from 50 seconds to ONLY 30 seconds. The point-to-point link type and explicit proposal-agreement process of RSTP permits convergence in less than 1 second.

Task 3

Use debug output to demonstrate the RSTP convergence process as SW3's F0/21 returns to its role as root port, SW3 F0/19 transitions from root port back to designated port, and SW4's F0/19 transitions from designated port back to discarding.

- Clear the logs on SW1, SW3 and SW4 and verify spanning-tree debugging is enabled on these switches.
- No-shut (activate) SW3 F0/21 and analyze the results.

This task builds on Task 2 and restores the original RSTP topology.

Clear logs and verify spanning-tree debugging on SW1, SW3 and SW4:

```
SWx#clear logging
Clear logging buffer [confirm]

SWx#Show debugging

Spanning Tree:
  Spanning Tree event debugging is on
  Spanning Tree state sync support debugging is on
```

Activate SW3 F0/21 and analyze the results:

Below you see part of the resulting log entries for SW1 as SW4 activated its port. In just a few milliseconds it sent a proposal toward SW4, received an agreement an moved the port to forwarding state.

```
20:01:07.960: RSTP(12): initializing port Fa0/21
20:01:07.960: RSTP(12): Fa0/21 is now designated
20:01:07.969: RSTP(12): transmitting a proposal on Fa0/21
20:01:07.969: RSTP(12): received an agreement on Fa0/21
20:01:09.965: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to up
20:01:10.972: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/21,
changed state to up
```

On SW3

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3(config)#Int F0/21
SW3(config-if)#No shut
```

As you see below, SW3's F0/21 took from 07.955 to 07.971, about 16 milliseconds, to determine it should be the root port and transition to forwarding. As part of the synchronization process, SW3 temporarily blocks all other non-edge ports, including; in this example F0/19.

```
20:01:07.955: RSTP(12): initializing port Fa0/21
20:01:07.955: RSTP(12): Fa0/21 is now designated

20:01:07.959: RSTP(12): transmitting a proposal on Fa0/21

20:01:07.967: RSTP(12): updt roles, received superior bpdu on Fa0/21
20:01:07.967: RSTP(12): Fa0/21 is now root port
20:01:07.967: RSTP(12): Fa0/19 blocked by re-root
20:01:07.967: RSTP(12): synced Fa0/21
20:01:07.967: RSTP(12): Fa0/19 is now designated
20:01:07.971: RSTP(12): transmitting an agreement on Fa0/21 as a response to a
proposal
20:01:07.971: RSTP(12): transmitting a proposal on Fa0/19
20:01:07.975: RSTP(12): transmitting a proposal on Fa0/19
20:01:08.099: RSTP(12): transmitting a proposal on Fa0/19
20:01:08.439: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to up
20:01:09.439: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/21,
changed state to up
20:01:10.099: RSTP(12): transmitting a proposal on Fa0/19
20:01:12.099: RSTP(12): transmitting a proposal on Fa0/19
20:01:14.099: RSTP(12): transmitting a proposal on Fa0/19
20:01:16.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:18.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:20.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:22.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:22.968: RSTP(12): Fa0/19 fdwhile Expired
20:01:24.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:26.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:28.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:30.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:32.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:34.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:36.100: RSTP(12): transmitting a proposal on Fa0/19
20:01:37.968: RSTP(12): Fa0/19 fdwhile Expired
```

F0/19 sends a proposal every Hello time (2 seconds). Since SW4 F0/19 will go to discarding state, it does not reply with an agreement. Without an explicit agreement, SW3's F0/19 must go through 15 seconds Listening state and 15 seconds Learning state before it can move to forwarding.

The output on SW4 is quite brief: It received the superior BPDU from SW3 on F0/19, immediately transitioned the port from designated to discarding (alternate), and ignored the incoming proposals.

```
SW4#Show log
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
[output removed for brevity]

20:01:07.970: RSTP(12): updt roles, received superior bpdu on Fa0/19
20:01:07.970: RSTP(12): Fa0/19 is now alternate
```

In this task you saw that RSTP rapid convergence depends on the explicit proposal/agreement process. Links could still take 30 seconds to become active.

Task 4

Use debug output to demonstrate the importance of link type in RSTP convergence.

- On BOTH SW3 and SW4:
  - Clear the logs
  - Verify spanning-tree debugging is enabled
  - Configure the shared link-type for F0/19
- Shutdown SW3 F0/21 and observe the result.

In Task 2 you saw that the blocking port F0/19 on SW4 went to forwarding state almost immediately when SW3 lost its root port. In Task 3 you saw that rapid convergence depends on the explicity proposal/agreement process.

The proposal/agreement process can only be used on point-to-point links, where switches can be certain of having only one possible neighboring switch. By default, switches associate the point-to-point link type with full-duplex links and the shared link type with half-duplex links. However the link type can be manually configured.

On SW3 and SW4 "Clear logs" and change F0/19 link type to shared:

```
SWx#clear logging
Clear logging buffer [confirm]

Switch#Show debugging

Spanning Tree:
  Spanning Tree event debugging is on
  Spanning Tree state sync support debugging is on

SWx(config-if)#int f0/19
SWx(config-if)#spanning-tree link-type shared
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Verify the configuration:

```
SW3#Sh spanning-tree vlan 12 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ----
Fa0/19           Desg FWD 19        128.19   Shr Peer(STP)
Fa0/21           Root FWD 19        128.21   P2p

SW4#Show spanning-tree vlan 12 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----

Fa0/19           Altn BLK 19        128.19   Shr Peer(STP)
Fa0/23           Root FWD 19        128.23   P2p
```

Shut SW3 F0/21 and observe the results:

```
SW3(config)#int f0/21
SW3(config-if)#shu

SW3#Show log
[output omitted for brevity]

21:00:41.578: RSTP(12): updt roles, root port Fa0/21 going down
21:00:41.578: RSTP(12): we become the root bridge
21:00:41.594: RSTP(12): updt roles, received superior bpdu on Fa0/19
21:00:41.594: RSTP(12): Fa0/19 is now root port

SW3#Show spanning-tree vlan 12 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Root FWD 19        128.19   Shr Peer(STP)
```

As you see in the output above, SW3 rapidly changed the role of its F0/19 port from designated to root port. But you will recall that in Task 2 SW3 responded to SW4's proposals with an agreement. With a shared link instead of the point-to-point link, SW3 ignores the proposals. Below you see that SW4 F0/19 did eventually transition from blocked to designated. Without an explicit agreement from SW3, however, SW4 has to wait 30 seconds before it can move the port to forwarding state.

```
SW4#Sh log
[output removed for brevity]

21:00:41.589: RSTP(12): updt roles, received superior bpdu on Fa0/19
21:00:41.589: RSTP(12): Fa0/19 is now designated
```


© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

```
21:00:41.593: RSTP(12): transmitting a proposal on Fa0/19
21:00:43.161: RSTP(12): transmitting a proposal on Fa0/19
21:00:45.161: RSTP(12): transmitting a proposal on Fa0/19
21:00:47.162: RSTP(12): transmitting a proposal on Fa0/19
21:00:49.162: RSTP(12): transmitting a proposal on Fa0/19
21:00:51.162: RSTP(12): transmitting a proposal on Fa0/19
21:00:53.162: RSTP(12): transmitting a proposal on Fa0/19
21:00:55.162: RSTP(12): transmitting a proposal on Fa0/19
21:00:56.590: RSTP(12): Fa0/19 fdwhile Expired
21:00:57.162: RSTP(12): transmitting a proposal on Fa0/19
21:00:59.162: RSTP(12): transmitting a proposal on Fa0/19
21:01:01.162: RSTP(12): transmitting a proposal on Fa0/19
21:01:03.162: RSTP(12): transmitting a proposal on Fa0/19
21:01:05.162: RSTP(12): transmitting a proposal on Fa0/19
21:01:07.162: RSTP(12): transmitting a proposal on Fa0/19
21:01:09.162: RSTP(12): transmitting a proposal on Fa0/19
21:01:11.163: RSTP(12): transmitting a proposal on Fa0/19
21:01:11.591: RSTP(12): Fa0/19 fdwhile Expired

SW4#Show spanning-tree vlan 12 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/19           Desg FWD 19        128.19   Shr Peer(STP)
Fa0/23           Root FWD 19        128.23   P2p
```

Task 5

Using debug output, explain how RSTP switches interoperate with switches in legacy, 802.1D mode.

- Reactivate F0/21 on SW3 and remove the "link-type" commands from F0/19 on SW3 and SW4.
- Clear the logs on SW3 and SW4 and verify spanning-tree debugging
- Enter the following commands on SW2
  - clear logging
  - debug spanning-tree events
  - debug spanning-tree synchronization
- Reactivate SW3 interface F0/23 and observe the result
- Clear the logs on SW2, reactivate SW4 interface F0/21 and observe the result

Recall that CAT-2 is in the default, PVST mode.

Prepare SW3, SW4 and SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3(config)#int f0/21
SW3(config-if)#No shut

SW3(config)#int f0/19
SW3(config-if)#No spanning-tree link-type shared

SW3#Clear logg
Clear logging buffer [confirm]

SW4#Clear logg
Clear logging buffer [confirm]

SW4(config)#int f0/19
SW4(config-if)#No spanning-tree link-type shared

SW2#Clear logg
Clear logging buffer [confirm]


SW2#Debug spann event
Spanning Tree event debugging is on

SW2#Debug spanning-tree sync
Spanning Tree state sync support debugging is on
```

Reactivate SW3 F0/23 and observe the result:

```
SW3(config)#int f0/23
SW3(config-if)#No shut

22:08:04.457: %LINK-3-UPDOWN: Interface FastEthernet0/23, changed state to down
22:08:04.789: RSTP(12): initializing port Fa0/23

22:08:04.789: RSTP(12): Fa0/23 is now designated
22:08:04.793: RSTP(12): transmitting a proposal on Fa0/23
22:08:05.445: RSTP(12): transmitting a proposal on Fa0/23

22:08:06.789: %LINK-3-UPDOWN: Interface FastEthernet0/23, changed state to up

22:08:07.445: RSTP(12): transmitting a proposal on Fa0/23
22:08:07.789: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/23,
changed state to up
22:08:19.790: RSTP(12): Fa0/23 fdwhile Expired
22:08:34.790: RSTP(12): Fa0/23 fdwhile Expired
```

Note that SW3 did send a few proposals on the link, but it received no replies, and had to wait 30 seconds for the Listening and learning states to expire before it could transition F0/23 to forwarding state. Below you see the final state:

© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

```
SW3#Show spanning-tree vlan 12 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.19   P2p
Fa0/21           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.23   P2p Peer(STP)
```

Here is what the process looked like from the SW2 perspective. SW2 ignores the proposals sent by SW3 and goes through the normal 802.1D process for an activating root port, taking 30 seconds.

```
SW2#show log
[output omitted for brevity]

22:08:04.735: set portid: VLAN0012 Fa0/23: new port id 8019
22:08:04.735: STP: VLAN0012 Fa0/23 -> listening

22:08:06.739: %LINK-3-UPDOWN: Interface FastEthernet0/23, changed state to up
22:08:07.746: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/23,
changed state to up


.Apr 22 22:08:09.432: STP: VLAN0012 heard root  4108-001a.a174.1f00 on Fa0/23
22:08:09.432:     supersedes 16396-0019.e780.6d80
22:08:09.432: STP: VLAN0012 new root is 4108, 001a.a174.1f00 on port Fa0/23, cost 38
22:08:19.742: STP: VLAN0012 Fa0/23 -> learning

22:08:34.750: STP: VLAN0012 sent Topology Change Notice on Fa0/23
22:08:34.750: STP: VLAN0012 Fa0/23 -> forwarding
```

Reactivate SW4 F0/21 and observe the result:

```
SW2#clear log
Clear logging buffer [confirm]

SW4(config)#int f0/21
SW4(config-if)#No shut

22:26:20.284: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to down
22:26:20.737: RSTP(12): initializing port Fa0/21
22:26:20.737: RSTP(12): Fa0/21 is now designated
22:26:20.741: RSTP(12): transmitting a proposal on Fa0/21
22:26:21.389: RSTP(12): transmitting a proposal on Fa0/21
22:26:22.737: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to up
22:26:23.389: RSTP(12): transmitting a proposal on Fa0/21
22:26:23.737: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/21,
changed state to up
22:26:25.389: RSTP(12): transmitting a proposal on Fa0/21
22:26:35.737: RSTP(12): Fa0/21 fdwhile Expired
22:26:50.738: RSTP(12): Fa0/21 fdwhile Expired
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

SW4 interface f0/21 sent a few proposals, but waited the full 30 seconds before going to forwarding state. Below you see its final state:

```
SW4#Show spanning-tree vlan 12 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.19   P2p
Fa0/21           Desg FWD 19        128.21   P2p Peer(STP)
Fa0/23           Root FWD 19        128.23   P2p
```

And finally, you see the spanning-tree operations on SW2 as its F0/21 interface became active and its final state:

```
SW2#Show log

[output removed for brevity]

22:26:20.747: set portid: VLAN0012 Fa0/21: new port id 8017
22:26:20.747: Found no corresponding dummy port for instance 12, port_id 128.25
22:26:20.747: STP: VLAN0012 Fa0/21 -> listening

22:26:22.752: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to up
22:26:23.758: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/21,
changed state to up
22:26:27.382: STP: VLAN0012 Fa0/21 -> blocking

SW2#Show spanning-tree vlan 12 | B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/2            Desg FWD 19        128.4    P2p
Fa0/21           Altn BLK 19        128.23   P2p
Fa0/23           Root FWD 19        128.25   P2p
```

In this task you saw that RSTP and PVST mode switches do interoperate, and provide convergence similar to that in a purely PVST environment.

Task 6

In your final RSTP lab task enable SW2 for RSTP mode. Make sure that none of the edge ports flap when the switched infrastructure reconverges. Use debugging output on SW2 to explain your results.

- Shut SW2 F0/21 and F0/23.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

- Enable the required RSTP configurations
- Clear the logs on SW2
- Reactivate SW2 interfaces F0/21 and F0/23 at the same time and observe the results

Shut SW2 F0/21 and F0/23:

```
SW2(config)#interface range f0/21, f0/23
SW2(config-if-range)#shut
```

Enable RSTP mode on SW2 and clear the logs:

```
SW2(config)#spanning-tree mode rapid-pvst

SW2#clear logg
Clear logging buffer [confirm]
```

Reactivate SW2 F0/21 and F0/23 at the same time and observe:

```
SW2(config)#interface range f0/21, f0/23
SW2(config-if-range)#NO shut

SW2#Sh log
[output removed for brevity]

22:52:11.005: RSTP(12): initializing port Fa0/21

22:52:11.005: RSTP(12): Fa0/21 is now designated
22:52:11.013: RSTP(12): transmitting a proposal on Fa0/21
22:52:11.047: RSTP(12): initializing port Fa0/23
22:52:11.047: RSTP(12): Fa0/23 is now designated
22:52:11.055: RSTP(12): transmitting a proposal on Fa0/23
22:52:11.449: RSTP(12): updt roles, received superior bpdu on Fa0/21
22:52:11.449: RSTP(12): Fa0/21 is now root port
22:52:11.449: RSTP(12): syncing port Fa0/2
22:52:11.449: RSTP(12): syncing port Fa0/23
22:52:11.449: RSTP(12): synced Fa0/21
22:52:11.458: RSTP(12): transmitting an agreement on Fa0/21 as a response to a
proposal
22:52:11.458: RSTP(12): transmitting a proposal on Fa0/2
22:52:11.458: RSTP(12): transmitting a proposal on Fa0/23
22:52:11.466: RSTP(12): updt roles, received superior bpdu on Fa0/23
22:52:11.466: RSTP(12): Fa0/23 is now root port
22:52:11.466: RSTP(12): Fa0/21 blocked by re-root
```

Down to this point everything looks as it should. Within less than half a second RSTP started forwarding on root port F0/23 and blocked F0/21. The problem can be seen in the very next message.

```
Apr 22 22:52:11.466: RSTP(12): syncing port Fa0/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Apr 22 22:52:11.466: RSTP(12): synced Fa0/23
Apr 22 22:52:11.475: RSTP(12): transmitting an agreement on Fa0/23 as a response to a
proposal
Apr 22 22:52:11.483: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:11.550: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:13.010: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to up
Apr 22 22:52:13.052: %LINK-3-UPDOWN: Interface FastEthernet0/23, changed state to up
Apr 22 22:52:13.563: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:15.577: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:17.590: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:19.603: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:21.617: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:23.630: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:25.643: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:26.457: RSTP(12): Fa0/2 fdwhile Expired
Apr 22 22:52:27.656: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:29.670: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:31.683: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:33.696: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:35.710: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:37.723: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:39.736: RSTP(12): transmitting a proposal on Fa0/2
Apr 22 22:52:41.464: RSTP(12): Fa0/2 fdwhile Expired
```

Before F0/23 can send an agreement it must sync all non-edge ports. In this case, RSTP disabled F0/2, which connects to R2 for 30 seconds. To avoid this, enable portfast on all edge ports.

Enable portfast on SW1 F0/1 and SW2 F0/1:

```
SW1(config)#int f0/1
SW1(config-if)#spanning-tree portfast

SW2(config)#int f0/2
SW2(config-if)#spanning-tree portfast
```

Verification:

On SW1:

```
SW1#Show spanning-tree vlan 12 B Inter

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/1            Desg FWD 19        128.3    Edge P2p
Fa0/21           Desg FWD 19        128.23   P2p
Fa0/23           Desg FWD 19        128.25   P2p

SW2#Sh spanning-tree vlan 12 B Inter
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------
Fa0/2            Desg FWD 19        128.4    Edge P2p
Fa0/21           Altn BLK 19        128.23   P2p
Fa0/23           Root FWD 19        128.25   P2p
```

If you now shut and then reactivate SW2 F0/21 and F0/23, you will see that F0/2 no longer synchronizes and remains active.

```
SW2(config)#int range f0/21, f0/23
SW2(config-if-range)#shut

SW2(config)#int range f0/21, f0/23
SW2(config-if-range)#No shut

23:18:27.671: RSTP(12): initializing port Fa0/23
23:18:27.671: RSTP(12): Fa0/23 is now designated
23:18:27.679: RSTP(12): transmitting a proposal on Fa0/23
23:18:27.696: RSTP(12): updt roles, received superior bpdu on Fa0/23
23:18:27.696: RSTP(12): Fa0/23 is now root port
23:18:27.696: RSTP(12): synced Fa0/23
23:18:27.704: RSTP(12): transmitting an agreement on Fa0/23 as a response to a
proposal
23:18:27.754: RSTP(12): initializing port Fa0/21
23:18:27.754: RSTP(12): Fa0/21 is now designated
23:18:27.763: RSTP(12): transmitting a proposal on Fa0/21

23:18:27.771: RSTP(12): updt roles, received superior bpdu on Fa0/21
23:18:27.771: RSTP(12): Fa0/21 is now alternate

23:18:29.675: %LINK-3-UPDOWN: Interface FastEthernet0/23, changed state to up
23:18:29.759: %LINK-3-UPDOWN: Interface FastEthernet0/21, changed state to up
23:18:30.682: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/23,
changed state to up

23:18:30.766: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/21,
changed state to up
```

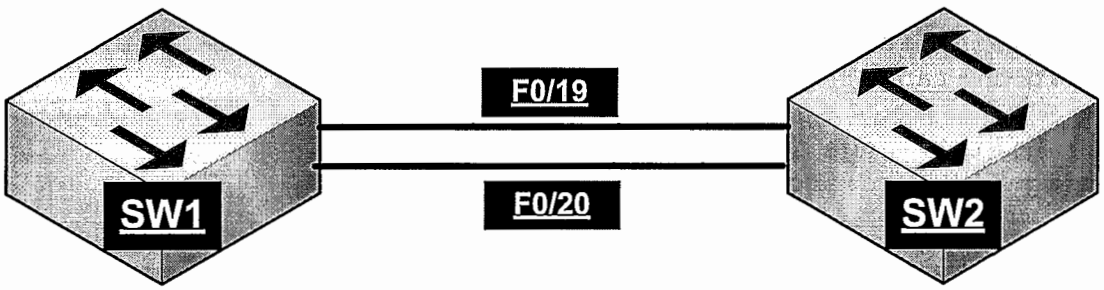
Task 7

Erase the "config.text" and "vlan.dat" and reload the switches before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 15
Multiple Spanning Tree Protocol (802.1s)



Task 1

The first Catalyst switch should be configured with a hostname of SW1 and the second Catalyst should have a hostname of SW2; The third and the fourth switch should have all their ports in Administratively down state.

On the first Switch:

```
Switch(config)#Hostname SW1
```

On the Second Switch:

```
Switch(config)#Hostname SW2
```

On the third and the fourth switch:

```
Switch(config)#Interface range f0/1-24
Switch(config-if-range)#Shut
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 2

Configure all the ports except F0/19 and F0/20 in shutdown mode.

On Both Switches:

```
SWx(config)#int range f0/1-18 , F0/21-24
SWx(config-if-range)#Shut
```

Task 3

Ports F0/19-20 on both switches should be in trunking mode, these ports should use an industry standard protocol to establish the trunk.

On Both Switches:

```
(config)#int range f0/19-20
(config-if-range)#Switchport trunk encapsulation dot1q
(config-if-range)#Switchport mode trunk
```

Task 4

Create VLANs 12, 34, 56, and 90 on SW1 and ensure that these VLANs are propagated to SW2 via VTP messages. Use any name to accomplish this task.

On SW1:

```
SW1(config)#vlan 12,34,56,90
SW1(config-vlan)#exit
```

To verify the configuration:

On SW1:

```
SW1#Sh vlan br | exc unsup
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/21, Fa0/22
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active
34   VLAN0034                         active
56   VLAN0056                         active
90   VLAN0090                         active
```

On SW2:

```
SW2#Sh vlan br | exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/21, Fa0/22
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
```

Note: None of the switches are in VTP transparent mode, and yet the VLANs are not getting propagated from SW1 to SW2. This is because the VTP domain name is not configured, if the VTP domain name is not configured, the switches will NOT propagate their VLAN information across the trunk links. Since the task states to use any name to accomplish this task, a VTP domain name of "TST" is configured:

On SW1:

```
SW1(config)#Vtp domain TST
```

Note the above command configures a VTP domain name, if the other switch does not have a domain name configured and a trunk has been established between the two switches, SW1 will convey the domain name via VTP messages and the two switches will synch up their VLAN information based on the highest VTP rev number. In this task, since a name has not been specified, a domain name of "TST" has been configured.

To verify the configuration:

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2#Sh 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, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/21, Fa0/22

                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
12   VLAN0012                         active
34   VLAN0034                         active
56   VLAN0056                         active
90   VLAN0090                         active
```

You can see that the VLANs are now propagated.

Task 5

Configure Multi-instance of Spanning Tree on these two switches using the following policy:

1. There should be two instances of STP, instance 1 and 2
2. The revision number should be 1
3. The MST region name should be "CCIE"
4. Instance 1 should handle VLANs 12 and 34
5. Instance 2 should handle VLAN 56 and 90
6. All future VLANs should use instance 0
7. Instance 1 should use F0/19
8. Instance 2 should use F0/20
9. SW1 should be the root bridge for the first instance
10. SW2 should be the root bridge for the second instance

On Both Switches:

The default mode for spanning-tree is PVST, the output of the following Show command verifies this information:

```
SWx#Show spanning-tree summary | Inc mode

Switch is in pvst mode
```

← **The default mode of Spanning-tree**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To change the spanning-tree mode to MST:

On Both Switches:

```
SWx(config)#Spanning-tree mode mst
```

To verify the configuration:

On Both Switches:

```
SWx#Sh spanning-tree summary | Inc mode

Switch is in mst mode (IEEE Standard)
```

To configure MST on the switches:

On Both Switches:

```
SWx(config)#Spanning-tree mst configuration
```

The above command enters the MST configuration mode

```
SWx (config-mst)#Revision 1
```

The above command sets the MST configuration revision number to 1. The range for this number is 1-65535.

```
SWx (config-mst)#Name CCIE
```

The above command configured the name of the region to be "CCIE"

```
SWx (config-mst)#Instance 1 vlan 12,34
SWx (config-mst)#Instance 2 vlan 56,90
SWx (config-mst)#exit
```

MST supports 16 instances, once the spanning-tree mode is changed to MST and the MST configuration mode is entered, instance 0 is created automatically and all VLANs are mapped to that instance.

The above commands map the requested VLAN(s) to the specified instances, and by default all the future VLANs or VLAN/s that are not statically mapped to a given instance will be assigned to instance 0, instance 0 is the Catch-all instance.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify this configuration:

On both Switches:

```
SWx #Show spanning-tree mst configuration

Name      [CCIE]
Revision  1     Instances configured 3

Instance  Vlans mapped
--------  ----------------------------------------
0         1-11,13-33,35-55,57-89,91-4094

1         12,34
2         56,90
-------------------------------------------
```

Note: VLANs 12 and 34 are mapped to instance 1, VLANs 56 and 90 are mapped to instance 2, and the rest of the VLANs are mapped to instance 0.

To Verify the configuration before configuring the next portion of the task:

On SW1:

```
SW1#Show spanning-tree bridge

                                                    Hello  Max  Fwd
MST Instance                Bridge ID               Time  Age  Dly  Protocol
--------------- ---------------------------------- -----  ---  ---  --------
MST0            32768 (32768,    0) 0012.7f40.9380    2   20   15  mstp
MST1            32769 (32768,    1) 0012.7f40.9380    2   20   15  mstp
MST2            32770 (32768,    2) 0012.7f40.9380    2   20   15  mstp
```

Note this command displays the BID for your switch (This is NOT the BID of the root Bridge), and instead of assigning a BID to each VLAN, there is a BID for each instance, the priority is incremented based on the instance number.

For example:

For Instance 0, 0 is added to the default priority of 32768
For Instance 1, 0 is added to the default priority of 32769
For Instance 2, 0 is added to the default priority of 32770

To see the root bridge for a given instance:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1#Show spanning-tree root
                                        Root    Hello Max Fwd
MST Instance          Root ID           Cost    Time  Age Dly  Root Port
---------------- -------------------- --------- ----- --- ---  -----
MST0             32768 0012.7f40.9380         0    2   20  15
MST1             32769 0012.7f40.9380         0    2   20  15
MST2             32770 0012.7f40.9380         0    2   20  15
```

On SW1:

```
SW1#Show ver | Inc Base

Base ethernet MAC Address       : 00:12:7F:40:93:80
```

On SW2:

```
SW2#Show spanning-tree root

                                        Root    Hello Max Fwd
MST Instance          Root ID           Cost    Time  Age Dly  Root Port
---------------- -------------------- --------- ----- --- ---  ------------
MST0             32768 0012.7f40.9380         0    2   20  15  Fa0/19
MST1             32769 0012.7f40.9380    200000    2   20  15  Fa0/19
MST2             32770 0012.7f40.9380    200000    2   20  15  Fa0/19
```

The above command displays the BID of the root bridge for different instances. The output may vary based on the Switch's BID, but in this case everything is pointing to SW1.

To configure the root bridge for instance 1 and 2:
The following commands will change the switch priority such that SW1 will be chosen as the root switch for instance 1 and SW2 will be chosen as the root bridge for instance 2.

By default the "Spanning-tree extend system-id" is configured as part of your startup configuration, because the extended system id is set, the priority must be configured in increments of 4096. Remember the lower value has higher preference.

On SW1:

```
SW1(config)#Spanning-tree mst 1 priority 0
SW1(config)#Spanning-tree mst 2 priority 4096
```

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2(config)#Spanning-tree mst 1 priority 4096
SW2(config)#Spanning-tree mst 2 priority 0
```

To verify the configuration:

On SW1:

```
SW1#Show spanning root
                                        Root    Hello Max Fwd
MST Instance            Root ID         Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ------------
MST0             32768 0012.7f40.9380         0     2   20  15
MST1                 1 0012.7f40.9380         0     2   20  15
MST2                 2 001d.e5d6.0000    200000     2   20  15  Fa0/19
```

The local switch (SW1) is the root bridge for instance 0 and 1. The column that specifies the Root ID shows the priority for MST1 and MST2 as 1 and 2 respectively, the priority is the sum of the instance number plus the Priority. Remember that this switch's priority is set to zero.

Note this switch is not the root for MST2. Another indication that this switch is NOT the root for instance 2 is the root port, remember that the root bridge does not have any ports set as root port for the VLANs or in this case Instances that it's the root bridge for.

On SW2:

```
SW2#Show spanning root
                                        Root    Hello Max Fwd
MST Instance            Root ID         Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ------------
MST0             32768 0012.7f40.9380         0     2   20  15  Fa0/19
MST1                 1 0012.7f40.9380    200000     2   20  15  Fa0/19
MST2                 2 001d.e5d6.0000         0     2   20  15
```

Note SW2 is the root bridge for instance 2, whereas, SW1 is the root for MST instances of 0 and 1.

To configure the last portion of this task, the existing state is displayed in the output of the following show command:

On SW1:

```
SW1#Show spanning int f0/19

Mst Instance      Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
MST0              Desg FWD 200000    128.21   P2p
MST1              Desg FWD 200000    128.21   P2p
MST2              Root FWD 200000    128.21   P2p

SW1#Show spanning int f0/20

Mst Instance      Role Sts Cost      Prio.Nbr Type
----------------- ---- --- --------- -------- -----
MST0              Desg FWD 200000    128.22   P2p
MST1              Desg FWD 200000    128.22   P2p
MST2              Altn BLK 200000    128.22   P2p
```

On SW2:

```
SW2#Show spanning int f0/19


Mst Instance      Role Sts Cost      Prio.Nbr Type
----------------- ---- --- --------- -------- -----
MST0              Root FWD 200000    128.21   P2p
MST1              Root FWD 200000    128.21   P2p
MST2              Desg FWD 200000    128.21   P2p

SW2#Show spanning int f0/20

Mst Instance      Role Sts Cost      Prio.Nbr Type
----------------- ---- --- --------- -------- -----
MST0              Altn BLK 200000    128.22   P2p
MST1              Altn BLK 200000    128.22   P2p
MST2              Desg FWD 200000    128.22   P2p
```

Note based on the output of the above show commands, traffic for all MST instances take port F0/19 and none of the instances are using port F0/20.

To configure items 7 and 8, the "Port-priority" command is used as follows:

On Both switches:

```
SWx(config)#Int F0/19
SWx(config-if)#Spanning-tree mst 1 port-priority 0
SWx(config-if)#Spanning-tree mst 2 port-priority 128

SWx(config)#Int F0/20
SWx(config-if)#Spanning-tree mst 1 port-priority 128
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SWx(config-if)#Spanning-tree mst 2 port-priority 0
```

In this task Port-priority is used when selecting an interface to put into the forwarding state for a given instance; a lower value has a higher priority.

In this case port F0/19 will be used by all the VLANs that are assigned to instances 0 & 1, because it has a higher priority (Lower value), and instance 2 will use port F0/20 because it has been configured with a higher priority (Lower value).

To verify the configuration:

On SW1:

```
SW1#Show spanning-tree int f0/19

Mst Instance        Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------
MST0                Desg FWD 200000    128.21   P2p

MST1                Desg FWD 200000      0.21   P2p
MST2                Altn BLK 200000    128.21   P2p

SW1#Sh spanning-tree int f0/20

Mst Instance        Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
MST0                Desg FWD 200000    128.22   P2p
MST1                Desg FWD 200000    128.22   P2p
MST2                Root FWD 200000      0.22   P2p
```

On SW2:

```
SW2#Show spanning-tree int f0/19

Mst Instance        Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
MST0                Root FWD 200000    128.21   P2p
MST1                Root FWD 200000      0.21   P2p
MST2                Desg FWD 200000    128.21   P2p

SW2#Sh spanning-tree int f0/20

Mst Instance        Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
MST0              Altn BLK 200000    128.22   P2p
MST1              Altn BLK 200000    128.22   P2p
MST2              Desg FWD 200000      0.22   P2p
```

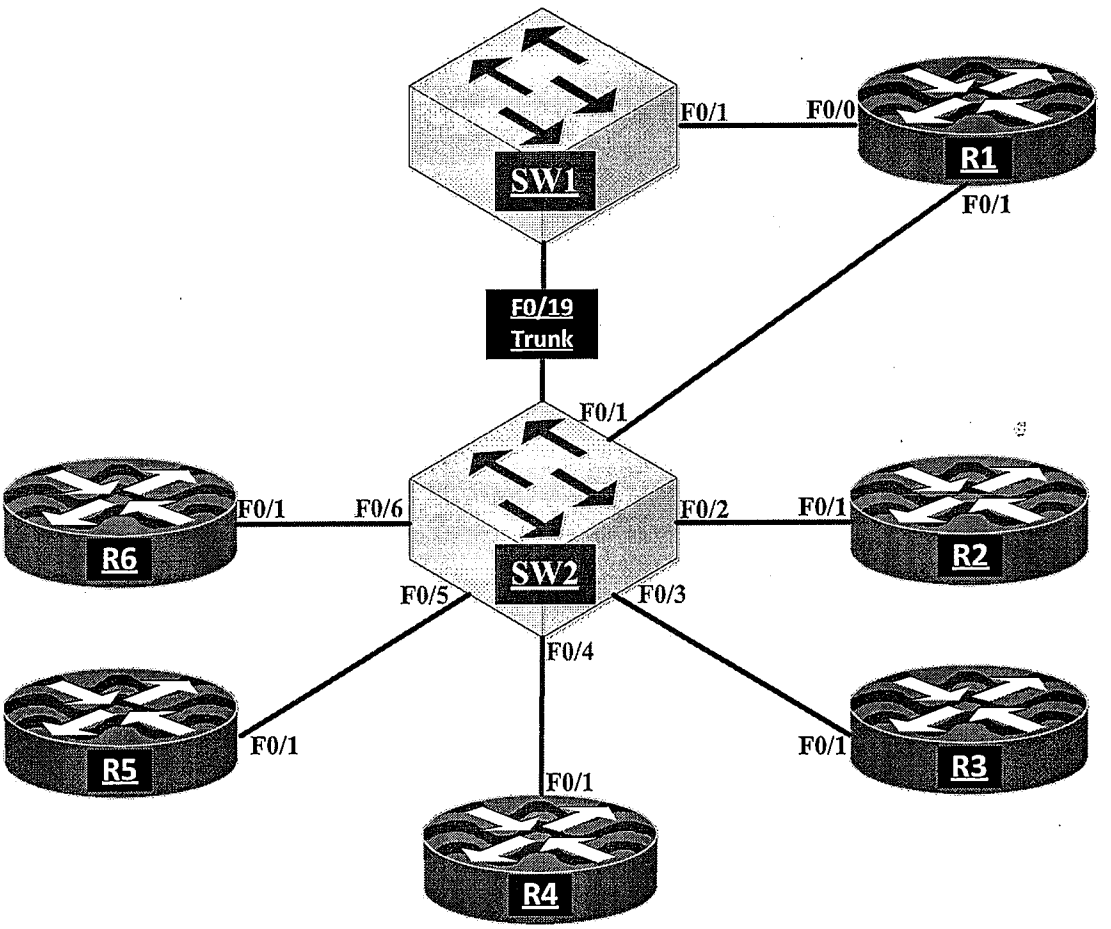
Note instances 0 & 1 use port F0/19 whereas, instance 2 uses port f0/20.

Task 6

Erase the startup configuration and vlan.dat before proceeding to the next lab

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 16
Spanning-Tree PortFast



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Shutdown all ports on SW1 and SW2, and configure an industry standard trunk between the two switches using port F0/19. ONLY the F0/19 interface should be in "No shutdown" state.

On SW1:

```
SW1(config)#Int range f0/1-24
SW1(config-if-range)#Shut

SW1(config)#Int f0/19
SW1(config-if)#Swi tru enca dot
SW1(config-if)#swi mode tru
SW1(config-if)#No shut
```

On SW2:

```
SW2(config)#Int range f0/1-24
SW2(config-if-range)#Shut

SW2(config)#int f0/19
SW2(config-if)#Swi tru enc dot
SW2(config-if)#swi mode tru
SW2(config-if)#No shut
```

To verify the configuration:

On SW2:

```
SW2#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

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      none
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2#Show interface status | Exc disable

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/19                       connected    trunk      a-full  a-100 10/100BaseTX
```

On SW1:

```
SW1#Show interface status | Exc disable

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/19                       connected    trunk      a-full  a-100 10/100BaseTX
```

Task 2

Configure the appropriate switch to implement the following policy:

- Configure the F0/0 interface of R1 in VLAN 100.
- The F0/0 interface of R1 should immediately transition into forwarding state bypassing the Spanning-tree listening and learning states.
- Test this configuration to ensure that the correct feature is enabled.
- You should use an interface configuration command to accomplish this task.

On SW1:

```
SW1(config)#Int f0/1
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc Vlan 100
SW1(config-if)#No shut
```

In order to accomplish this task, portfast should be enabled on SW1's F0/1 interface:

On SW1:

```
SW1(config)#Int f0/1
SW1(config-if)#Spanning-tree portfast
```

You should see the following console message:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%Warning: portfast should only be enabled on ports connected to a single
 host. Connecting hubs, concentrators, switches, bridges, etc... to this
 interface  when portfast is enabled, can cause temporary bridging loops.
 Use with CAUTION

%Portfast will be configured on FastEthernet0/1 but will only  have
effect when the interface is in a non-trunking mode.
```

To verify the configuration:

On SW1:

```
SW1#Show spanning interface f0/1 portfast

no spanning tree info available for FastEthernet0/1
```

Since the F0/0 interface of R1 is NOT enabled (It's in Shutdown state), you should get the following console message:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#No shut
```

On SW1:

```
SW1#Show spanning interface f0/1 portfast

VLAN0100            enabled
```

To test the configuration:

On SW1:

To test the STP PortFast, "Debug spanning-tree events" is enabled on SW1, and R1's F0/0 interface is "Shutdown" and then, "NO Shutdown", this can emulate a brand new connection to the switch port.

Once the port on R1 is enabled using the "No Shutdown" command, the spanning-tree states can be verified in the output of the debug command:

```
SW1#Debug spanning-tree events
Spanning Tree event debugging is on
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1(config)#int f0/0
R1(config-if)#Shut
R1(config-if)#No shut
```

On SW1:

```
00:58:53: set portid: VLAN0100 Fa0/1: new port id 8003

00:58:53: STP: VLAN0100 Fa0/1 ->jump to forwarding from blocking

00:58:55: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
00:58:56: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state to up
```

STP changes a mesh topology into a star topology which is a loop FREE topology. When a host is connected to a port, the link on the switch comes up and STP starts its calculation on that port, the result of the calculation is the transition of the port into forwarding state. This calculation can take upto 30 – 50 seconds.

The "Spanning PortFast" command can be enabled to allow immediate transition of a given port into forwarding state. Once this feature is enabled, the port transitions into forwarding state immediately.

Task 3

Enable the F0/1 interfaces of R1-R6, and configure SW2 such that all ports except the trunk port/s bypass spanning-tree listening and learning state. DO NOT use an interface configuration mode command on SW2 to accomplish this task.

On R1 – R6:

```
Rx(config)#Int F0/1
Rx(config-if)#Shut
Rx(config-if)#No Shut
```

On SW2:

```
SW2(config)#Spanning-tree portfast default
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

You should see the following console message:

```
%Warning: this command enables portfast by default on all interfaces. You
 should now disable portfast explicitly on switched ports leading to hubs,
 switches and bridges as they may create temporary bridging loops.
```

To verify the configuration:

On SW2:

```
SW2#Show spanning-tree interface F0/6 portfast

no spanning tree info available for FastEthernet0/6
```

Let us enable ports F0/1-6 on SW2 and try the show command again:

On SW2:

```
SW2(config)#Int range f0/1-6
SW2(config-if-range)#No shut

SW2#Show spanning-tree interface F0/6 portfast

VLAN0001         enabled
```

To test the configuration:

On SW2:

```
SW2#Debug spanning-tree events
Spanning Tree event debugging is on
```

On R6:

```
R6(config)#Int F0/1
R6(config-if)#Shut
R6(config-if)#No shut
```

On SW2:

```
01:24:08: set portid: VLAN0001 Fa0/6: new port id 8008

01:24:08: STP: VLAN0001 Fa0/6 ->jump to forwarding from blocking
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
01:24:10: %LINK-3-UPDOWN: Interface FastEthernet0/6, changed state to up
01:24:12: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/6,
changed state
to up
```

Spanning-tree PortFast can be enabled on an interface basis, or it can be enabled globally, when Spanning-tree PortFast is enabled globally, the IOS will apply the Spanning-tree PortFast configuration to all interfaces that are in access mode. To verify this information:

On SW2:

```
SW2#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

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1

SW2#Show spanning-tree interface f0/19 portfast

VLAN0001            disabled
```

NOTE: Spanning-tree portfast is disabled on interface/s that are configured as trunks.

Task 4

Configure VLANs 100 and 200 on SW2.

Configure SW2's F0/5 as a trunk and ONLY allow VLAN 100 and 200 to traverse through this trunk.

R5's F0/1 interface should be configured to provide InterVlan routing using 100.1.1.100 /24 and 200.1.1.100 /24 as the gateway for the hosts in VLANs 100 and 200 respectively.

The above configuration should NOT modify the PortFast feature implemented in the previous task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2(config)#Vlan 100,200
SW2(config-vlan)#Exit

SW2(config)#int f0/5
SW2(config-if)#Swit trunk encapsulation dot1q
SW2(config-if)#Swit mode trunk
SW2(config-if)#Swit Trunk allowed VLAN 100,200
```

To verify the configuration:

On SW2:

```
SW2#Show interface f0/5 trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/5       on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/5       100,200

Port        Vlans allowed and active in management domain
Fa0/5       100,200

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/5       100,200
```

To configure R5:

On R5:

```
R5(config)#Int F0/1.100
R5(config-subif)#Encapsulation Dot1q 100
R5(config-subif)#Ip address 100.1.1.100 255.255.255.0

R5(config-subif)#Int F0/1.200
R5(config-subif)#Encapsulation Dot1q 200
R5(config-subif)#Ip address 200.1.1.100 255.255.255.0
```

NOTE: The last section of this task states "The above configuration should NOT modify the PortFast feature implemented in the previous task". Since in the last task the "Spanning-tree Portfast" was enabled on all interfaces except the trunks, you must ensure that "Spanning-tree Portfast" is still enabled after this configuration. To verify portfast on F0/5 interface of SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2#Show spanning-tree interface F0/5 portfast

VLAN0100            disabled
VLAN0200            disabled
```

NOTE: PortFast is disabled on the F0/5 interface of SW2, this is because F0/5 interface is configured as a trunk. The global "Spanning-tree portfast default" command does NOT affect the ports that are configured as trunk.

Configure the following to enable the PortFast feature on the F0/5 interface of SW2, even though it's configured as a trunk:

On SW2:

```
SW2(config)#Int f0/5
SW2(config-if)#Spanning-tree PortFast Trunk
```

You should see the following console message:

```
%Warning: portfast should only be enabled on ports connected to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to this
interface  when portfast is enabled, can cause temporary bridging loops.
Use with CAUTION
```

To verify the configuration:

On SW2:

```
SW2#Sh spannin inter f0/5 portfast

VLAN0100            enabled
VLAN0200            enabled
```

To test the configuration:

On R5:

```
R5(config)#int f0/1
R5(config-if)#Shut
R5(config-if)#No shut
```

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
set portid: VLAN0100 Fa0/5: new port id 8007
STP: VLAN0100 Fa0/5 ->jump to forwarding from blocking

set portid: VLAN0200 Fa0/5: new port id 8007

STP: VLAN0200 Fa0/5 ->jump to forwarding from blocking

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

Task 5

Disable Spanning-tree PortFast on F0/1 interface of SW2, DO NOT remove the "spanning-tree portfast default" command to accomplish this task.

To verify PortFast state on F0/1 interface of SW2:

On SW2:

```
SW2#Sh spanning-tree inter f0/1 portfast

VLAN0001            enabled
```

To disable PortFast on a given port:

On SW2:

```
SW2(config)#Int F0/1
SW2(config-if)#Spanning-tree portf disable
```

To verify the configuration:

On SW2:

```
SW2#Sh spannin inter f0/1 portfast

VLAN0001            disabled
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

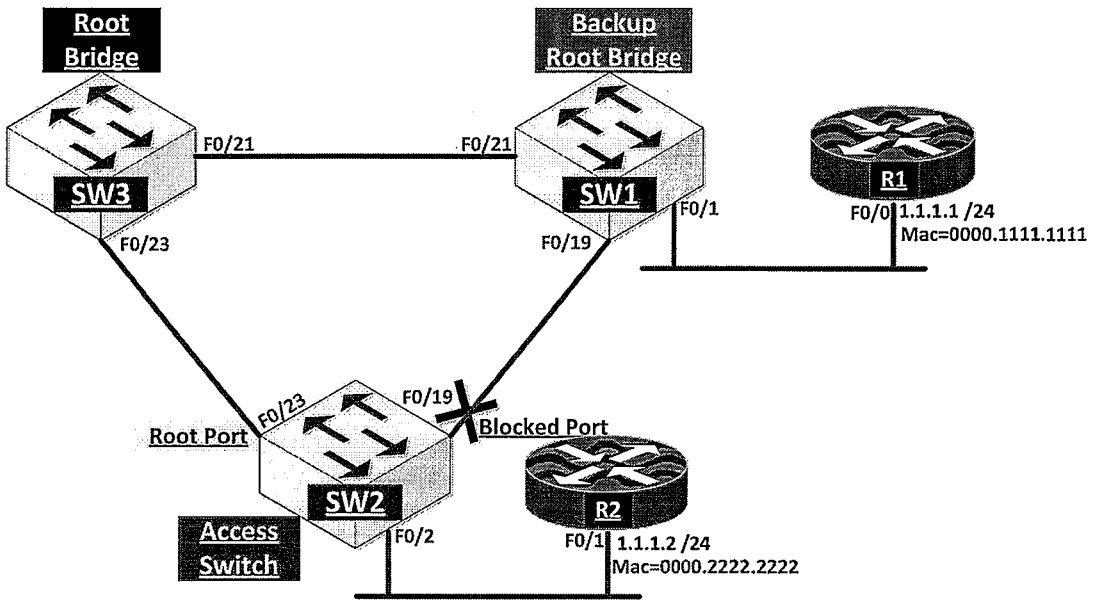
Task 6

Erase the configuration on all routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 17
Configuring UplinkFast



Task 1

Configure the above topology using the interfaces that are identified, the unused interfaces should be in "shutdown" state.

On R1:

```
R1(config-if)#Int F0/0
R1(config-if)#Ip address 1.1.1.1 255.255.255.0
R1(config-if)#mac-address 0000.1111.1111
R1(config-if)#No shu
```

On R2:

```
R2(config)#int f0/1
R2(config-if)#Ip address 1.1.1.2 255.255.255.0
R2(config-if)#mac-address 0000.2222.2222
R2(config-if)#No shu
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW1:

```
SW1(config)#Int range f0/1-24
SW1(config-if-range)#Shut

SW1(config)#Int F0/1
SW1(config-if)#No shut

SW1(config)#Int range f0/19 , F0/21
SW1(config-if-range)#Swi tru enc dot
SW1(config-if-range)#Swi mode trunk
SW1(config-if-range)#No Shut

SW1(config)#Spanning-tree vlan 1 root secondary
```

On SW2:

```
SW2(config)#Int range f0/1-24
SW2(config-if-range)#Shut

SW2(config)#Int F0/2
SW2(config-if)#No shut

SW2(config)#Int range f0/19 , F0/23
SW2(config-if-range)#Swi tru enc dot
SW2(config-if-range)#Swi mode tru
SW2(config-if-range)#No shut
```

On SW3:

```
SW3(config)#Int range f0/1-24
SW3(config-if-range)#Shut

SW3(config)#Int range f0/21 , F0/23
SW3(config-if-range)#Swi tru enc dot
SW3(config-if-range)#Swi mode tru
SW3(config-if-range)#No shut

SW3(config)#Spanning-tree vlan 1 root primary
```

To verify the configuration:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     0011.928e.9c00
              Cost        19
              Port        23 (FastEthernet0/21)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    28673  (priority 28672 sys-id-ext 1)
              Address     0014.6a8d.2280
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
Fa0/1               Desg FWD 19        128.3    P2p
Fa0/19              Desg FWD 19        128.21   P2p
Fa0/21              Root FWD 19        128.23   P2p
```

On SW2:

```
SW2#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     0011.928e.9c00
              Cost        19
              Port        25 (FastEthernet0/23)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     0015.c621.f980
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -----
Fa0/2               Desg FWD 19        128.4    P2p Edge
Fa0/19              Altn BLK 19        128.21   P2p
Fa0/23              Root FWD 19        128.25   P2p
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW3:

```
SW3#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     0011.928e.9c00
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24577  (priority 24576 sys-id-ext 1)
              Address     0011.928e.9c00
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

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

Task 2

Configure SW2 such that if its root port (F0/23) goes down, the alternate port (F0/19) transitions into forwarding state immediately.

On SW2:

```
SW2(config)#spanning-tree uplinkfast
```

UplinkFast is a Cisco proprietary feature that improves convergence time when there is an uplink failure. This feature is designed to run when the local switch has at least one alternate root port, in which case it's in blocking state.

UplinkFast is an access layer solution and it works based on an Uplink Group; an uplink group consists of the root port which is in forwarding state and one or more backup connections to the root bridge which are in blocking state.

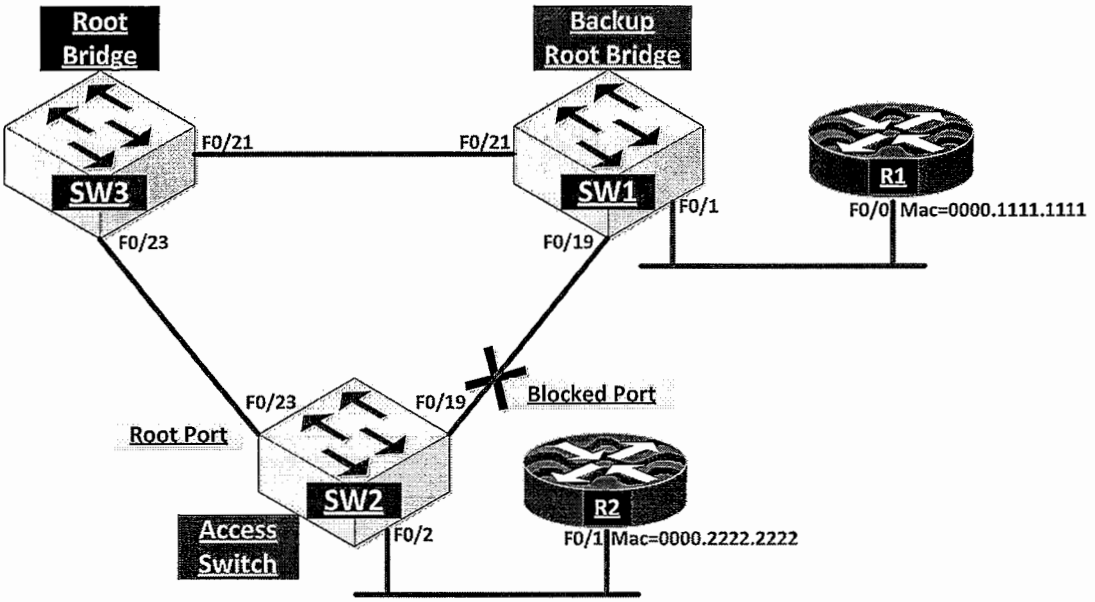
If the root port fails, a port with the next lowest cost from the uplink group will transition into forwarding state and it will become the root port.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

If the uplinkFast is NOT enabled and the primary port on the access switch fails, STP must recalculate and eventually unblock the backup uplink port. This can take up to 30 seconds. With UplinkFast enabled, the same recovery can occur in one second.

The following outlines the events that occur when the UplinkFast is NOT enabled:



- Port F0/23 on the access switch (SW2) goes down.
- The access switch declares its Uplink to SW1 as down.
- The access switch looks at its F0/19 interface, which is still receiving BPDUs from the root bridge through SW1.
- The access switch needs to transition its F0/19 interface into forwarding state, this process means that the port has to go through "Listening" and "learning" states. Each state takes 15 seconds by default unless the "forward-delay" command or the "Diagram" option has been modified.

The following outlines the events that occur when the UplinkFast is configured:

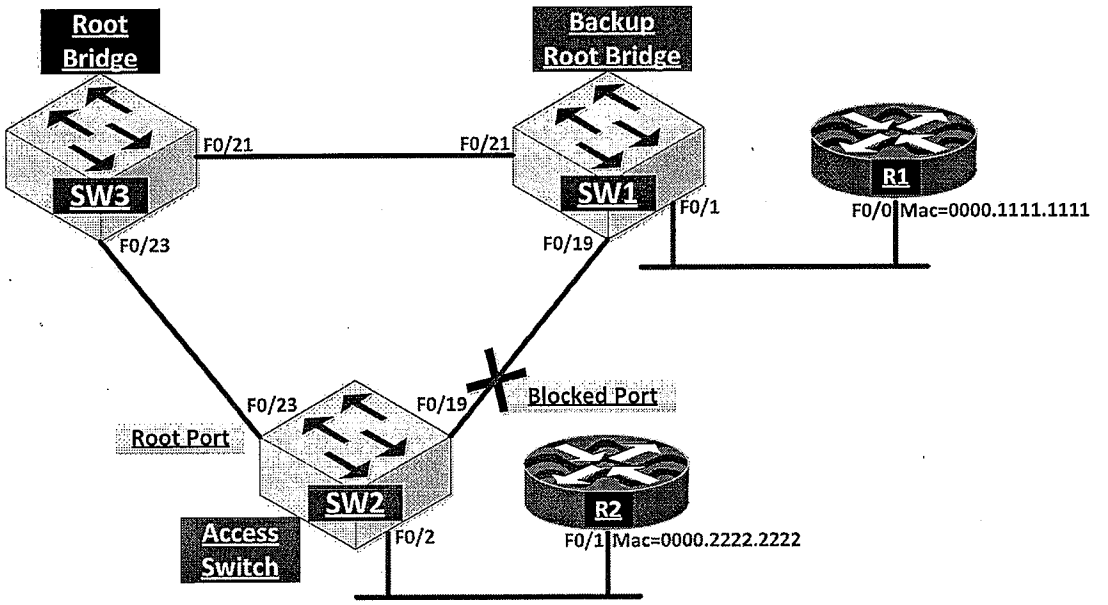
The UplinkFast feature is based on the definition of an uplink group. The Uplink group consists of the root port and all the ports that provide an Alternate/backup connection to the root bridge (Blocked Ports). If the root port fails, a port with the next lowest cost from the group is placed into immediate forwarding state. The uplink group of the Access switch consists of F0/19 and F0/23 interfaces.

- The F0/23 interface of the Access-switch fails.
- The Access-switch detects a link down on its port F0/23.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

- The Access-switch immediately knows that the link to the root bridge is lost, and that the other path that is in the uplink group is in blocked state.
- The access switch places the F0/19 interface in forwarding state, this is done immediately.
- The CAM table of the other switches are updated.

NOTE: If the CAM of the other switches are NOT updated, the CAM table of the other switches will be invalid for a short period of time. The following outlines the reason for this behavior:



In order for R1 to reach R2, it has to take the following path:

SW1 → SW3 → SW2 → R2

The CAM table of the switches before the failure is as follows:

On SW3 - F0/23 interface is associated to R2, whereas, F0/21 is associated to R1.

To generate some traffic, R1 is configured to Ping R2 with a high repeat count:

On R1:

```
R1#Ping 1.1.1.2 rep 1000000
```

On SW3:

```
SW3#Sh mac address-table dynamic
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

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

   1    0000.1111.1111    DYNAMIC     Fa0/21 → R1
   1    0000.2222.2222    DYNAMIC     Fa0/23 → R2
   1    0014.6a8d.2297    DYNAMIC     Fa0/21 → SW1's F0/21 interface
   1    0015.c621.f999    DYNAMIC     Fa0/23 → SW2's F0/23 interface

Total Mac Addresses for this criterion: 4
```

On SW1 - **F0/1 interface is associated to R1, and F0/21 interface is associated to R2.**

On SW1:

```
SW1#Sh mac address-table dynamic

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

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    0000.1111.1111    DYNAMIC     Fa0/1 → R1
   1    0000.2222.2222    DYNAMIC     Fa0/21 → R2
   1    0011.928e.9c15    DYNAMIC     Fa0/21 → SW3's F0/21 interface
   1    0015.c621.f995    DYNAMIC     Fa0/19 → SW2's F0/19 interface
Total Mac Addresses for this criterion: 4
```

On SW2 - **On SW2 (The Access switch), R1 is associated to interface F0/23 and R2 is associated to F0/2 interface as well.**

On SW2:

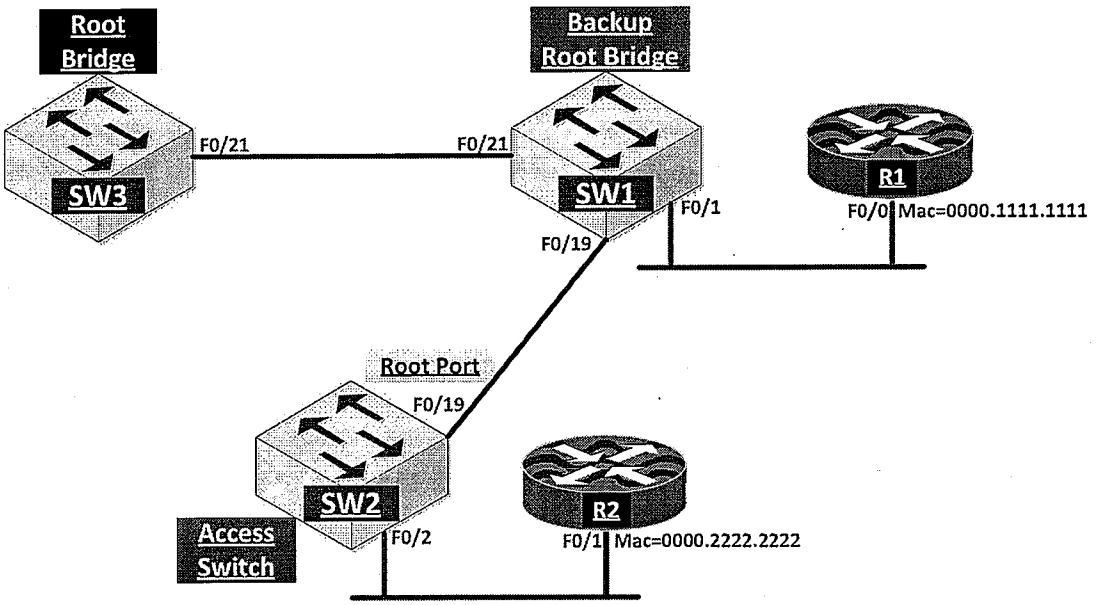
```
SW2#Sh mac address-table dynamic
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    0000.1111.1111    DYNAMIC     Fa0/23 → R1
   1    0000.2222.2222    DYNAMIC     Fa0/2 → R2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   1    0011.928e.9c17    DYNAMIC     Fa0/23 → SW3's F0/23 interface
Total Mac Addresses for this criterion: 3
```

But once STP converges, the actual topology looks like the following:



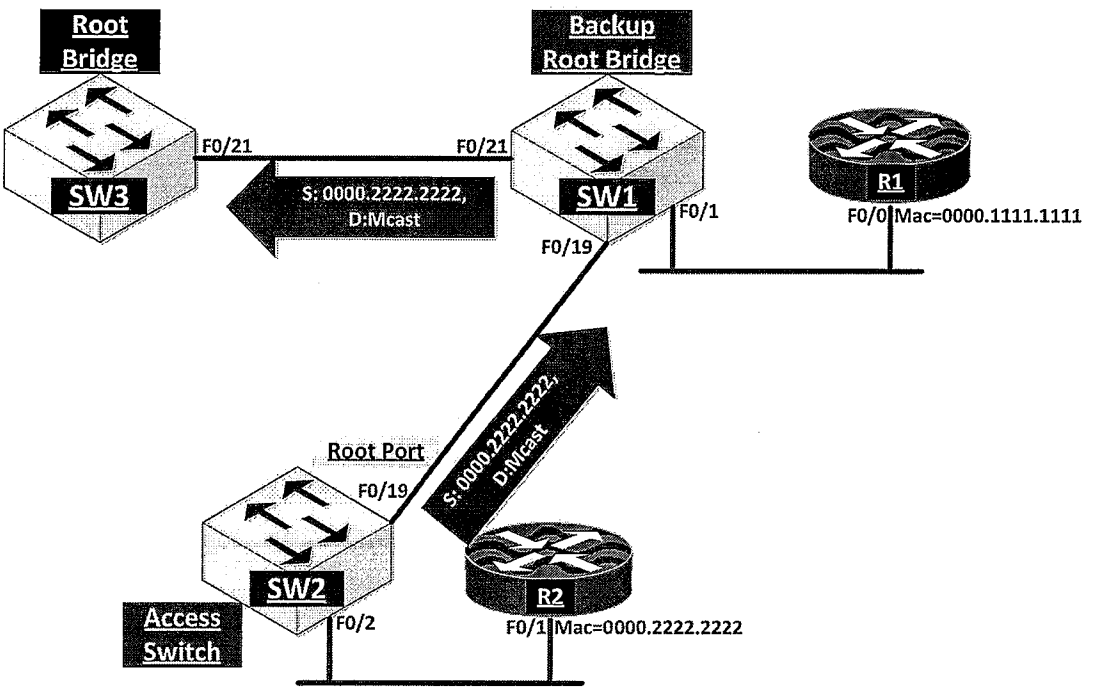
The backup link is brought up quickly, But from CAM's perspective, things have not changed yet and therefore, it will be inaccurate.

In order to solve this problem, the access switch (SW2) begins to flood dummy packets with different MAC addresses that it has in its CAM; it uses those MAC addresses as the source, and with Cisco's proprietary Multicast MAC address as the destination.

This method will ensure that the dummy packets are flooded on the entire network and the CAM of all switches are updated as follows:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



The CAM table of the switches will look like the following after the flooding:

Before we see the Mac address-table of the switches, let's shutdown the F0/23 interface of SW2:

On SW2:

```
SW2(config)#Int F0/23
SW2(config-if)#Shut
```

To verify the configuration:

On SW3:

```
SW3#Show mac address-table dynamic

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

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    0000.1111.1111    DYNAMIC     Fa0/21 → R1
   1    0014.6a8d.2297    DYNAMIC     Fa0/21 → SW1's F0/21 interface
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Total Mac Addresses for this criterion: 2
```

On SW1:

```
SW1#Sh mac address-table dynamic
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    0000.1111.1111    DYNAMIC     Fa0/1 → R1
   1    0000.2222.2222    DYNAMIC     Fa0/19 → R2
   1    0011.928e.9c15    DYNAMIC     Fa0/21 → SW3's F0/21 interface
   1    0015.c621.f995    DYNAMIC     Fa0/19 → SW2's F0/19 interface
Total Mac Addresses for this criterion: 4
```

On SW2:

```
SW2#Sh mac address-table dynamic
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    0000.1111.1111    DYNAMIC     Fa0/19 → R1
   1    0000.2222.2222    DYNAMIC     Fa0/2 → R2
   1    0014.6a8d.2295    DYNAMIC     Fa0/19 → SW1's F0/19 interface
Total Mac Addresses for this criterion: 3
```

To verify the operation of STP before the configuration of UplinkFast, the "UplinkFast" feature is disabled and the "Debug Spanning-tree events" is enabled:

On SW2:

To see if the UplinkFast feature is enabled:

```
SW2#Show spanning-tree summary | Inc Uplink

UplinkFast                   is enabled
UplinkFast statistics
```

To Disable the UplinkFast feature:

```
SW2(config)#No Spanning-tree Uplinkfast
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To see if the UplinkFast feature is disabled:

```
SW2#Show spanning-tree summary | Inc Uplink

UplinkFast                   is disabled

SW2#Debug spanning-tree events
 Spanning Tree event debugging is on
```

Shutdown the root port on SW2:

```
SW2(config)#Int F0/23
SW2(config-if)#No Shut

SW2(config-if)#Shut
```

You should see the following debug output:

```
00:43:57: STP: VLAN0001 new root port Fa0/19, cost 38
00:43:57: STP: VLAN0001 Fa0/19 -> listening

*Mar  1 00:43:59.173: %LINK-5-CHANGED: Interface FastEthernet0/23,
changed state to administratively down

00:43:59: STP: VLAN0001 sent Topology Change Notice on Fa0/19

*Mar  1 00:44:00.180: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/23, changed state to down

00:44:12: STP: VLAN0001 Fa0/19 -> learning

00:44:27: STP: VLAN0001 sent Topology Change Notice on Fa0/19
00:44:27: STP: VLAN0001 Fa0/19 -> forwarding
```

NOTE: It took 30 seconds for STP to converge; let's enable "Spanning-tree uplinkfast" and test the same scenario again:

On SW2:

```
SW2(config)#Int F0/23
SW2(config-if)#No shut
```

To configure UplinkFast:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW2:

NOTE: The STP UplinkFast can ONLY be configured in the global configuration mode and as a result of that, it effects the entire switch.

```
SW2(config)#Spanning-tree uplinkfast
```

To verify the primary and the alternate Uplink ports:

```
SW2#Show spanning-tree vlan 1 | Inc Fa0

Fa0/2               Desg FWD 3019      128.4    P2p Edge
Fa0/19              Altn BLK 3019      128.21   P2p
Fa0/23              Root FWD 3019      128.25   P2p
```

With the "Debug spanning-tree events" still configured, the F0/23 interface of SW2 is shutdown and the convergence time is verified:

```
SW2(config)#Int F0/23
SW2(config-if)#Shut
```

You should see the following debug output:

```
01:07:37: STP: VLAN0001 new root port Fa0/19, cost 3038

01:07:37.418: %SPANTREE_FAST-7-PORT_FWD_UPLINK: VLAN0001 FastEthernet0/19
moved to Forwarding (UplinkFast).

01:07:39.406: %LINK-5-CHANGED: Interface FastEthernet0/23, changed state
to administratively down

01:07:39: STP: VLAN0001 sent Topology Change Notice on Fa0/19

01:07:40.413: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/23, changed state to down

SW2#Show spanning-tree vlan 1 | Inc Fa0

Fa0/2               Desg FWD 3019      128.4    P2p Edge
Fa0/19              Root FWD 3019      128.21   P2p
```

You can see that it ONLY took a second for the F0/19 interface to transition into forwarding state.

To ensure stability and to protect itself from constant link flaps, when port F0/23 is enabled, STP does

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

NOT converge immediately. Besides, SW2's F0/23 interface does NOT transition into forwarding state before it transitions through "listening" and "learning" states, which can take up to 30 seconds.

This switchover is delayed by (2*forward-delay + 5 seconds), which means 35 seconds by default. The reason the 5 additional seconds is added is to allow other protocols like DTP time to negotiate, if any is used.

It is important to be aware of some of the changes that UplinkFast applies. Once UplinkFast is configured, some of the STP parameters are modified because of the following reasons:

The Bridge priority of the switch is increased significantly, this ensures that the switch will NOT be elected as the root bridge. To verify this information:

Before the UplinkFast feature was enabled:

```
SW2#Show spanning-tree vlan 1 | Inc Priority

  Root ID     Priority    24577
  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
```

After the UplinkFast feature was enabled:

```
SW2#Show spanning-tree vlan 1 | Inc Priority

  Root ID     Priority    24577
  Bridge ID   Priority    49153  (priority 49152 sys-id-ext 1)
```

To ensure that ports on the local switch do NOT get elected as the designated ports, all the ports on the local switch will have their cost increased by 3000.

Before the UplinkFast feature was enabled:

```
SW2#Show spanning-tree vlan 1 | B Interface
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -------
Fa0/2               Desg FWD 19        128.4    P2p Edge
Fa0/19              Altn BLK 19        128.21   P2p
Fa0/23              Root FWD 19        128.25   P2p
```

After the UplinkFast feature was enabled:

```
SW2#Show spanning-tree vlan 1 | B Interface
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- -------
Fa0/2               Desg FWD 3019      128.4    P2p Edge
Fa0/19              Altn BLK 3019      128.21   P2p
Fa0/23              Root FWD 3019      128.25   P2p
```

Remember that the above changes can seriously alter your current STP topology.

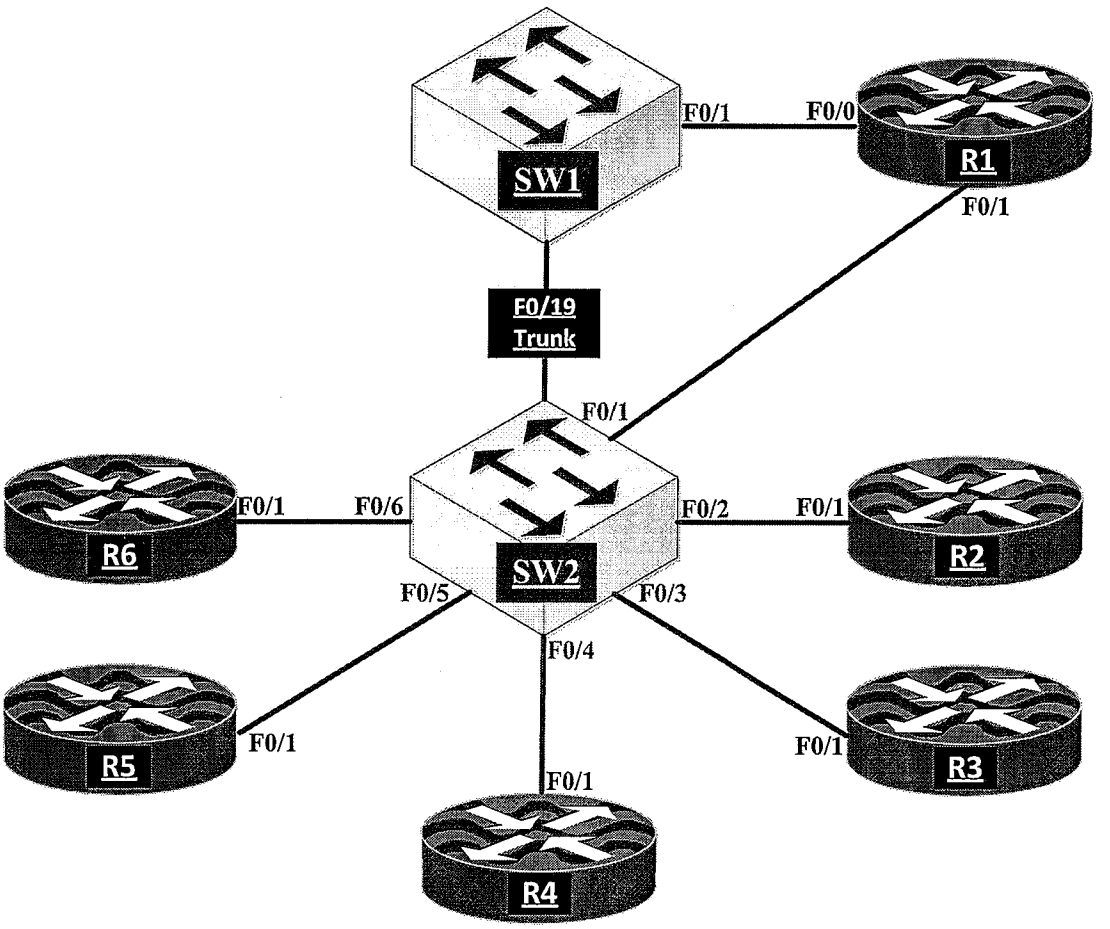
Task 3

Erase the configuration on all routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 18
Configuring BPDUGuard



Task 1

Shutdown all ports on SW1 and SW2, and configure an industry standard trunk between the two switches using port F0/19. ONLY the F0/19 interface should be in "No shutdown" state.

On SW1:

```
SW1(config)#Int range f0/1-24
SW1(config-if-range)#Shut
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1(config)#Int f0/19
SW1(config-if)#Swi tru enca dot
SW1(config-if)#swi mode tru
SW1(config-if)#No shut
```

On SW2:

```
SW2(config)#Int range f0/1-24
SW2(config-if-range)#Shut

SW2(config)#int f0/19
SW2(config-if)#Swi tru enc dot
SW2(config-if)#swi mode tru
SW2(config-if)#No shut
```

Task 2

Configure the F0/1 interface of SW1 such that if it detects any BPDUs, it transitions into "err-disable" mode.

Remember the F0/1 interface of SW1 is connected to R1's F0/0 interface and both ports are in administratively down state.

On SW1:

```
SW1#Show interface f0/1 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/24                       disabled     1            auto   auto 10/100BaseTX
```

The "BPDUGuard" is the feature that can accomplish this task. BPDUGuard can be enabled in two different ways:

- **Globally**
- **On a per interface basis**

Globally:

If the "BPDUGuard" feature is enabled globally, it will affect ports that are Portfast enabled. If a port that is configured as portfast receives BPDUs, the "BPDUGuard' feature will shutdown the port in "Err-isabled"

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

state.

To enable the "BPDUGuard" feature globally, you must use the "Spanning-tree PortFast BpduGuard Default" global configuration command.

On a per-interface basis:

If the "BPDUGuard" feature is enabled on a given interface, and the port receives BPDUs, it will once again shut the port down in "Err-Disable" mode.

NOTE: In this mode the "PortFast" feature does not need to be enabled.

To enable the "BpduGuard" feature on a per interface basis, you must use the "Spanning-tree BpduGuard Enable" interface configuration mode.

On SW1:

```
SW1(config)#int f0/1
SW1(config-if)#Spanning-tree bpduguard enable
SW1(config-if)#No shut
```

To test and verify the configuration:

You must wait for the F0/1 interface of SW1 to transition into forwarding state.

On SW1:

```
SW1#Debug Spanning-tree events
Spanning Tree event debugging is on
```

In order to test this feature, R1 is configured as a bridge sending BPDUs:

On R1:

```
R1(config)#Bridge 1 protocol ieee

R1(config)#Int F0/0
R1(config-if)#Bridge-group 1
R1(config-if)#No Shut
```

On SW1:

```
set portid: VLAN0001 Fa0/1: new port id 8003
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
STP: VLAN0001 Fa0/1 -> listening

%SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port Fa0/1 with BPDU Guard
enabled.

Disabling port.
%PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting Fa0/1 in
err-disable state
```

To verify the result:

On SW1:

```
SW1#Show interface F0/1 Status

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

SW1#Show errdisable detect | Inc Err|bpdu

ErrDisable Reason       Detection     Mode
bpduguard               Enabled       port
```

To verify the state of the port:

On SW1:

```
SW1#Show interface F0/24 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/24                       err-disabled 1            auto   auto 10/100BaseTX
```

Since the port is NOT in "Administratively down" state, a "No Shutdown" will NOT enable the interface; the "No Shutdown" command ONLY works if the port's state is in "Administratively down" state. In order to enable a port that is in "err-disable" mode, the port must be placed in "Administratively down" state first, using the "Shutdown" interface configuration command, and then, it can be enabled by using the "No Shutdown" command.

Task 3

Configure SW2 such that the existing and futue port/s that are "PortFast" enabled transition into "ERR-DISABLE" mode upon detection of BPDUs.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

The following command enables the "BPDUGuard" feature globally on SW2. When this feature is enabled globally, it must also include the "PortFast" feature, therefore, affecting existing and future port/s that are PortFast enabled:

```
SW2(config)#Spanning-tree PortFast bpduguard default
```

To verify the configuration:

On SW2:

```
SW2#Show spanning-tree summary | Inc BPDU Guard

PortFast BPDU Guard Default  is enabled
```

To test the configuration:

SW2's F0/1 interface is configured with PortFast, and the "Debug spanning-tree events" is enabled:

On SW2:

```
SW2(config)#Int F0/1
SW2(config-if)#Spanning-tree PortFast
SW2(config)#No Shut

SW2#Debug spanning-tree events
Spanning Tree event debugging is on
```

On R1:

Since R1 is already configured as a Bridge, you don't need to configure the "Brigde 1 protocol ieee" again; all you need to do is to apply the bridge configuration to F0/1 interface of R1 using the "Bridge-group 1" command, and finally the interface is enabled.

```
R1(config)#Int f0/1
R1(config-if)#bridge-group 1
R1(config-if)#N0 Shut
```

You should see the following debug output on SW2:

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
set portid: VLAN0001 Fa0/1: new port id 8003
STP: VLAN0001 Fa0/1 ->jump to forwarding from blocking

%SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port Fa0/1 with BPDU Guard
enabled.

Disabling port.

%PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting Fa0/1 in
err-disable state
```

To verify the configuration:

On SW2:

```
SW2#Show interface F0/1 status

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

Task 4

Configure SW1's F0/1 interface to automatically recover from errdisable mode caused by BPDUGuard. This recovery should occur every 30 seconds, and if the interface of the switch detects BPDUs, it should once again transition back into errdisable mode, but if it does NOT detect BPDUs, it should remain in up/up state.

On SW1:

```
SW1(config)#errdisable recovery cause bpduguard
SW1(config)#errdisable recovery interval 30
```

To verify the configuration:

On SW1:

```
SW1#Show errdisable detect | Inc Err|bpdu

ErrDisable Reason    Detection status
bpduguard            Enabled
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

You should see the following console messages every 30 seconds as the switch tries to recover from the errdisable mode:

Every 30 seconds the port tries to recover from BPDUGuard on the F0/1 interface:

```
00:54:19: %PM-4-ERR_RECOVER: Attempting to recover from bpduguard err-disable state on Fa0/1
```

The interface comes up:

```
00:54:21: set portid: VLAN0001 Fa0/1: new port id 8003
00:54:21: STP: VLAN0001 Fa0/1 -> listening
00:54:23: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
```

The interface once again detects BPDUs:

```
00:54:23: %SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port
```

The interface (F0/1) has BPDUGuard enabled and transitions into errdisable mode:

```
FastEthernet0/1 with BPDU Guard enabled. Disabling port.

00:54:23: %PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting
Fa0/1 in err-disable state

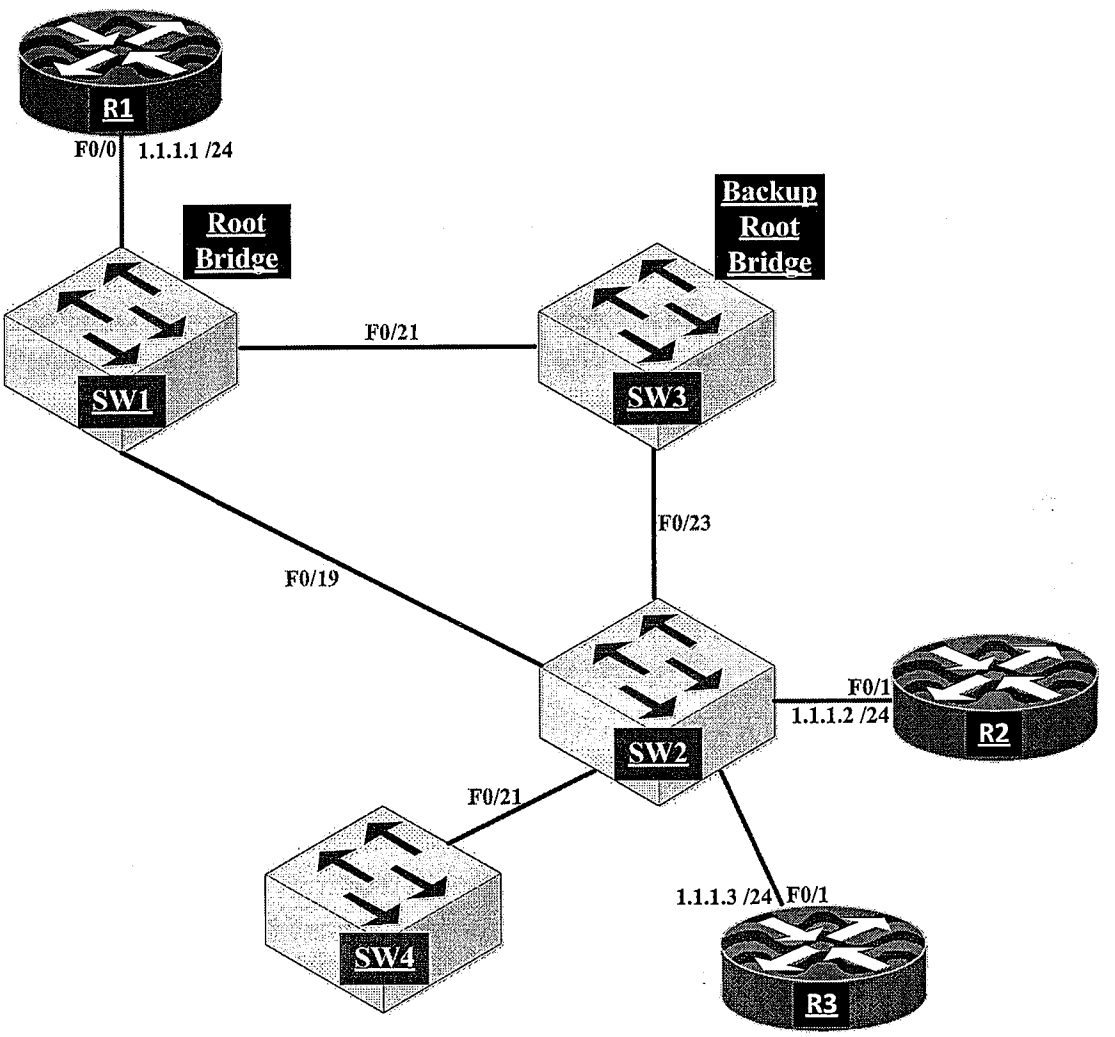
00:54:25: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
```

Task 5

Erase the configuration on all routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 19
Configuring BPDUFilter



Task 1

Configure the routers and the switches according to the above diagram. The inter-switch connection should be configured as trunk using an industry standard encapsulation. The unused ports should be configured in "Administratively Shutdown" mode.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1

```
R1(config)#int f0/0
R1(config-if)#ip addr 1.1.1.1 255.255.255.0
R1(config-if)#No shut
```

On R2:

```
R2(config)#int f0/1
R2(config-if)#ip addr 1.1.1.2 255.255.255.0
R2(config-if)#No shut
```

On R3:

```
R3(config)#int f0/1
R3(config-if)#ip addr 1.1.1.3 255.255.255.0
R3(config-if)#No shut
```

On SW1:

```
SW1(config)#Int range f0/1-24
SW1(config-if-range)#Shut

SW1(config)#Int range f0/19 , f0/21
SW1(config-if-range)#Swi tru enc dot
SW1(config-if-range)#Swi mode trunk
SW1(config-if)#No shut
```

On SW2:

```
SW2(config)#Int range f0/1-24
SW2(config-if-range)#Shut

SW2(config)#Int range f0/19 , f0/21 , f0/23
SW2(config-if-range)#Swi tru enc dot
SW2(config-if-range)#Swi mode trunk
SW2(config-if)#No shut
```

On SW3:

```
SW3(config)#Int range f0/1-24
SW3(config-if-range)#Shut
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW3(config)#Int range f0/21 , f0/23
SW3(config-if-range)#Swi tru enc dot

SW3(config-if-range)#Swi mode trunk
SW3(config-if)#No shut
```

On SW4:

```
SW4(config)#int f0/21
SW4(config-if)#Swi trunk enc dot
SW4(config-if)#Swi mode trunk
SW4(config-if)#No shut
```

On SW1:

```
SW1(config)#Spanning-tree vlan 1 root primary

SW1(config)#Int F0/1
SW1(config-if)#No shut
```

On SW3:

```
SW3(config)#Spanning-tree vlan 1 root secondary
```

On SW2:

```
SW2(config)#Int range F0/2-3
SW2(config-if)#No shut
```

To verify the configuration:

On SW1:

```
SW1#Show inter trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/21      on           802.1q         trunking      1
(The rest of the output is omitted)

SW1#Show spanning-tree vlan 1 | Inc Root ID|Bridge ID

  Root ID     Priority    24577
  Bridge ID   Priority    24577  (priority 24576 sys-id-ext 1)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Based on the output of the above show command, you can see that this switch is the root.

On SW2:

```
SW2#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      1
Fa0/21      on           802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1
(The rest of the output is omitted)

SW2#Show spanning-tree vlan 1 | Inc Root ID|Bridge ID

   Root ID     Priority    24577
   Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
```

On SW3:

```
SW3#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/21      on           802.1q         trunking      1
Fa0/23      on           802.1q         trunking      1
(The rest of the output is omitted)

SW3#Show spanning-tree vlan 1 | Inc Root ID|Bridge ID

   Root ID     Priority    24577
   Bridge ID   Priority    28673  (priority 28672 sys-id-ext 1)
```

Based on the output of the above show command, you can see that this switch is the backup root or root secondary.

On SW4:

```
SW4#Show interface trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/21      on           802.1q         trunking      1
(The rest of the output is omitted)

SW4#Show spanning-tree vlan 1 | Inc Root ID|Bridge ID
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Root ID     Priority    24577
Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
```

Task 2

Configure SW2's F0/21 interface using the following policy:

The F0/21 interface of SW2 should be configured to filter sending or receiving BPDUs. If this port receives BPBUs, it should NOT transition into Err-Disable state.

If BPDU Filtering is enabled on a given interface, the incoming BPDU packets are filtered. Unlike BPDUGuard, it does NOT place the interface of the switch into err-disable state.

Let us check the incoming and outgoing BPDUs before configuring this task:

On SW2:

```
SW2#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 339, received 0

SW2#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 345, received 0
```

Remember that BPDUs are generated by the root bridge and they are forwarded downstream, since SW1 is the root, it sends the BPDUs to SW2, and SW2 is the switch that forwards it downstream to SW4.

On SW4:

```
SW4#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 0, received 210

SW4#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 0, received 220
```

Remember that BPDUs are received from root ports and propagated to downstream switches.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To configure the task:

On SW2:

```
SW2(config)#Int f0/21
SW2(config-if)#Spanning-tree bpdufilter enable
```

To verify the task:

On SW2:

```
SW2#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 581, received 0

SW2#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 581, received 0
```

Since SW1 is the root bridge, SW2 refloods the BPDU messages that it receives from SW1. Since SW4 is NOT the root bridge, it will NOT generate BPDUs.

Once BPDUFilter is configured on SW2's F0/21 interface, SW4 will no longer see the BPDUs, and it will claim to be the root bridge, and it will start generating BPDUs, let us verify this information:

On SW4:

```
SW4#Show spanning-tree interface F0/21 Detail | Inc BPDU

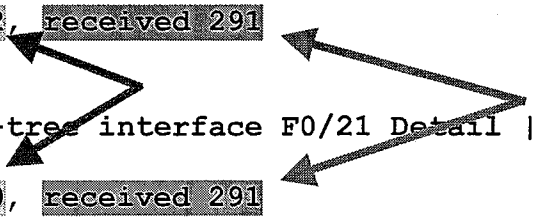
    BPDU: sent 342, received 291

SW4#Show spanning-tree interface F0/21 Detail | Inc BPDU

    BPDU: sent 350, received 291

SW4#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    32769
              Address     000c.ced1.fb80
```



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     000c.ced1.fb80
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

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

NOTE: SW4 is no longer receiving BPDUs from SW2, therefore, it claims to be the root bridge, and as a result of that, it starts generating BPDUs.

But since SW2 is filtering the incoming BPDUs, the receive BPDU counter on SW2 does not increment. Therefore, BPDUFilter filters/Blocks the incoming and outgoing BPDU messages.

To verify the information:

On SW2:

```
SW2#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 585, received 0

SW2#Show spanning-tree interface F0/21 Detail | Inc BPDU

   BPDU: sent 585, received 0
```

Task 3

On the existing topology, demonstrate and prove as to why and how the "BPDUFilter" feature can cause forwarding loop.

To verify the communication between R1 – R3:

On R1:

```
R1#Ping 1.1.1.2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 1.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

On R2:

```
R2#Ping 1.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Let us find out the switch interface that is in Blocking state. SW1 is the root bridge therefore, all of its interfaces should be in forwarding state, let us check:

On SW1:

```
SW1#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------
Number of blocked ports (segments) in the system : 0
```

On SW3:

```
SW3#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------
Number of blocked ports (segments) in the system : 0
```

On SW2:

```
SW2#Show spanning-tree blockedports
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Name                 Blocked Interfaces List
-------------------- ------------------------------------
VLAN0001             Fa0/23

Number of blocked ports (segments) in the system : 1
```

Well.....it was obvious, since SW1 was configured to be the root bridge and SW3 was configured to be the backup root bridge, the blocked port had to be on SW2; but since the F0/19 interface of SW2 is the closest interface to the root bridge (Root port), the only other interface (F0/23) must be in blocking state.

If BPDUFilter is configured on this interface (F0/23), all interfaces on SW2 will be in forewarding state as well, therefore, we will see a spanning-tree forwarding loop. With BPDUFilter configured, the Spanning-tree protocol is disabled. The following demonstrates this fact:

On SW2:

```
SW2(config)#Int F0/23
SW2(config-if)#Spanning-tree bpdufilter enable
```

To verify the configuration:

On SW1:

```
SW1#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------
Number of blocked ports (segments) in the system : 0
```

On SW3:

```
SW3#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------
Number of blocked ports (segments) in the system : 0
```

On SW2:

```
SW2#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------
Number of blocked ports (segments) in the system : 0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Well.....if none of the switches are blocking any of their ports, then, how is STP performing loop avoidance?

Well...it is NOT, and this is why this command/feature must be configured carefully.

You should see the following console messages:

NOTE: Depending on the IOS version you may or may not see the following console messages.

On SW2:

```
%SW_MATM-4-MACFLAP_NOTIF: Host 0019.e8d7.6821 in vlan 1 is flapping
between port Fa0/19 and port Fa0/23

%SW_MATM-4-MACFLAP_NOTIF: Host 0007.b35d.81a1 in vlan 1 is flapping
between port Fa0/23 and port Fa0/19
```

To test the configuration:

On R1, R2 and R3:

```
R1#Clear arp
```

On R1:

```
R1#Ping 1.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R1#Ping 1.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let us remove the command before proceeding to the next task.

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2(config)#Int F0/23
SW2(config-if)#Spanning-tree bpdufilter disable
```

To verify the configuration:

On SW2:

```
SW2#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ---------------------------------
VLAN0001             Fa0/23

Number of blocked ports (segments) in the system : 1

R1#Ping 1.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 4

Configure SW2 based on the following Policy:

- Configure the F0/2 interface of this switch to bypass Spanning-tree listening and learning state.
- This Switch should be configured such that upon detection of BPDU messages on existing and/or future portfast enabled ports, it loses its spanning-tree PortFast status, but the port does NOT transition into err-disable mode.

On SW2:

```
SW2(config)#Spanning-tree portfast bpdufilter default
```

To verify the configuration:

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show spanning-tree summary | Inc BPDU Filter

Portfast BPDU Filter Default is enabled

SW2(config)#Int f0/2
SW2(config-if)#Spanning-tree Portfast
```

To verify the configuration:

On SW2:

```
SW2#Show spanning-tree int f0/2 portfast

VLAN0001            enabled
```

To test and verify the configuration:

On SW2:

```
SW2#Debug Spanning-tree events
Spanning Tree event debugging is on
```

On R2:

```
R2(config)#Bridge 1 protocol ieee

R2(config)#Int F0/1
R2(config-if)#Shut
R2(config-if)#Bridge-group 1
R2(config-if)#No Shut
```

You should see the following console messages on SW2:

On SW2:

```
set portid: VLAN0001 Fa0/2: new port id 8004
STP: VLAN0001 Fa0/2 ->jump to forwarding from blocking

%LINK-3-UPDOWN: Interface FastEthernet0/2, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/2, changed
state to up

STP: VLAN0001 heard root 32768-0011.9279.1f71 on Fa0/2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
STP: VLAN0001 Topology Change rcvd on Fa0/2
STP: VLAN0001 sent Topology Change Notice on Fa0/19
```

You MUST wait for the above message on your console before entering the following show command:

```
SW2#Show spanning-tree int f0/2 portfast

VLAN0001            disabled

SW2#Show interface F0/2 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/2                        connected    1          a-full  a-100 10/100BaseTX

SW2#Show run int F0/2 | B interface

interface FastEthernet0/2
 spanning-tree portfast

SW2#Show run | Inc spanning-tree portfast bpduf

spanning-tree portfast bpdufilter default
```

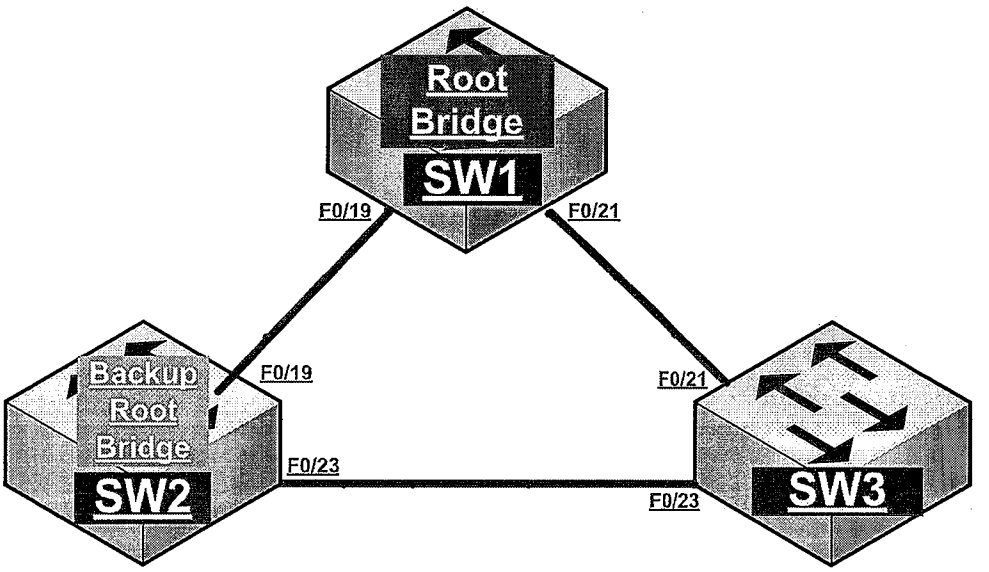
Task 5

Erase the configuration on all routers and switches and reload them before proceeding to the next lab

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 20
Spanning-tree BackBoneFast



Task 1

Configure the switches according to the above diagram. The unused ports must be in "Shutdown" state.

On All Switches:

```
SWx(config)#int range f0/1-24
SWx(config-if-range)#Shut
```

On SW1:

```
SW1(config)#Spanning-tree vlan 1 root primary

SW1(config)#Int range f0/19 , f0/21
SW1(config-if-range)#Swi tru enc dot
SW1(config-if-range)#Swi mode tru
SW1(config-if-range)#No shut
```

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2(config)#Spanning-tree vlan 1 root secondary

SW2(config)#Int range f0/19 , f0/23
SW2(config-if-range)#Swi tru enc dot
SW2(config-if-range)#Swi mode tru
SW2(config-if-range)#No shut
```

On SW3:

```
SW3(config)#Int range f0/21 , f0/23
SW3(config-if-range)#Swi tru enc dot
SW3(config-if-range)#Swi mode trunk
SW3(config-if-range)#No Shut
```

To verify the configuration:

On SW3:

```
SW3#Show Spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ------------
VLAN0001         24577 0014.6a8d.2280        19    2   20  15  Fa0/21
```

The output of the above show command reveals that SW3 can NOT be the root bridge. The question is why?

A root bridge can not have a root port, we can see that the root port is F0/21, with a cost of 19. So what happened to the F0/23 interface?

Well.....the output of the following command reveals that the F0/23 interface of this switch is in blocking state.

```
SW3#Show Spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------
VLAN0001             Fa0/23

Number of blocked ports (segments) in the system : 1
```

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2#Show Spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ----------
VLAN0001         24577 0014.6a8d.2280        19    2   20  15   Fa0/19
```

Once again we can see a root port in the above output, therefore, this switch can NOT be the root. Let us check the state of the F0/23 interface of this switch:

```
SW2#Show Spanning-tree block

Name                 Blocked Interfaces List
-------------------- --------------------------------
Number of blocked ports (segments) in the system : 0
```

Well....it is not in blocking state, let us use another command to determine the state of the F0/23 interface of this switch:

```
SW2#Show spanning-tree vlan 1 | 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
```

We can see that the F0/23 interface of this switch is the designated port.

On SW1:

```
SW1#Show Spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ---------
VLAN0001         24577 0014.6a8d.2280         0    2   20  15
```

Well....we can see that this switch (SW1) is the root bridge, because it has no root port. This switch should NOT have any ports in blocking state, let us verify this information:

```
SW1#Show Spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- --------------------------------
Number of blocked ports (segments) in the system : 0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 2

Configure the appropriate switche/s such that if an indirectly connected link fails, the STP convergence time is improved, DO NOT manually configure any timers to accomplish this task.

To resolve this task, the "Spanning-tree BackBoneFast" feature must be enabled. BackboneFast is a spanning-tree feature that optimizes the convergence by expiring the max-age timer. **BackboneFast MUST be enabled on all switches.**

With "UplinkFast", the switches detected a failure of a directly connected link, with "BackboneFast", the switches detect a failure of an indirectly connected link.

BackboneFast works in two stages:

The first stage:

When a switch receives inferior BPDUs through an NDP (Non Designated Port), it knows that the switch that sent the inferior BPDUs has lost its connection to the Root Bridge.

When the local switch receives the inferior BPDUs it does the following:

- It verifies if the source that generated these messages is from a local segment; if it is from a local segment, it knows that an indirectly connected link failure has occurred.
- If the source of the inferior BPDUs is from a switch that is NOT on a local segment, the local switch will ignore them.

The second stage:

In this stage the local switch goes through a verification process; the switch uses a request and response protocol, this process queries other switches to determine if the connection to the Root Bridge is lost.

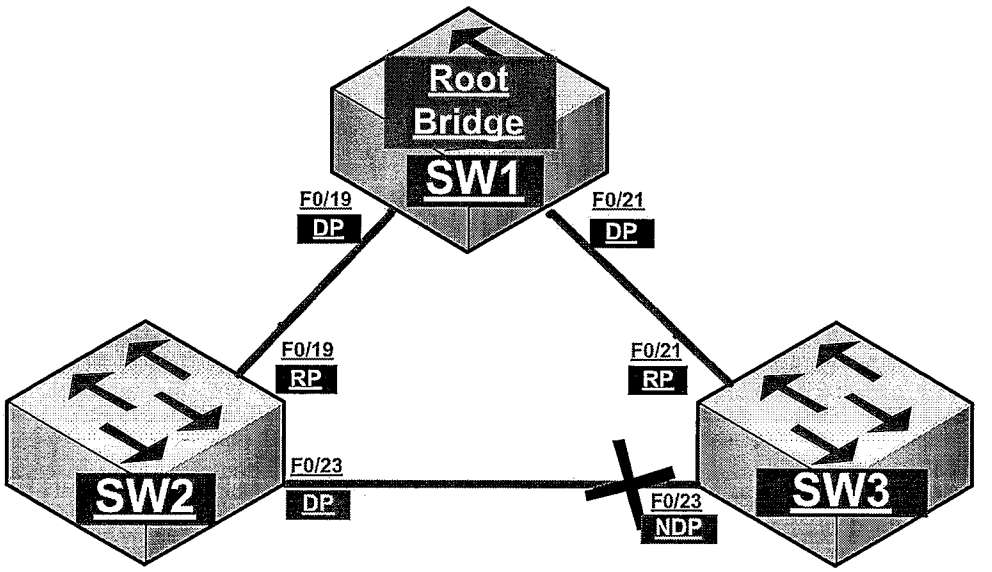
The verification process is done by the local switch that received the inferior BPDUs; The switch generates RLQ (Root Link Query) requests, these messages are sent out of the root port/s.

These messages are sent to query the upstream switch/s if their connection to the root bridge is up. The receiving switch sends RLQ responses to reply to the RLQ requests ONLY if the connection to the root bridge is up. When the switch receives an RLQ on its root port, it knows that the connection to the root bridge is good, therefore, it expires the max-age timer in order to speed up the convergence.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The following outlines the steps that occur when BackboneFast is enabled:



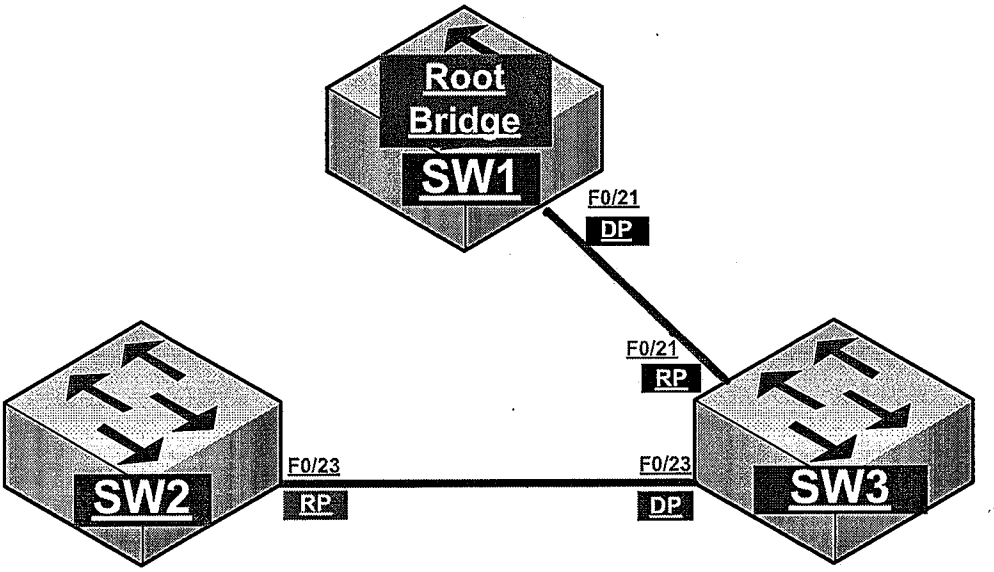
- SW2 loses its connection to SW1, basically interface F0/19 goes down.

- SW2 immediately withdraws interface F0/19 as its Root port and starts generating configuration BPDUs announcing itself as the Root Bridge, these messages are sent out of F0/23 interface.
- SW3's F0/23 receives these C-BPDUs and realizes that they are inferior BPDUs from a switch on a local segment.
- SW3 sends RLQ requests out of its F0/21 interface, in order to verify a stable connection to the Root Bridge.
- SW1 receives the RLQs from SW3, and since it's the root bridge, it replies with RLQ responses.

- SW3 receives the RLQ responses through its Root Port, and immediately expires its max-age timer (20 seconds by default).
- SW3 starts sending C-BPDUs out F0/23 and goes through listening and learning states.

- SW2 receives the C-BPDUs and it knows that it is NOT the Root Bridge, and changes the state of its F0/23 as the RP.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets



To configure BackboneFast:

On All Switches:

```
SWx(config)#Spanning-tree backbonefast
```

To test and verify the configuration:

On SW3:

```
SW3#Debug spanning-tree events
Spanning Tree event debugging is on
```

On SW2:

```
SW2(config)#Int F0/19
SW2(config-if)#Shut
```

You should see the following output of the debug on SW3:

© 2014 Narbik Kocharians, All rights reserved

https://t.me/learningnets

On SW3:

```
00:41:21: STP: VLAN0001 heard root 28673-001b.2be5.1200 on Fa0/23
00:41:21: STP: VLAN0001 Fa0/23 -> listening

00:41:22: STP: VLAN0001 Topology Change rcvd on Fa0/23
00:41:22: STP: VLAN0001 sent Topology Change Notice on Fa0/21

00:41:36: STP: VLAN0001 Fa0/23 -> learning

00:41:51: STP: VLAN0001 sent Topology Change Notice on Fa0/21
00:41:51: STP: VLAN0001 Fa0/23 -> forwarding
```

To verify the result:

On SW1:

```
SW1#Show spanning-tree vlan 1 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- ----------------
Fa0/21           Desg FWD 19        128.23   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

SW3#Show spanning-tree summary | Inc Back

BackboneFast                  is enabled
BackboneFast statistics
```

On SW2:

```
SW2#Show spanning-tree vlan 1 | B Interface

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


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

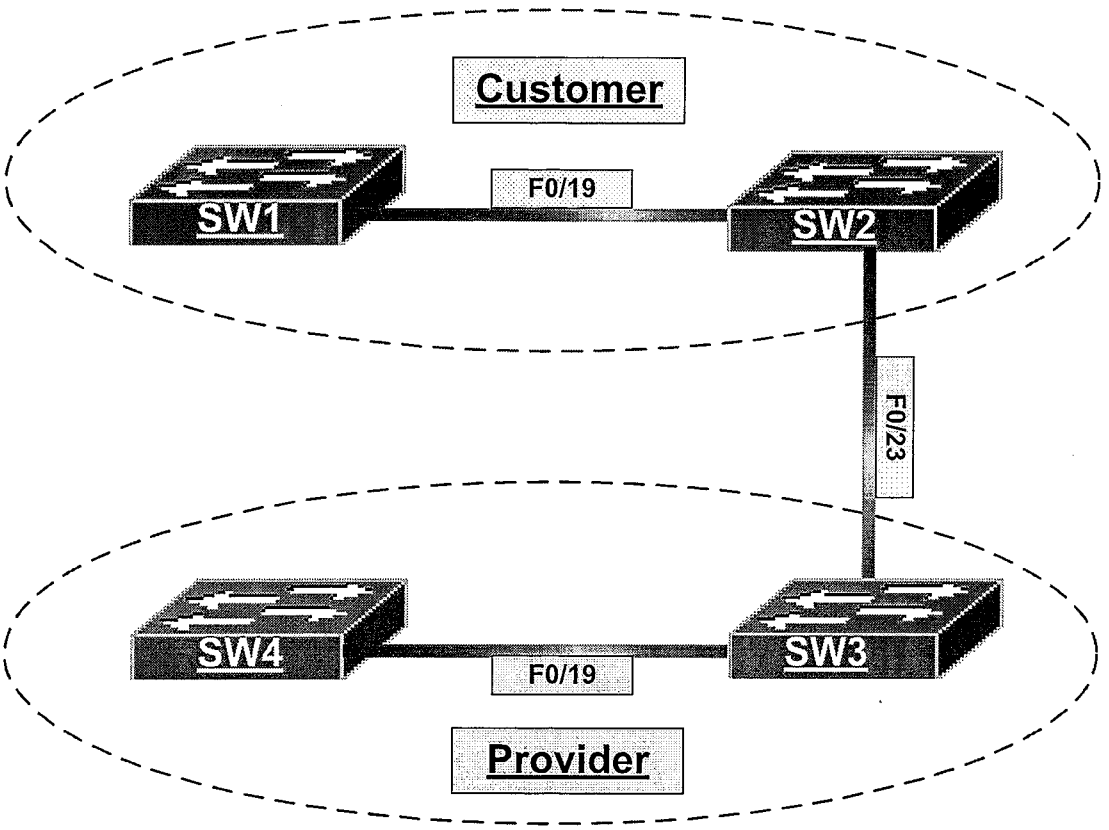
Task 3

Erase the configuration on all routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 21
Spanning-tree RootGuard



Task 1

Configure the switches according to the above diagram.

On All Switches:

```
SW4(config)#Int range f0/1-24
SW4(config-if-range)#Shut
```

On SW1 and SW4:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SWx(config)#Int F0/19
SWx(config-if)#Swi tru enc do
SWx(config-if)#Swi mode trunk
SWx(config-if)#No Shut
```

On SW2 and SW3:

```
SW2(config)#Int range f0/19 , F0/23
SW2(config-if-range)#Swi tru enc dot
SW2(config-if-range)#Swi mode trunk
SW2(config-if-range)#No shut
```

Task 2

Configure the appropriate switch such that the Customer edge switch is never elected as the Root bridge of the provider switched network, no matter how it's configured. DO NOT use BPDUFilter to accomplish this task.

This feature is typically seen in the service provider's network. An SP with layer 2 network can have many switch connections to the customer network; Spanning-tree can reconfigure itself based on the credentials of the customer's switch, and the customer's switch can become the Root Bridge for the SP's switched network.

This situation can be avoided by using the "Root guard" feature; root guard must be enabled on the interface facing the customer's switch.

When this feature enabled, if STP calculations cause the customer's switch to become the root bridge, root guard will immediately and automatically place the interface facing the customer in the "root-inconsistent" state (Which means that the port is in blocking state) to prevent the customer's switch from becoming the root switch.

To test this configuration, the provider's edge switch (SW3) is set up as the root bridge, in this configuration the "Priority" command is used, but a macro can also be used to accomplish this:

On SW3:

```
SW3(config)#Spanning-tree vlan 1 priority 8192
```

To test and verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1#Show spanning vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    8193
              Address     000a.41f2.1080
              Cost        38
              Port        21 (FastEthernet0/19)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     0014.a934.b880
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

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

SW1#Show version | Inc Base

Base ethernet MAC Address       : 00:14:A9:34:B8:80
```

On SW2:

```
SW2#Show spanning vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    8193
              Address     000a.41f2.1080
              Cost        19
              Port        25 (FastEthernet0/23)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              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
---------------- ---- --- --------- -------- -----
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/23           Root FWD 19        128.25   P2p

SW2#Show version | Inc Base

Base ethernet MAC Address       : 00:1B:2B:E5:12:00
```



On SW3:

```
SW3#Show spanning vlan 1

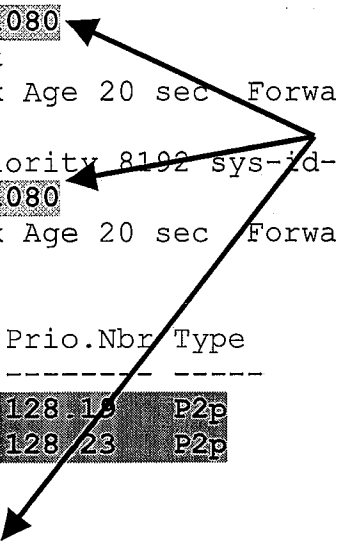
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    8193
              Address     000a.41f2.1080
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    8193   (priority 8192 sys-id-ext 1)
              Address     000a.41f2.1080
              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        128.19   P2p
Fa0/23           Desg FWD 19        128.23   P2p

SW3#Show version | Inc Base

Base ethernet MAC Address: 00:0A:41:F2:10:80
```



On SW4:

```
SW4#Show spanning vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    8193
              Address     000a.41f2.1080
              Cost        19
              Port        19 (FastEthernet0/19)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

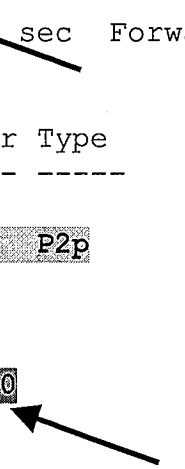
```
              Address     000a.8a3e.b200
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----

Fa0/19           Root FWD 19        128.19   P2p

SW4#Show version | Inc Base

Base ethernet MAC Address: 00:0A:8A:3E:B2:00
```



To see the existing situation:

On SW2 and SW3:

```
SWx#Debug spanning events
Spanning Tree event debugging is on
```

Let's lower the priority for VLAN 1 on SW2, the customer's edge switch:

On SW2:

```
SW2(config)#Spanning-tree vlan 1 priority 4096
```

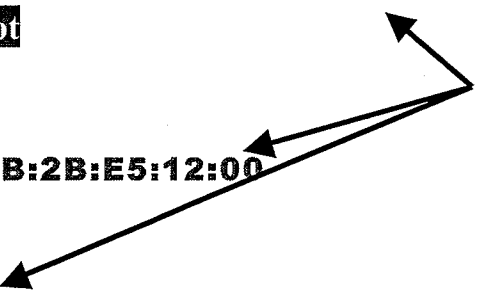
You should see the following debug output on SW2:

setting bridge id (which=1) prio 4097 prio cfg 4096 sysid 1 (on) id 1001.001b.2be5.1200

STP: VLAN0001 we are the spanning tree root

```
SW2#Show version | Inc Base

Base ethernet MAC Address       : 00:1B:2B:E5:12:00
```



On SW3:

STP: VLAN0001 heard root 4097-001b.2be5.1200 on Fa0/23
supersedes 8193-000a.41f2.1080
STP: VLAN0001 new root is 4097, 001b.2be5.1200 on port Fa0/23, cost 19

NOTE: If a switch connecting to your switch has a lower priority, the Spanning-tree calculation can elect

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

that switch to be the Root Bridge. Let's configure "Root Guard" and analyze its behavior:

To remove the configuration on SW2:

On SW2:

```
SW2(config)#NO Spanning-tree vlan 1 priority 4096
```

You should wait for STP to converge, and once it converges, SW3 will once again become the root bridge.

To verify the configuration:

On SW3:

```
SW3#Sh spanning-tree vlan 1 | Inc root
          This bridge is the root
```

To configure Root Guard:

On SW3:

```
SW3(config)#Int F0/23
SW3(config-if)#Spanning-tree guard root
```

You should see the following console message:

%SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port FastEthernet0/23.

NOTE: The root guard feature is implemented using the "Spanning-tree Guard root" interface configuration mode command.

To test and verify the configuration:

On SW2:

```
SW2(config)#Spanning-tree vlan 1 priority 4096
```

You should see the following debug output on SW2:

setting bridge id (which=1) prio 4097 prio cfg 4096 sysid 1 (on) id 1001.001b.2be5.1200
STP: VLAN0001 we are the spanning tree root

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW3:

You should see the following debug output on SW3:

The following message states that F0/23 interface of SW3 just heard that the switch on the other end of the interface (F0/23) is claiming to be the Root:

STP: VLAN0001 heard root 4097-001b.2be5.1200 on Fa0/23

The following debug output states that the switch that claims to be the root, has a better BID and it supersedes our's:

supersedes 8193-000a.41f2.1080

The last message states that Root Guard places the F0/23 interface in blocking state, but you may have missed this message, please stop the debug and scroll up to see the following message:

%SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port FastEthernet0/23 on VLAN0001.

If you notice every 2 seconds you will see the following debug output:

STP: VLAN0001 heard root 4097-001b.2be5.1200 on Fa0/23
supersedes 8193-000a.41f2.1080

Let us disable the Debug on SW3:

On SW3:

```
SW3#Undebug all
All possible debugging has been turned off
```

To test the configuration:

On SW2:

```
SW2#Show spanning vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    4097
              Address     001b.2be5.1200
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    4097   (priority 4096 sys-id-ext 1)

             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
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.25   P2p
```

On SW3:

```
SW3#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    8193
             Address     000a.41f2.1080
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8193   (priority 8192 sys-id-ext 1)
             Address     000a.41f2.1080
             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        128.19   P2p
Fa0/23           Desg BKN*19        128.23   P2p *ROOT_Inc
```

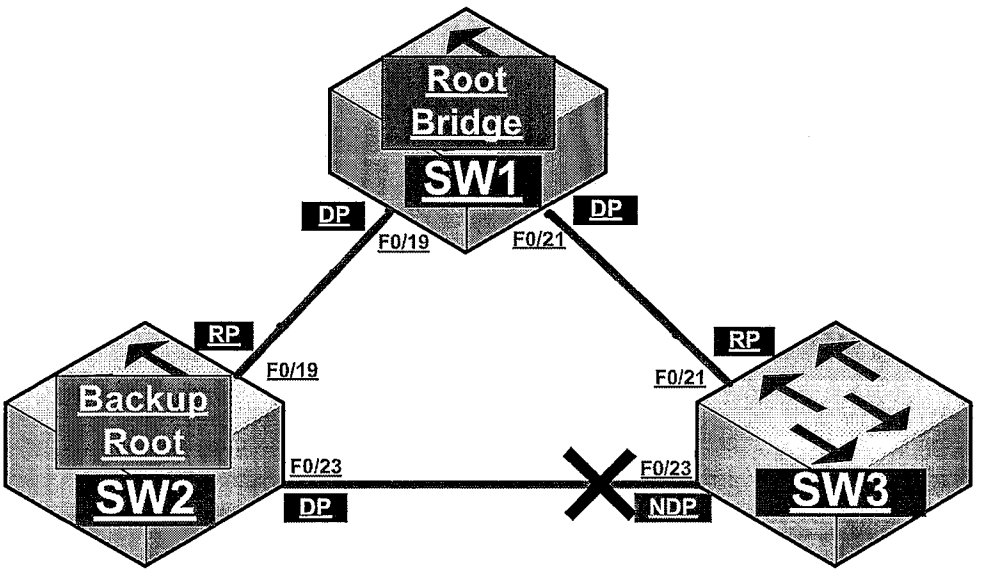
Task 3:

Erase the configuration on all routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 22
Spanning-tree LoopGuard



Task 1

Configure the switches according to the above diagram. The unused ports should be configured in "Shutdown" state.

On All Switches:

```
SW4(config)#Int range f0/1-24
SW4(config-if-range)#Shut
```

On SW1:

```
SW1(config)#Spanning-tree vlan 1 root primary

SW1(config)#Int range f0/19 , f0/21
SW1(config-if-range)#Swi tru enc dot
SW1(config-if-range)#Swi mode trunk
SW1(config-if-range)#No Shut
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW2:

```
SW2(config)#Spanning-tree vlan 1 Root Secondary

SW2(config)#Int range f0/19 , f0/23
SW2(config-if-range)#Swi tru enc dot
SW2(config-if-range)#Swi mode trunk
SW2(config-if-range)#No Shut
```

On SW3:

```
SW3(config)#Int range f0/21 , f0/23
SW3(config-if-range)#Swi tru enc dot
SW3(config-if-range)#Swi mode trunk
SW3(config-if-range)#No shut
```

Task 2

Configure F0/23 interface of SW3 such that if it does NOT receive BPDUs from SW2's F0/23 interface, Layer 2 loops do NOT occur.

The question is that if Spanning-tree protocol is running, why would we ever run into loops?

Well....let's think about the following scenario:

1. **SW1 is configured as the root bridge and SW2 is the backup root bridge.**
2. **SW3 does NOT receive BPDUs from SW2 due to unidirectional link failure.**
3. **The F0/23 interface of SW3 does NOT receive BPDUs from SW2, therefore, it will transition into Spanning-tree Listening and Learning state, and when the max-age timer expires, the F0/23 interface of SW3 will transition into forwarding state.**
4. **This situation creates a Spanning-tree LOOP.**

The Spanning-tree Loop Guard feature can prevent spanning-tree loops from occurring. the blocked port (F0/23) on SW3 will transition into "Loop-Inconsistent" state after the max-age timer expires, avoiding Spanning-tree loops.

To demonstrate the Loop, configure the switches as follows:

On SW2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

NOTE: SW2 is constantly sending BPDUs:

```
SW2#Sh spanning inter f0/23 det | Inc BPDU

    BPDU: sent 227, received 0
```

To stop SW2's F0/23 interface from sending BPDUs, the Spanning-tree BPDUFilter is configured:

```
SW2(config)#int f0/23
SW2(config-if)#Shut
SW2(config-if)#spanning-tree bpdufilter ena
SW2(config-if)#No Shut
```

Once the above command is entered, SW2 will no longer send any BPDUs out of its F0/23 interface, let's verify this fact:

```
SW2#Show spanning inter F0/23 det | Inc BPDU

   BPDU: sent 0, received 0

SW2#Show spanning inter F0/23 det | Inc BPDU

   BPDU: sent 0, received 0

SW2#Show spanning inter F0/23 det | Inc BPDU

   BPDU: sent 0, received 0
```

To test the configuration:

On SW1:

```
SW1#Show spanning-tree vlan 1 | Inc Fa

Fa0/19           Desg FWD 19        128.21   P2p
Fa0/21           Desg FWD 19        128.23   P2p
```

On SW2:

```
SW2#Show spanning-tree vlan 1 | Inc Fa0

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

On SW3:

```
SW3#Show spanning-tree vlan 1 | Inc Fa0

Fa0/21           Root FWD 19        128.21   P2p
Fa0/23           Desg FWD 19        128.23   P2p
```

NOTE: We have a Spanning-tree loop because all ports are in forwarding state.

Let's remove the BPDUFilter command from the F0/23 interface of SW2:

On SW2:

```
SW2(config)#int f0/23
SW2(config-if)#No spanning-tree bpdufilter ena
```

To verify the configuration:

On SW3:

```
SW3#Show spanning-tree vlan 1 | Inc Fa0

Fa0/21           Root FWD 19        128.21   P2p
Fa0/23           Altn BLK 19        128.23   P2p
```

Let's configure the Spanning-tree Loop Guard on SW3's F0/23 interface and enable the "Debug spanning-tree events":

On SW3:

```
SW3#Debug spanning-tree event

SW3(config)#Int F0/23
SW3(config-if)#Shut
SW3(config-if)#Spanning-tree Guard Loop
SW3(config-if)#No Shut
```

To verify the configuration:

On SW3:

```
SW3#Show spanning-tree vlan 1 | Inc Fa0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/21              Root FWD 19        128.21   P2p
Fa0/23              Altn BLK 19        128.23   P2p
```

Let's configure BPDUFilter on the F0/23 interface of SW2:

On SW2:

```
SW2(config)#int f0/23
SW2(config-if)#spanning-tree bpdufilter ena
```

You should see the following console message on SW3 within 30 seconds:

On SW3:

%SPANTREE-2-LOOPGUARD_BLOCK: Loop guard blocking port FastEthernet0/23 on VLAN0001

To verify the configuration:

On SW3:

```
SW3#Show spanning-tree vlan 1 | Inc Fa0

Fa0/21              Root FWD 19        128.21   P2p
Fa0/23              Desg BKN*19        128.23   P2p *LOOP_Inc
```

SW3 places the F0/23 interface in a Loop Inconsistant.

Loop Guard can also be configured in the global configuration mode, once it's configured, the feature protection is applied to all interfaces on the local switch, this can be achieved by using the following command:

```
SWx(config)#Spanning-tree loopguard default
```

Task 3

Erase the configuration on all routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Advanced
CCIE Routing & Switching
v5.0

www.MicronicsTraining.com

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

DMVPN

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

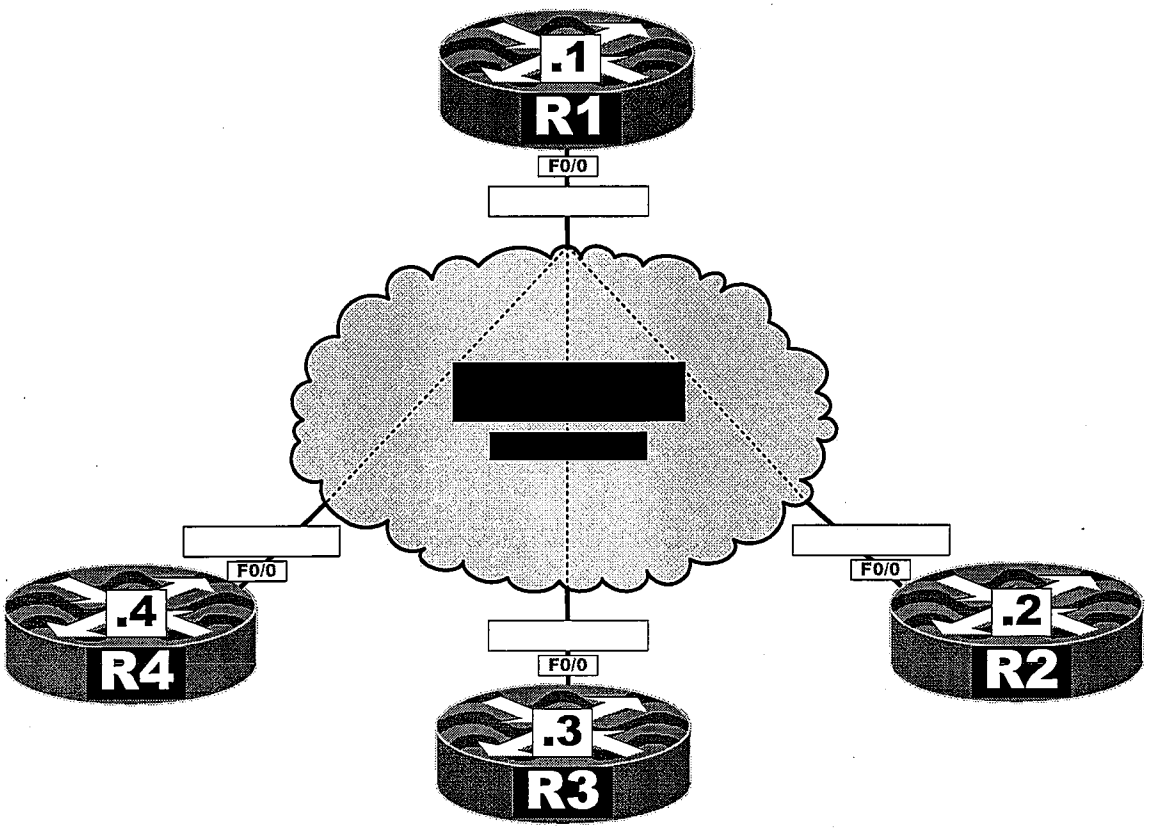
DMVPN:

Lab-1:
Configuring DMVPN Phase - 1

Lab-2:
Configuring DMVPN Phase – 2 & 3

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 1 - DMVPN – Phase #1 with Static Mapping



Task 1

SW1 represents the Internet; configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the F0/0 interfaces of all routers in this topology. The switch interface to which the routers are connected to should have a ".10" in the host portion of the IP address for that subnet.

Let's configure SW1's interfaces for these routers. Since in this lab SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers .

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1(config)#Int range f0/1-4
SW1(config-if-range)#No switchport

SW1(config)#Int F0/1
SW1(config-if)#ip address 192.1.1.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#ip address 192.1.2.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#ip address 192.1.3.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/4
SW1(config-if)#ip address 192.1.4.10 255.255.255.0
SW1(config-if)#No shut
```

Let's NOT forget to enable "IP routing" or else the switch will not be able to route from one subnet to another.

```
SW1(config)#IP routing
```

Let's configure the routers:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 192.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 192.1.1.10
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#ip addr 192.1.2.2 255.255.255.0
R2(config-if)#No shut

R2(config)#ip route 0.0.0.0 0.0.0.0 192.1.2.10
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 192.1.3.3 255.255.255.0
R3(config-if)#No shut

R3(config)#ip route 0.0.0.0 0.0.0.0 192.1.3.10
```

On R4:

```
R4(config)#Int f0/0
R4(config-if)#ip addr 192.1.4.4 255.255.255.0
R4(config-if)#No shut

R4(config)#ip route 0.0.0.0 0.0.0.0 192.1.4.10
```

To verify the configuration:

On R1:

```
R1#Ping 192.1.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R2:

```
R2#Ping 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 2

Configure DMVPN Phase 1 such that R1 is the HUB, and R2, R3, and R4 are configured as the SPOKES. You should use 10.1.1.x /24, where "x" is the router number. If this configuration is performed correctly, these routers should have reachability to all tunnel end points. You should configure static mapping to accomplish this task.

DMVPN:

DMVPN is a combination of mGRE and NHRP (Next Hop Resolution Protocol) and IPsec (Optional). DMVPN can be implemented as Phase 1, Phase 2, or Phase 3.

There are two GRE flavors:

- **GRE**
- **mGRE**

GRE which is a point-to-point logical link is configured with a Tunnel source, Tunnel destination, and Tunnel encapsulation. When Tunnel destination is configured, it ties the Tunnel to a specific end point which makes these tunnels a point-to-point tunnel, this means that if there are 200 endpoints, each endpoint needs to configure 199 GRE Tunnels.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

With "mGRE" (Multipoint Generic Routing Encapsulation) the configuration includes the Tunnel source, and Tunnel mode, the tunnel destination is NOT configured, therefore, the tunnel can have any or many endpoints and only a single tunnel interface is utilized. The endpoints can be configured as GRE, or mGRE.

But what if the spokes need to communicate with each other especially with the NBMA nature of mGRE? How would we accomplish that? In a hub and spoke Frame-Relay, if a spoke needs to communicate with another spoke, a Frame-Relay mapping needs to be configured, is there a mapping that we need to configure in mGRE?

Well, mGRE does not have that capability and this is why another protocol is incorporated, it's called "NHRP", which stands for Next Hop Resolution Protocol.

NHRP:
NHRP is defined in RFC 2332, provides a layer two address resolution protocol and caching services, very much like ARP or an Inverse-arp.
NHRP is used by the spokes connected to an NBMA network to determine the NBMA IP address of the next-hop router. With NHRP we can map a tunnel IP address to an NBMA IP address either statically or dynamically.

The NBMA IP address in this scenario is the IP address that was acquired from the service provider, the Tunnel IP address is the IP address that WE assigned to the Tunnel interface, typically an RFC 1918 addressing.

In NHRP, the routers are configured as NHC (NHRP Client/s) or NHS (The NHRP Server). The NHS acts as a mapping agent and stores all registered mappings performed by the NHC/s so it can reply to the queries made by NHC/s. NHCs send a query to the NHS if they need to communicate with another NHC.

NHRP is like ARP protocol, why is it like ARP protocol?
Because it allows NHCs to dynamically register their NBMA to Tunnel IP addresses, this allows the NHCs to join the NBMA network without having to configure and reconfigure the NHS. This means that when a new NHC is added to the NBMA network, none of the NHCs or the NHS/es need to be configured.

Let's look at a scenario where the NHC/s have a dynamic physical IP address, or the NHC is behind a NAT device. Now, how would you configure the NHS and what IP are you going to use for the NHCs?
This is the reason that dynamic registration and queries are very useful, because it is almost impossible to preconfigure the logical VPN-IP to the physical NBMA-IP mapping for the NHCs on the NHS. Therefore, NHRP is a resolution protocol that allows the NHCs to dynamically discover the logical-IP to physical-IP mapping for other NHCs within the same NBMA network.

Without this discovery, packets must traverse through the hub to reach other spokes, this can negatively impact the CPU and the bandwidth consumption of the hub router.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

There are three phases in DMVPN configuration, Phase 1, 2 and 3.

Important Points to remember on DMVPN Phase – 1:

- mGRE is configured on the Hub,and GRE is configured on the Spokes.
- Multicast or unicast traffic can ONLY flow between the hub and the spokes and NOT spoke to spoke.
- This can be configured statically or have the NHCs (Spokes) register themselves dynamically with the NHS.

Let's configure R1 (The hub router) with static mappings:

The tunnel configuration, whether static or dynamic, can be broken down into two configuration phases; in the first phase the mGRE configuration is completed, this includes three commands: the IP address of the tunnel, the Tunnel source, and the Tunnel mode:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#IP address 10.1.1.1 255.255.255.0
R1(config-if)#Tunnel source 192.1.1.1
R1(config-if)#Tunnel mode gre multipoint
```

In the second phase of our configuration, the NHRP is configured, this configuration includes three NHRP commands: The NHRP network-id which enables NHRP on that tunnel interface, NHRP mapping that maps the Tunnel IP address of the spoke/s to the physical IP (NBMA-IP) address of the spoke/s, this needs to be done for each spoke, and an optional configuration of NHRP mapping of multicast to the physical IP address of the spokes which enables Multicasting and allows the IGPs that use Multicasting over the tunnel interface (Does this remind you of the Frame-Relay days "Broadcast" keyword at the end of the frame-relay map statement?). In this task the mapping of Multicast to the NBMA-IP is not configured because the task did not ask for it.

```
R1(config-if)#IP NHRP Network-id 111
R1(config-if)#IP NHRP map 10.1.1.2 192.1.2.2
R1(config-if)#IP NHRP map 10.1.1.3 192.1.3.3
R1(config-if)#IP NHRP map 10.1.1.4 192.1.4.4
```

To verify the configuration:

```
R1#Show ip nhrp

10.1.1.2/32 via 10.1.1.2
   Tunnel1 created 00:05:20, never expire
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   Type: static, Flags:
   NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
   Tunnel1 created 00:05:12, never expire
   Type: static, Flags:
   NBMA address: 192.1.3.3
10.1.1.4/32 via 10.1.1.4
   Tunnel1 created 00:05:05, never expire
   Type: static, Flags:
   NBMA address: 192.1.4.4
```

On R2:

Since in DMVPN phase #1 configuration the spoke routers should be configured as point-to-point, the configuration includes the tunnel source and the tunnel destination, and because the tunnel destination is

configured, it ties that tunnel to that destination only, which makes the tunnel a point-to-point tunnel and NOT a multipoint tunnel. Once the tunnel commands are configured, the next step or the last step is to configure "NHRP", in this configuration, NHRP is enabled first, and then a single mapping is configured for the hub's tunnel IP address:

```
R2(config)#Int tunnel 1
R2(config-if)#IP addr 10.1.1.2 255.255.255.0
R2(config-if)#Tunnel source 192.1.2.2
R2(config-if)#Tunnel destination 192.1.1.1
R2(config-if)#IP nhrp network-id 222
R2(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
```

To verify the configuration:

```
R2#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:04:03, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

On R3:

```
R3(config)#Int tunnel 1
R3(config-if)#IP addr 10.1.1.3 255.255.255.0
R3(config-if)#Tunnel source F0/0
R3(config-if)#Tunnel destination 192.1.1.1
R3(config-if)#IP nhrp network-id 333
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
```

On R4:

```
R4(config)#Int tunnel 1
R4(config-if)#IP addr 10.1.1.4 255.255.255.0
R4(config-if)#Tunnel source F0/0
R4(config-if)#Tunnel destination 192.1.1.1
R4(config-if)#IP nhrp network-id 444
R4(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
```

To test the configuration:

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 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R1#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R1#Ping 10.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

On R2:

```
R2#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
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

```
R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R2#Ping 10.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

To see the traffic path between the spokes:

```
R2#Traceroute 10.1.1.3

Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.1 4 msec 4 msec 4 msec
  2 10.1.1.3 0 msec *  0 msec

R2#Traceroute 10.1.1.4

Type escape sequence to abort.
Tracing the route to 10.1.1.4
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.1 4 msec 4 msec 0 msec
  2 10.1.1.4 4 msec *  0 msec
```

On R3:

```
R3#Ping 10.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Traceroute 10.1.1.4

Type escape sequence to abort.
Tracing the route to 10.1.1.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.1 0 msec 4 msec 4 msec
  2 10.1.1.4 0 msec *  0 msec
```

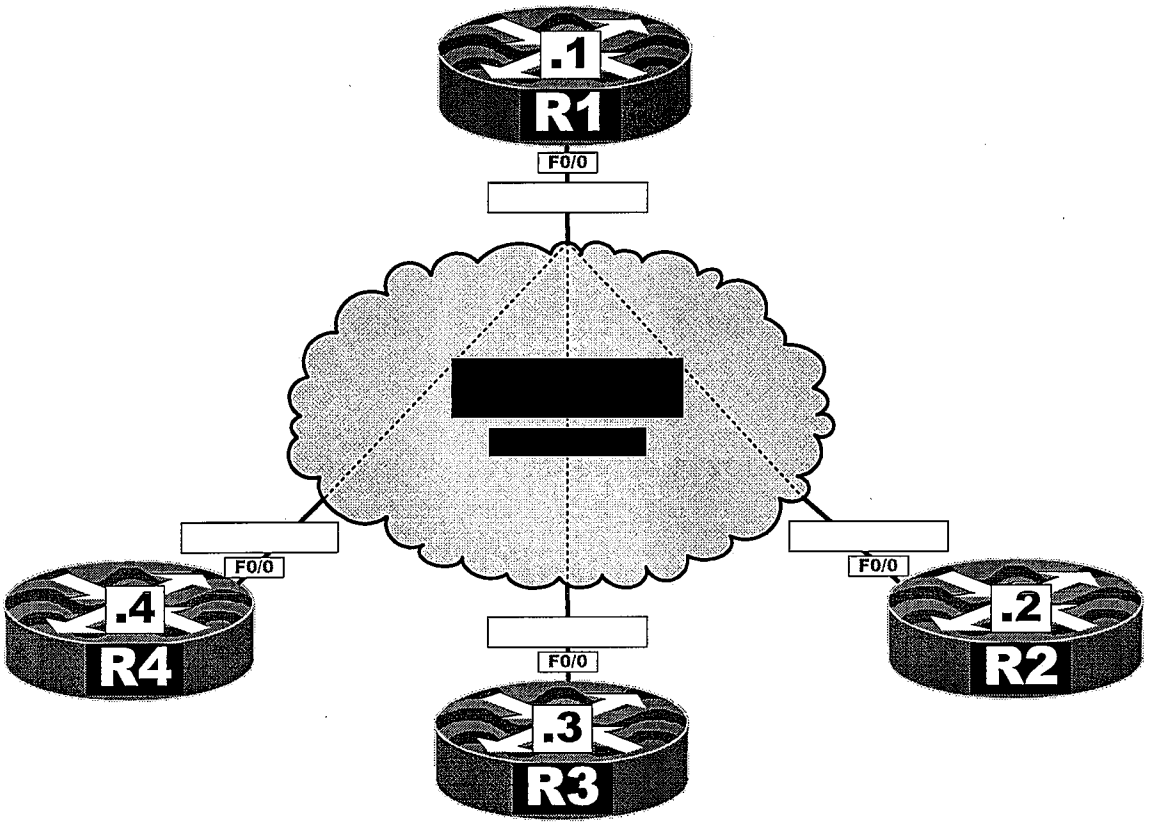
Since the spokes are configured in a point-to-point manner, there is no need to map Multicast traffic to the NBMA-IP of a given endpoint.

Task 3

Erase the startup configuration of the routers and the switch and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 2 - DMVPN – Phase #1 with Dynamic Mapping



Task 1

SW1 represents the Internet; configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the F0/0 interfaces of all routers in this topology. The switch interface to which the routers are connected to should have a ".10" in the host portion of the IP address for that subnet.

Let's configure SW1's interfaces for these routers. Since in this lab SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers .

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1(config)#Int range f0/1-4
SW1(config-if-range)#No switchport

SW1(config)#Int F0/1
SW1(config-if)#ip address 192.1.1.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#ip address 192.1.2.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#ip address 192.1.3.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/4
SW1(config-if)#ip address 192.1.4.10 255.255.255.0
SW1(config-if)#No shut
```

Let's NOT forget to enable "IP routing" or else the switch will not be able to route from one subnet to another.

```
SW1(config)#IP routing
```

Let's configure the routers:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 192.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 192.1.1.10
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#ip addr 192.1.2.2 255.255.255.0
R2(config-if)#No shut

R2(config)#ip route 0.0.0.0 0.0.0.0 192.1.2.10
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 192.1.3.3 255.255.255.0
R3(config-if)#No shut

R3(config)#ip route 0.0.0.0 0.0.0.0 192.1.3.10
```

On R4:

```
R4(config)#Int f0/0
R4(config-if)#ip addr 192.1.4.4 255.255.255.0
R4(config-if)#No shut

R4(config)#ip route 0.0.0.0 0.0.0.0 192.1.4.10
```

To verify the configuration:

On R1:

```
R1#Ping 192.1.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R2:

```
R2#Ping 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 2

Configure DMVPN Phase 1 such that R1 is the HUB, and R2, R3, and R4 are configured as the SPOKES. You should use 10.1.1.x /24, where "x" is the router number. If this configuration is performed correctly, these routers should have reachability to all tunnel end points. You should NOT configure static mappings on the hub router to accomplish this task.

Since the task refers to DMVPN Phase #1, we know that the hub router must be configured in a multipoint manner and the spokes should be configured in a point-to-point manner. The task also specifies that the hub router should NOT be configured with static mappings; this means that the hub should be configured such that the mappings are performed dynamically. This configuration is more scalable than configuring the hub with static mappings. If the hub is configured with static mappings for each spoke, a newly added spoke will require some configuration on the hub for that given spoke, whereas, if the mappings are performed dynamically, the spokes can be added dynamically without any configuration performed on the hub.
How does that work?

When a spoke router initially connects to the DMVPN network, it registers its Tunnel-IP to its NBMA-IP mapping with the hub router. This registeration enables the mGRE interface on the hub router to build a dynamic GRE tunnel back to the registering spoke router, this means that the spoke router/s MUST know

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

and be configured with the tunnel IP address of the hub and a mapping of that tunnel-IP to the NBMA-IP of the hub; besides this configuration, the spokes should be configured with the tunnel IP address of the NHS (NHRP-Server), or else they won't know where to go to register their Tunnel-IP to NBMA-IP mapping.

Let's configure the routers starting with the hub (R1). Remember that this configuration should be done in two steps, in the first step, the tunnel source and tunnel mode is configured:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#IP address 10.1.1.1 255.255.255.0
R1(config-if)#Tunnel source 192.1.1.1
R1(config-if)#Tunnel mode gre multipoint
```

In the second phase of our configuration, the NHRP is configured, this configuration includes three NHRP commands: The NHRP network-id which enables NHRP on that tunnel interface, then, NHRP mapping should be configured, this mapping can be configured dynamically or statically, the previous lab demonstrated the static mapping and in this lab the dynamic mapping is demonstrated, in dynamic mapping configuration nothing else is done, as long as NHRP is enabled on that interface the mappings will occur dynamically, let's configure the routers and verify:

```
R1(config-if)#IP NHRP Network-id 111
```

On R2:

Since in DMVPN phase #1 configuration the spoke routers should be configured as point-to-point, the configuration includes the tunnel source and the tunnel destination, and because the tunnel destination is configured, it ties that tunnel to that destination only, which makes the tunnel a point-to-point tunnel and NOT a multipoint tunnel. Once the tunnel commands are configured, the next step or the last step is to configure "NHRP", in this configuration, NHRP is enabled first, and then a single mapping is configured for the hub's tunnel IP address:

```
R2(config)#Int tunnel 1
R2(config-if)#IP addr 10.1.1.2 255.255.255.0
R2(config-if)#Tunnel source F0/0
R2(config-if)#Tunnel destination 192.1.1.1
R2(config-if)#IP nhrp network-id 222
R2(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R2(config-if)#ip nhrp nhs 10.1.1.1
```

To verify the configuration:

```
R2#Show ip nhrp
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:04:03, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

On R1:

```
R1#Show ip nhrp

10.1.1.2/32 via 10.1.1.2
   Tunnel1 created 00:01:05, expire 01:58:54
   Type: dynamic, Flags: unique registered
   NBMA address: 192.1.2.2
```

We can see that once R2 was connected to the DMVPN network, the hub router, R1, created a dynamic entry for R2. Let's configure the other spoke routers:

On R3:

```
R3(config)#Int tunnel 1
R3(config-if)#IP addr 10.1.1.3 255.255.255.0
R3(config-if)#Tunnel source F0/0
R3(config-if)#Tunnel destination 192.1.1.1
R3(config-if)#IP nhrp network-id 333
R3(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R3(config-if)#ip nhrp nhs 10.1.1.1
```

On R4:

```
R4(config)#Int tunnel 1
R4(config-if)#IP addr 10.1.1.4 255.255.255.0
R4(config-if)#Tunnel source F0/0
R4(config-if)#Tunnel destination 192.1.1.1
R4(config-if)#IP nhrp network-id 444
R4(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R4(config-if)#ip nhrp nhs 10.1.1.1
```

To verify the configuration:

On R1:

```
R1#Show ip nhrp

10.1.1.2/32 via 10.1.1.2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   Tunnel1 created 00:05:53, expire 01:54:06
   Type: dynamic, Flags: unique registered
   NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
   Tunnel1 created 00:02:24, expire 01:57:35
   Type: dynamic, Flags: unique registered
   NBMA address: 192.1.3.3
10.1.1.4/32 via 10.1.1.4
   Tunnel1 created 00:00:21, expire 01:59:38
   Type: dynamic, Flags: unique registered
   NBMA address: 192.1.4.4
```

To test the configuration:

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 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R1#Ping 10.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 10.1.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

On R3:

```
R3#Ping 10.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R3#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:
```


© 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

R3#Ping 10.1.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Traceroute 10.1.1.2

Type escape sequence to abort.
Tracing the route to 10.1.1.2
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.1 0 msec 4 msec 4 msec
  2 10.1.1.2 4 msec *  0 msec
```

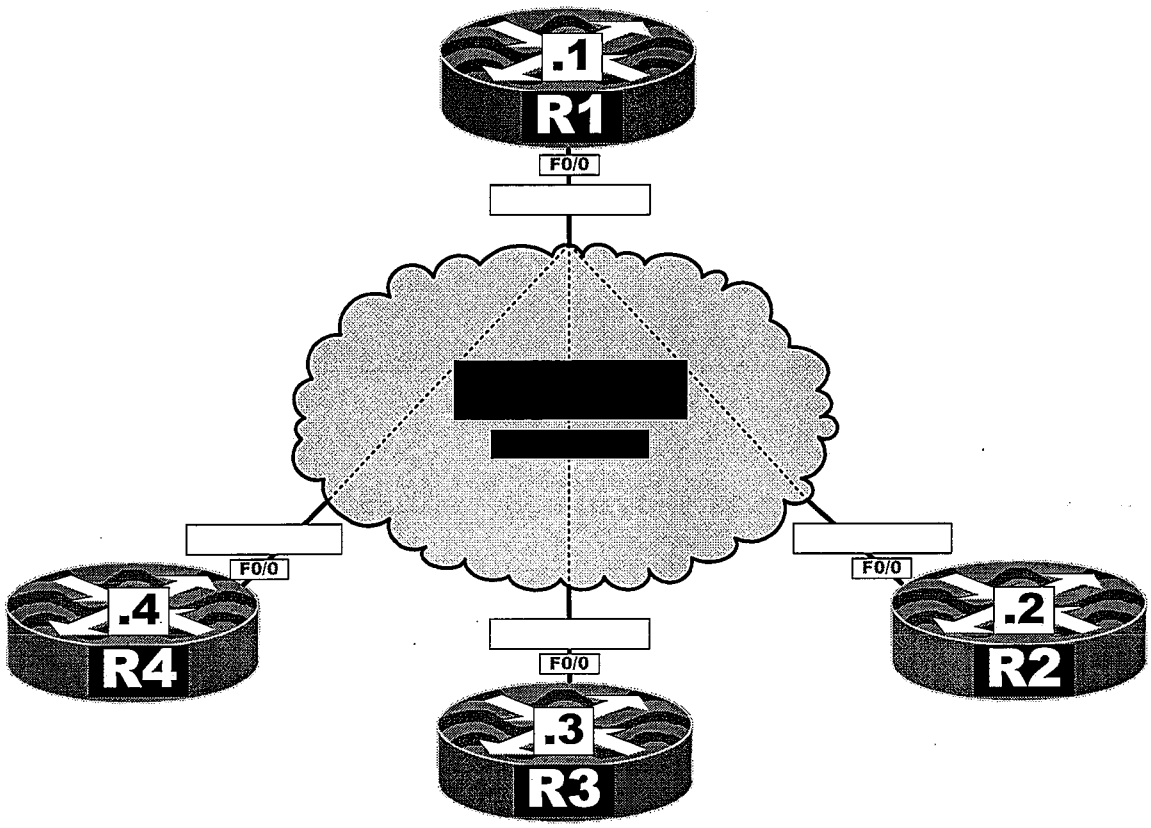
We can see that in DMVPN phase one, no matter how it is configured (Static mapping or Dynamic mapping) the spokes will go through the hub to connect to the other spokes.

Task 3

Erase the startup configuration of the routers and the switch and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 3 – DMVPN – Phase #2 Using Static Mappings



Task 1

SW1 represents the Internet; configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the F0/0 interfaces of all routers in this topology. The switch interface to which the routers are connected to should have a ".10" in the host portion of the IP address for that subnet.

Let's configure SW1's interfaces for these routers. Since in this lab SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers .

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1(config)#Int range f0/1-4
SW1(config-if-range)#No switchport

SW1(config)#Int F0/1
SW1(config-if)#ip address 192.1.1.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#ip address 192.1.2.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#ip address 192.1.3.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/4
SW1(config-if)#ip address 192.1.4.10 255.255.255.0
SW1(config-if)#No shut
```

Let's NOT forget to enable "IP routing" or else the switch will not be able to route from one subnet to another.

```
SW1(config)#IP routing
```

Let's configure the routers:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 192.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 192.1.1.10
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#ip addr 192.1.2.2 255.255.255.0
R2(config-if)#No shut

R2(config)#ip route 0.0.0.0 0.0.0.0 192.1.2.10
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 192.1.3.3 255.255.255.0
R3(config-if)#No shut

R3(config)#ip route 0.0.0.0 0.0.0.0 192.1.3.10
```

On R4:

```
R4(config)#Int f0/0
R4(config-if)#ip addr 192.1.4.4 255.255.255.0
R4(config-if)#No shut

R4(config)#ip route 0.0.0.0 0.0.0.0 192.1.4.10
```

To verify the configuration:

On R1:

```
R1#Ping 192.1.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R2:

```
R2#Ping 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 2

Configure DMVPN Phase 2 such that R1 is the HUB, and R2, R3, and R4 are configured as the SPOKES. You should use 10.1.1.x /24, where "x" is the router number. If this configuration is performed correctly, these routers should have reachability to all tunnel end points. You should configure static mappings on the hub router to accomplish this task.

Since the task refers to DMVPN Phase #2, the hub and all the spokes are configured in multipoint GRE, in this phase the spokes have reachability to the other spokes directly.

Let's configure the routers and verify:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#IP address 10.1.1.1 255.255.255.0
R1(config-if)#Tunnel source 192.1.1.1
R1(config-if)#Tunnel mode gre multipoint

R1(config-if)#ip nhrp Network-id 111
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-if)#ip nhrp map 10.1.1.2 192.1.2.2
R4(config-if)#ip nhrp map 10.1.1.3 192.1.3.3
R4(config-if)#ip nhrp map 10.1.1.4 192.1.4.4
```

On R2:

```
R2(config)#Int tunnel 1
R2(config-if)#IP addr 10.1.1.2 255.255.255.0
R2(config-if)#Tunnel source F0/0
R2(config-if)#Tunnel mode gre multipoint

R2(config-if)#ip nhrp network-id 222
R2(config-if)#ip nhrp map 10.1.1.1 192.1.1.1
R2(config-if)#ip nhrp map 10.1.1.3 192.1.3.3
R2(config-if)#ip nhrp map 10.1.1.4 192.1.4.4
```

On R3:

```
R3(config)#Int tunnel 1
R3(config-if)#IP addr 10.1.1.3 255.255.255.0
R3(config-if)#Tunnel source F0/0
R3(config-if)#Tunnel mode gre multipoint

R3(config-if)#IP nhrp network-id 333
R3(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R3(config-if)#ip nhrp map 10.1.1.2 192.1.2.2
R3(config-if)#ip nhrp map 10.1.1.4 192.1.4.4
```

On R4:

```
R4(config)#Int tunnel 1
R4(config-if)#IP addr 10.1.1.4 255.255.255.0
R4(config-if)#Tunnel source F0/0
R4(config-if)#Tunnel mode gre multipoint

R4(config-if)#IP nhrp network-id 444
R4(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R4(config-if)#ip nhrp map 10.1.1.2 192.1.2.2
R4(config-if)#ip nhrp map 10.1.1.3 192.1.3.3
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip nhrp

10.1.1.2/32 via 10.1.1.2
   Tunnel1 created 00:19:16, never expire
   Type: static, Flags: used
   NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
   Tunnel1 created 00:19:08, never expire
   Type: static, Flags:
   NBMA address: 192.1.3.3
10.1.1.4/32 via 10.1.1.4
   Tunnel1 created 00:01:09, never expire
   Type: static, Flags:
   NBMA address: 192.1.4.4
```

On R2:

```
R2#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:15:34, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.3/32 via 10.1.1.3
   Tunnel1 created 00:15:27, never expire
   Type: static, Flags:
   NBMA address: 192.1.3.3
10.1.1.4/32 via 10.1.1.4
   Tunnel1 created 00:15:19, never expire
   Type: static, Flags:
   NBMA address: 192.1.4.4
```

On R3:

```
R3#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:15:19, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1 created 00:15:11, never expire
   Type: static, Flags:
   NBMA address: 192.1.2.2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
10.1.1.4/32 via 10.1.1.4
   Tunnel1 created 00:15:02, never expire
   Type: static, Flags:
   NBMA address: 192.1.4.4
```

On R4:

```
R4#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:15:13, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1 created 00:15:04, never expire
   Type: static, Flags:
   NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
   Tunnel1 created 00:14:57, never expire
   Type: static, Flags:
   NBMA address: 192.1.3.3
```

To test the configuration:

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 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R1#Ping 10.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 10.1.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
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

On R3:

```
R3#Ping 10.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R3#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 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Ping 10.1.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Traceroute 10.1.1.2

Type escape sequence to abort.
Tracing the route to 10.1.1.2
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.2 4 msec *  0 msec
```

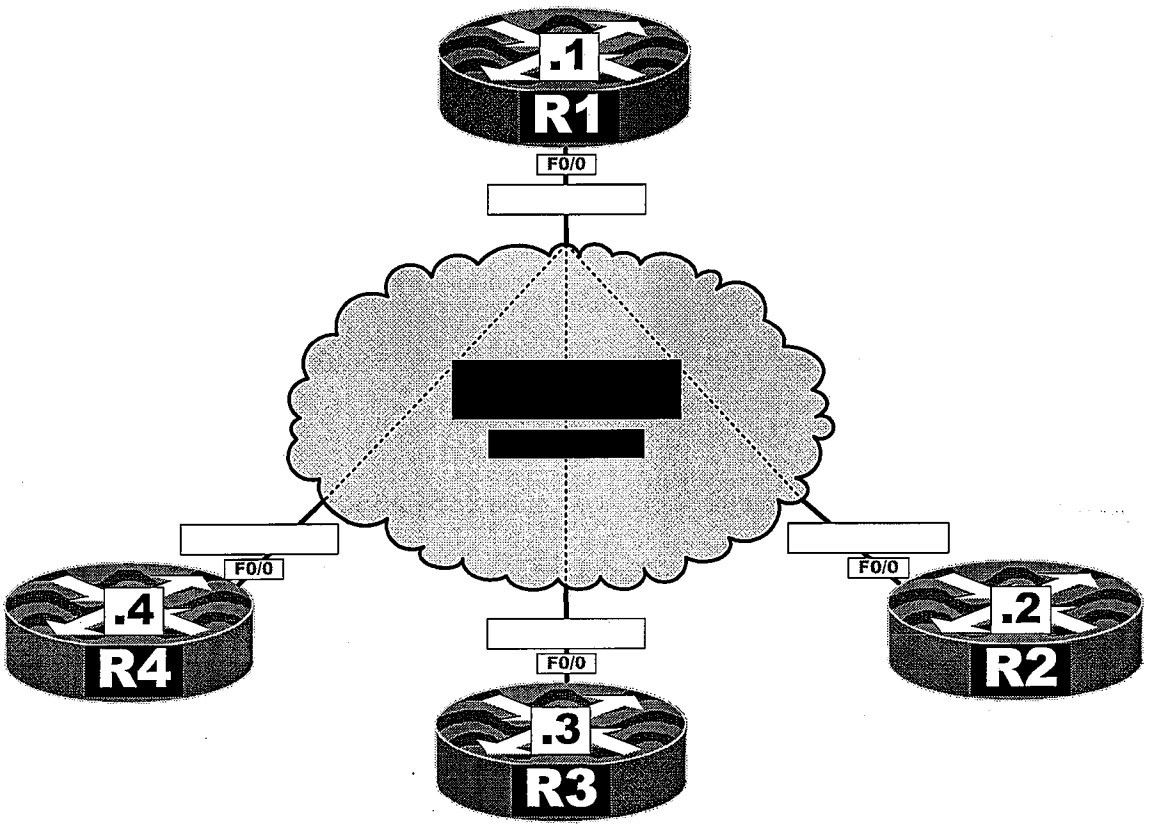
We can see that in DMVPN phase two, the spokes can communicate and reach other spokes directly.

Task 3

Erase the startup configuration of the routers and the switch and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 4 – DMVPN – Phase #2 Using Dynamic Mappings



Task 1

SW1 represents the Internet; configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the F0/0 interfaces of all routers in this topology. The switch interface to which the routers are connected to should have a ".10" in the host portion of the IP address for that subnet.

Let's configure SW1's interfaces for these routers. Since in this lab SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers .

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1(config)#Int range f0/1-4
SW1(config-if-range)#No switchport

SW1(config)#Int F0/1
SW1(config-if)#ip address 192.1.1.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#ip address 192.1.2.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#ip address 192.1.3.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/4
SW1(config-if)#ip address 192.1.4.10 255.255.255.0
SW1(config-if)#No shut
```

Let's NOT forget to enable "IP routing" or else the switch will not be able to route from one subnet to another.

```
SW1(config)#IP routing
```

Let's configure the routers:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 192.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 192.1.1.10
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#ip addr 192.1.2.2 255.255.255.0
R2(config-if)#No shut

R2(config)#ip route 0.0.0.0 0.0.0.0 192.1.2.10
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 192.1.3.3 255.255.255.0
R3(config-if)#No shut

R3(config)#ip route 0.0.0.0 0.0.0.0 192.1.3.10
```

On R4:

```
R4(config)#Int f0/0
R4(config-if)#ip addr 192.1.4.4 255.255.255.0
R4(config-if)#No shut

R4(config)#ip route 0.0.0.0 0.0.0.0 192.1.4.10
```

To verify the configuration:

On R1:

```
R1#Ping 192.1.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R2:

```
R2#Ping 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 192.1.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 2

Configure DMVPN Phase 2 such that R1 is the HUB, and R2, R3, and R4 are configured as the SPOKES. You should use 10.1.1.x /24, where "x" is the router number. If this configuration is performed correctly, these routers should have reachability to all tunnel end points. You should NOT configure static mappings on the hub router to accomplish this task.

Since the task refers to DMVPN Phase #2, the hub and all the spokes are configured in multipoint GRE, in this phase the spokes have reachability to the other spokes directly.

NOTE: The hub router is configured with tunnel source, tunnel mode, and ip nhrp network-id which enables NHRP on this tunnel interface:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#IP address 10.1.1.1 255.255.255.0
R1(config-if)#Tunnel source 192.1.1.1
R1(config-if)#Tunnel mode gre multipoint
R1(config-if)#ip nhrp Network-id 111
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The spokes are configured with tunnel source, tunnel mode gre multipoint, IP nhrp network-id which enables NHRP, ip nhrp map 10.1.1.1 192.1.1.1 which maps the hub's tunnel IP to the hub's NBMA-IP and the last command identifies the "NHS" which is the tunnel IP address of the hub router.

On R2:

```
R2(config)#Int tunnel 1
R2(config-if)#IP addr 10.1.1.2 255.255.255.0
R2(config-if)#Tunnel source F0/0
R2(config-if)#Tunnel mode gre multipoint

R2(config-if)#ip nhrp network-id 222
R2(config-if)#ip nhrp map 10.1.1.1 192.1.1.1
R2(config-if)#ip nhrp nhs 10.1.1.1
```

Before we configure R3, let's enable "Debug nhrp packet" and "Debug nhrp cache" on both R1 and R3 and verify the output:

On R3:

```
R3(config)#Int tunnel 1
R3(config-if)#IP addr 10.1.1.3 255.255.255.0
R3(config-if)#Tunnel source F0/0
R3(config-if)#Tunnel mode gre multipoint

R3(config-if)#IP nhrp network-id 333
R3(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R3(config-if)#ip nhrp nhs 10.1.1.1
```

The following debug output was produced once the "IP nhrp map 10.1.1.1 192.1.1.1" command was configured:

```
NHRP: Tunnel1: Cache add for target 10.1.1.1/32 next-hop 10.1.1.1
           192.1.1.1
```

The following debug output sends a registration request to the NHS, this happened because we configured the NHS using the "IP nhrp nhs 10.1.1.1" command:

```
NHRP: Send Registration Request via Tunnel1 vrf 0, packet size: 92
 src: 10.1.1.3, dst: 10.1.1.1
```

On R1:

R1 received the registration request and added the entry to its cache:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
NHRP: Receive Registration Request via Tunnel1 vrf 0, packet size: 92
 (F) afn: IPv4(1), type: IP(800), hop: 255, ver: 1
     shtl: 4(NSAP), sstl: 0(NSAP)
     pktsz: 92 extoff: 52
 (M) flags: "unique nat ", reqid: 65537
     src NBMA: 192.1.2.2
     src protocol: 10.1.1.2, dst protocol: 10.1.1.1

NHRP: Tunnel1: Cache update for target 10.1.1.2/32 next-hop 10.1.1.2
           192.1.2.2
```

Let's configure R4:

On R4:

```
R4(config)#Int tunnel 1
R4(config-if)#IP addr 10.1.1.4 255.255.255.0
R4(config-if)#Tunnel source F0/0
R4(config-if)#Tunnel mode gre multipoint

R4(config-if)#IP nhrp network-id 444
R4(config-if)#IP nhrp map 10.1.1.1 192.1.1.1
R4(config-if)#ip nhrp nhs 10.1.1.1
```

To verify and test the configuration:

Let's see the existing NHRP cache:

```
R4#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:00:24, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
```

The following traceroute should go through the hub router to reach 10.1.1.2, since this is the initial communication between the two spokes; R4 needs to send a resolution request to the HUB router, R1. Once R4 gets the resolution reply back, it can go directly to R2 for subsequent traffic. As you can see the second traceroute was sent directly to R2.

```
R4#Traceroute 10.1.1.2

Type escape sequence to abort.
Tracing the route to 10.1.1.2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.1 8 msec 0 msec 4 msec
  2 10.1.1.2 8 msec *  0 msec

R4#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1 created 00:00:44, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1 created 00:00:07, expire 01:59:52
   Type: dynamic, Flags: router used
   NBMA address: 192.1.2.2
10.1.1.4/32 via 10.1.1.4
   Tunnel1 created 00:00:07, expire 01:59:52
   Type: dynamic, Flags: router unique local
   NBMA address: 192.1.4.4
    (no-socket)
```

We can see that in DMVPN phase two, the spokes can communicate and reach other spokes directly.

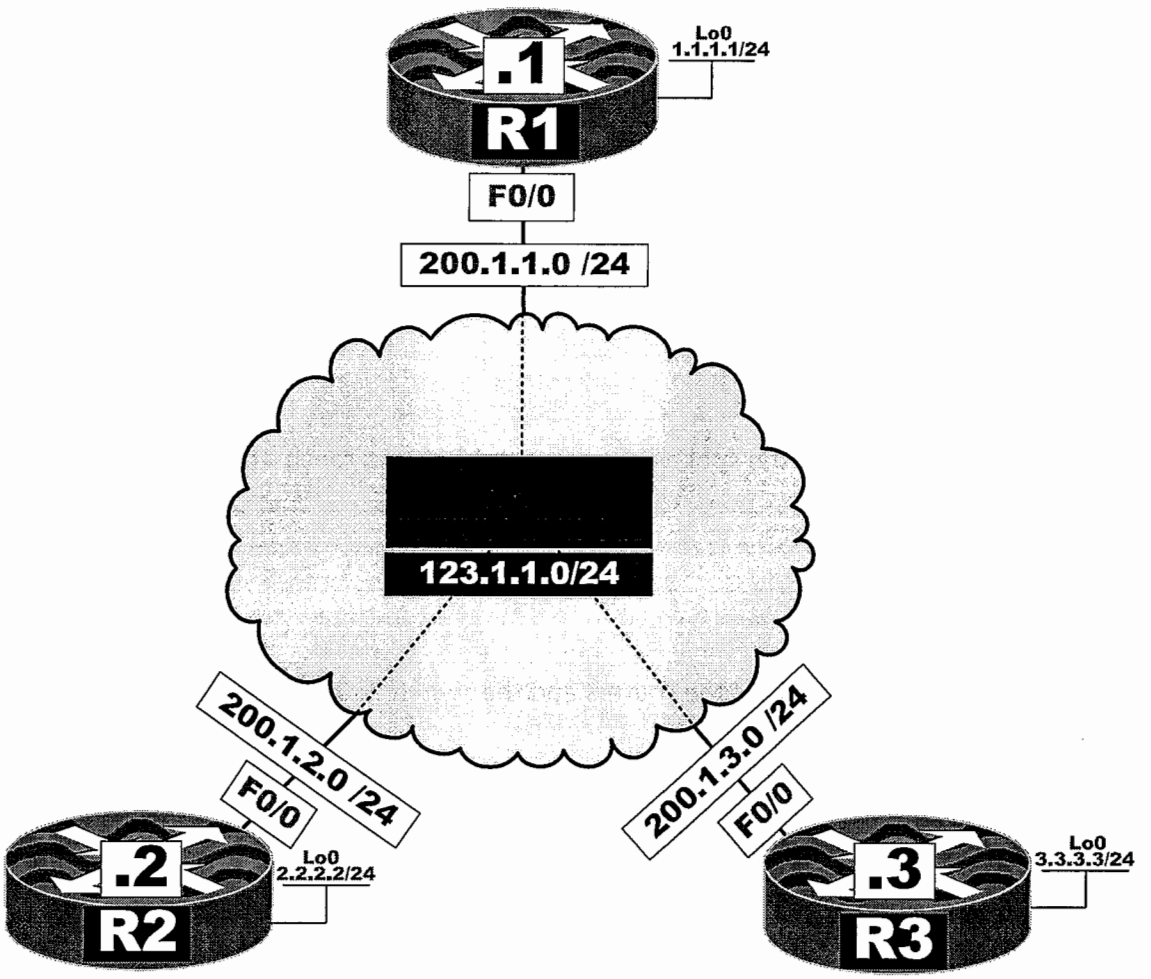
Task 3

Erase the startup configuration of the routers and the switch and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 5 – Running Routing Protocols on DMVPN Phase #1



Task 1

SW1 represents the Internet; configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these routers should be able to ping and have reachability to the F0/0 interfaces of all routers

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

in this topology. The switch interface to which the routers are connected to should have a ".10" in the host portion of the IP address for that subnet.

Let's configure SW1's interfaces for these routers. Since in this lab SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers.

On SW1:

```
SW1(config)#Int range f0/1-3
SW1(config-if-range)#No switchport

SW1(config)#Int F0/1
SW1(config-if)#ip address 200.1.1.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#ip address 200.1.2.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#ip address 200.1.3.10 255.255.255.0
SW1(config-if)#No shut
```

Let's NOT forget to enable "IP routing" or else the switch will not be able to route from one subnet to another.

```
SW1(config)#IP routing
```

Let's configure the routers:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 200.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 200.1.1.10
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#ip addr 200.1.2.2 255.255.255.0
R2(config-if)#No shut
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#ip route 0.0.0.0 0.0.0.0 200.1.2.10
```

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 200.1.3.3 255.255.255.0
R3(config-if)#No shut

R3(config)#ip route 0.0.0.0 0.0.0.0 200.1.3.10
```

To verify the configuration:

On R1:

```
R1#Ping 200.1.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R2:

```
R2#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 200.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.3.3, timeout is 2 seconds:
!!!!!
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

Task 2

Configure DMVPN Phase 1 such that R1 is the HUB, and R2 and R3 are configured as the Spokes. You should use 123.1.1.x /24, where "x" is the router number. If this configuration is performed correctly, these routers should have reachability to the tunnel end points. You should configure static mapping to accomplish this task.

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#IP address 123.1.1.1 255.255.255.0
R1(config-if)#Tunnel source 200.1.1.1
R1(config-if)#Tunnel mode gre multipoint
R1(config-if)#IP NHRP Network-id 111
R1(config-if)#IP NHRP map 123.1.1.2 200.1.2.2
R1(config-if)#IP NHRP map 123.1.1.3 200.1.3.3
```

The following command provides Multicast capabilities so the routing protocols can function using their reserved Multicast address:

```
R1(config-if)#IP NHRP map multicast 200.1.2.2
R1(config-if)#IP NHRP map multicast 200.1.3.3
```

On R2:

```
R2(config)#Int Tunnel 1
R2(config-if)#IP address 123.1.1.2 255.255.255.0
R2(config-if)#Tunnel source 200.1.2.2
R2(config-if)#Tunnel destination 200.1.1.1
R2(config-if)#IP NHRP network-id 222
R2(config-if)#IP NHRP map 123.1.1.1 200.1.1.1
```

NOTE: The IP NHRP Network-id does NOT have to match the other routers connected to the same NBMA network. The "IP NHRP map" statement maps the Tunnel-IP to the NBMA-IP of the hub router. Since the spokes are configured in a point-to-point manner and the point-to-point networks have the capability to have Unicast and/or Multicast traffic, there is no need to configure the "ip nhrp map multicast" command on the spokes.

On R3:

```
R3(config)#Int Tunnel 1
R3(config-if)#IP address 123.1.1.3 255.255.255.0
R3(config-if)#Tunnel source 200.1.3.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-if)#Tunnel destination 200.1.1.1
R3(config-if)#IP NHRP network-id 333
R3(config-if)#IP NHRP map 123.1.1.1 200.1.1.1
```

To verify the configuration:

On R1:

```
R1#Show ip nhrp detail

123.1.1.2/32 via 123.1.1.2
    Tunnel1 created 03:41:27, never expire
    Type: static, Flags:
    NBMA address: 200.1.2.2
123.1.1.3/32 via 123.1.1.3
    Tunnel1 created 03:41:27, never expire
    Type: static, Flags:
    NBMA address: 200.1.3.3
```

The output of the above show command displays the NHRP mapping, it reveals the mapping of Tunnel IP to NBMA IP for each spoke or NHC.

On R2:

```
R2#Show ip nhrp detail

123.1.1.1/32 via 123.1.1.1
    Tunnel1 created 00:08:27, never expire
    Type: static, Flags:
    NBMA address: 200.1.1.1
```

NOTE: The spokes have a single mapping to the hub.

On R3:

```
R3#Show ip nhrp detail

123.1.1.1/32 via 123.1.1.1
    Tunnel1 created 00:06:12, never expire
    Type: static, Flags:
    NBMA address: 200.1.1.1
```

Let's test the reachability by using Ping and Traceroute:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
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 = 1/3/4 ms

R1#Ping 123.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

On R2:

```
R2#Ping 123.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R2#Ping 123.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R3:

```
R3#Ping 123.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

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 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Question:
Why is R2 able to Ping R3 when R2 or R3 have a single mapping pointing to the hub (R1)?

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's "Traceroute" and see the what happens:

On R2:

```
R2#Traceroute 123.1.1.3 numeric

Type escape sequence to abort.
Tracing the route to 123.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 4 msec 4 msec 4 msec
  2 123.1.1.3 0 msec *  0 msec
```

NOTE: The spoke routers can reach each other through the hub routers.

```
R2#Traceroute 123.1.1.3 numeric

Type escape sequence to abort.
Tracing the route to 123.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 4 msec 4 msec
  2 123.1.1.3 0 msec *  0 msec

R2#Traceroute 123.1.1.3 numeric

Type escape sequence to abort.
Tracing the route to 123.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 0 msec 4 msec
  2 123.1.1.3 4 msec *  0 msec
```

We can see that the spoke-to-spoke communication has to traverse through R1, the hub router.

Task 3

Configure all three routers to advertise their loopback 0 interfaces using RIPv2. If this configuration is performed correctly, every router should have reachability to every loopback0 interface.

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 123.0.0.0
R1(config-router)#Network 1.0.0.0
```

On R2:

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 123.0.0.0
R2(config-router)#Network 2.0.0.0
```

On R3:

```
R3(config)#Router rip
R3(config-router)#No au
R3(config-router)#Ver 2
R3(config-router)#Network 123.0.0.0
R3(config-router)#Network 3.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is 200.1.1.10 to network 0.0.0.0

      2.0.0.0/24 is subnetted, 1 subnets
R        2.2.2.0 [120/1] via 123.1.1.2, 00:00:02, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/1] via 123.1.1.3, 00:00:01, Tunnel1
```

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:06, Tunnel1
```

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip route rip | B Gate
Gateway of last resort is 200.1.3.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:04, Tunnel1
```

The hub router can see the networks that the spokes are advertising, but the spokes can only see the network that the hub is advertising. Why don't they see each other's network?

The answer is "Split-horizon". The rule for Split-horizon states that "I will NOT advertise a route out of the interface through which I learned the route".

Let's disable the "IP Split-horizon" and verify the result:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#No ip split-horizon
```

To verify the configuration:

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:22, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/2] via 123.1.1.3, 00:00:22, Tunnel1
```

This is perfect, it worked, but the next-hop IP address is pointing to 123.1.1.3 which is R3. Does this mean that R2 goes directly to R3 to reach network 3.0.0.0/8?

Lets verify using the traceroute:

```
R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 4 msec 4 msec 0 msec
  2 123.1.1.3 0 msec *  0 msec
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 4 msec 4 msec
  2 123.1.1.3 4 msec *  0 msec

R2#Traceroute 3.3.3.3 source 2.2.2.2

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 4 msec 4 msec
  2 123.1.1.3 0 msec *  0 msec

R2#Traceroute 3.3.3.3 source 2.2.2.2

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 4 msec 4 msec 4 msec
  2 123.1.1.3 0 msec *  0 msec
```

This should convince us that R2 has to traverse through R1 to reach R3 or the networks that R3 is advertising. The routing table is pointing to 123.1.1.3 via the tunnel 1 interface, so when the packet exits the tunnel 1 interface R1 gets it (Remember that the tunnel on the spokes are point-to-point GRE), when R1 receives the packet it checks its routing table and the routing table is pointing to 123.1.1.3, then it looks up the NHRP mapping and it knows that the NBMA IP address is 200.1.3.3 so it sends the packet to R3.

Task 4

Remove the "No ip split-horizon" and provide another solution. DO NOT use PBR to accomplish this task.

Once the split-horizon is enabled, the spokes will loose the networks advertised by the other spokes; summarization can be used to achieve the same goal. When summarization is performed the next-hop is set to the IP address of the router that performed the summary. Let's configure and test this solution:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Int tunnel 1
R1(config-if)#ip split-horizon

R1(config)#Int tunnel 1
R1(config-if)#IP summary-address rip 0.0.0.0 0.0.0.0

R1#Clear ip route *
```

To verify the configuration:

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:01:25, Tunnel1
```

Why can't we see the default route in RIP?
Because we already have a static default route that we configured in Task 1.

What we should do is remove the default route and add two specific static routes instead, one static route should be configured for R1's NBMA IP address, and the second one should be configured for R3's NBMA IP address. Let's configure this and verify:

On R2:

```
R2(config)#No ip route 0.0.0.0 0.0.0.0 200.1.2.10

R2(config)#IP route 200.1.1.0 255.255.255.0 200.1.2.10
R2(config)#IP route 200.1.3.0 255.255.255.0 200.1.2.10
```

On R3:

```
R3(config)#No ip route 0.0.0.0 0.0.0.0 200.1.3.10

R3(config)#IP route 200.1.1.0 255.255.255.0 200.1.3.10
R3(config)#IP route 200.1.2.0 255.255.255.0 200.1.3.10
```

Let's verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2#Show ip route rip | B Gate
Gateway of last resort is 123.1.1.1 to network 0.0.0.0

R*    0.0.0.0/0 [120/1] via 123.1.1.1, 00:00:23, Tunnel1
```

On R3:

```
R3#Show ip route rip | B Gate
Gateway of last resort is 123.1.1.1 to network 0.0.0.0

R*    0.0.0.0/0 [120/1] via 123.1.1.1, 00:00:15, Tunnel1
```

To test the configuration:

On R2:

```
R2#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

R2#Ping 1.1.1.1 Source 2.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 3.3.3.3
Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R2#Ping 3.3.3.3 Source 2.2.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
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

Task 5

Remove the RIP routing protocol and configure EIGRP AS 100, if this configuration is successful, all the routers should have reachability to every network advertised in EIGRP routing domain.

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#No ip summary-address rip 0.0.0.0 0.0.0.0
```

On All Routers:

```
Rx(config)#No router rip
```

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 123.1.1.1 0.0.0.0
R1(config-router)#Network 1.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp 100
R2(config-router)#Network 123.1.1.2 0.0.0.0
R2(config-router)#Network 2.2.2.2 0.0.0.0
```

On R3:

```
R3(config)#Router eigrp 100
R3(config-router)#Network 123.1.1.3 0.0.0.0
R3(config-router)#Network 3.3.3.3 0.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is 200.1.1.10 to network 0.0.0.0

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/27008000] via 123.1.1.2, 00:01:19, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D        3.3.3.0 [90/27008000] via 123.1.1.3, 00:00:34, Tunnel1
```

On R2:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:02:25, Tunnel1
```

On R3:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:02:07, Tunnel1
```

Once again we see the same problem that we had with RIPv2, let's disable the split-horizon and verify:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#No ip split-horizon eigrp 100
```

To verify the configuration:

On R2:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:04:46, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
D        3.3.3.0 [90/28288000] via 123.1.1.1, 00:00:38, Tunnel1
```

On R3:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:04:20, Tunnel1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/28288000] via 123.1.1.1, 00:00:56, Tunnel1
```

Can we resolve this problem without disabling the Split-horizon in EIGRP? Yes, we should be able to, let's configure a default route like we configured in RIPv2 and verify:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#IP split-horizon eigrp 100
R1(config-if)#IP summary-address eigrp 100 0.0.0.0 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Ping 3.3.3.3 source 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R2#Ping 1.1.1.1 source 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

To see the routing table of the spoke routers:

On R2:

```
R2#Show ip route eigrp | B Gate
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Gateway of last resort is 123.1.1.1 to network 0.0.0.0

D*    0.0.0.0/0 [90/27008000] via 123.1.1.1, 00:12:57, Tunnel1
```

On R3:

```
R3#Show ip route eigrp | B Gate
Gateway of last resort is 123.1.1.1 to network 0.0.0.0

D*    0.0.0.0/0 [90/27008000] via 123.1.1.1, 00:14:02, Tunnel1
```

Task 6

Remove the EIGRP routing protocol and configure OSPF in Area 0. If this configuration is successful, all the routers in this topology should have reachability to every network advertised in this area.

On All Routers:

```
Rx(config)#No router eigrp 100
```

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Network 123.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 1.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Network 123.1.1.2 0.0.0.0 area 0
R2(config-router)#Network 2.2.2.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Network 123.1.1.3 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-router)#Network 3.3.3.3 0.0.0.0 area 0
```

Once we configure OSPF on all routers, you will notice that the HUB router gets an unbelievable amount of console messages stating the local router has established an adjacency with R2 and then it goes down and tries to establish an adjacency with R3 and the cycle keeps on repeating, this is due to the OSPF network type on GRE tunnels, OSPF does not read the Tunnel mode on the HUB router, therefore it sees the tunnel's network as a Point-to-point network type. To verify:

On R1:

```
R1#Show ip ospf interface Tunnel 1 | Inc Network

 Internet Address 123.1.1.1/24, Area 0, Attached via Network Statement
 Process ID 1, Router ID 0.0.0.1, Network Type POINT_TO_POINT, Cost: 1000
```

The following are the console messages on R1 stating that the adjacency keeps on flapping between R2 and R3:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Tunnel1 from EXCHANGE to DOWN,
Neighbor Down: Adjacency forced to reset

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Tunnel1 from EXCHANGE to DOWN,
Neighbor Down: Adjacency forced to reset
```

The question is which OSPF network type should we use?
Let's check our options:

Point-to-point on all three routers:
We know that point-to-point on all three routers will NOT work, this was the default and we had problems.

Broadcast on all three routers:
Since the next-hop IP address does NOT change on OSPF broadcast network types, this network type will work, we have to remember to configure an OSPF priority of zero on the spokes and allow Multicast traffic to go through.

Non-broadcast:
This network type should also work, we have to remember to set the OSPF priority to zero on the spoke routers (R2 and R3), and configure the "Neighbor" command on the Hub router. With this network type we don't need to enable Multicast traffic through, becasue once the neighbor commands are configured, OSPF will use Unicast.

Point-to-multipoint:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

This network type will also work and we don't need to worry about DR/BDR but once again we need to remember to allow multicast traffic through.

Let's configure OSPF "Point-to-multipoint" network type and verify:

On All Routers:

```
Rx(config)#Int tunnel 1
Rx(config-if)#IP OSPF network point-to-multipoint
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is 200.1.1.10 to network 0.0.0.0

      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/1001] via 123.1.1.2, 00:01:57, Tunnel1
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/1001] via 123.1.1.3, 00:01:57, Tunnel1
      123.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        123.1.1.2/32 [110/1000] via 123.1.1.2, 00:01:57, Tunnel1
O        123.1.1.3/32 [110/1000] via 123.1.1.3, 00:01:57, Tunnel1
```

On R2:

```
R2#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 [110/1001] via 123.1.1.1, 00:02:31, Tunnel1
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/2001] via 123.1.1.1, 00:02:31, Tunnel1
      123.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        123.1.1.1/32 [110/1000] via 123.1.1.1, 00:02:31, Tunnel1
O        123.1.1.3/32 [110/2000] via 123.1.1.1, 00:02:31, Tunnel1
```

On R3:

```
R3#Show ip route ospf | 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
O        1.1.1.1 [110/1001] via 123.1.1.1, 00:02:57, Tunnel1
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/2001] via 123.1.1.1, 00:02:47, Tunnel1
      123.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        123.1.1.1/32 [110/1000] via 123.1.1.1, 00:02:57, Tunnel1
O        123.1.1.2/32 [110/2000] via 123.1.1.1, 00:02:47, Tunnel1
```

To test the configuration:

On R3:

```
R3#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 = 1/3/4 ms

R3#Ping 1.1.1.1 source 3.3.3.3

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 3.3.3.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

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 = 1/3/4 ms

R3#Ping 2.2.2.2 source 3.3.3.3

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 3.3.3.3
!!!!!
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

Task 7

Remove OSPF and configure BGP instead, if this configuration is performed successfully the routers should have reachability to every route advertised in this topology. R1 should be configured in AS 100, R2 should be configured in AS 200, and R3 should be configured in AS 300.

On All Routers:

```
Rx(config)#No Router ospf 1

Rx(config)#Int tunnel 1
Rx(config-if)#No ip ospf network point-to-multipoint
```

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 123.1.1.2 remote-as 200
R1(config-router)#Neighbor 123.1.1.3 remote-as 300
R1(config-router)#Network 1.1.1.0 mask 255.255.255.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 123.1.1.1 remote-as 100
R2(config-router)#Neighbor 123.1.1.3 remote-as 300
R2(config-router)#Neighbor 123.1.1.3 ebgp-multihop 2
R2(config-router)#Network 2.2.2.0 mask 255.255.255.0
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Neighbor 123.1.1.1 remote-as 100
R3(config-router)#Neighbor 123.1.1.2 remote-as 200
R3(config-router)#Neighbor 123.1.1.2 ebgp-multihop 2
R3(config-router)#Network 3.3.3.0 mask 255.255.255.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Net
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*  2.2.2.0/24       123.1.1.2                              0 300 200 i
*>                  123.1.1.2                0             0 200 i
*  3.3.3.0/24       123.1.1.3                              0 200 300 i

*>                  123.1.1.3                0             0 300 i
```

We can see multiple paths for networks 2.2.2.0/24 and 3.3.3.0/24, this is because R2 has an adjacency with R3 and it advertises its loopback 0 to both neighbors and both neighbors advertise the same network to R1, but R1 takes the path by going directly to R2 for network 2.2.2.0/24 and R3 for network 3.3.3.0/24. Why? Well, you should look at the AS-Path attribute, remember the shortest as-path wins when the "Weight", "Local-Preference", are equal and the third item in the bestpath list is "Network, Aggregation and Redistribution" and after this the next attribute is the "AS-Path" attribute, since R1 has a shorter AS-Path through R2 for network 2.2.2.0/24, it chooses R2 as the best path, the same applied for network 3.3.3.0/24.

On R2:

```
R2#Show ip bgp | B Net
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.0/24       123.1.1.1                              0 300 100 i
*>                  123.1.1.1                0             0 100 i
*> 2.2.2.0/24       0.0.0.0                  0         32768 i
*  3.3.3.0/24       123.1.1.3                              0 100 300 i
*>                  123.1.1.3                0             0 300 i
```

On R3:

```
R3#Show ip bgp | B Net
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.0/24       123.1.1.1                              0 200 100 i
*>                  123.1.1.1                0             0 100 i
*> 2.2.2.0/24       123.1.1.2                0             0 200 i
*                   123.1.1.2                              0 100 200 i
*> 3.3.3.0/24       0.0.0.0                  0         32768 i
```

To verify and test the configuration:

On R3:

```
R3#Ping 1.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
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


R3#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 3.3.3.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

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 = 1/2/4 ms

R3#Ping 2.2.2.2 source 3.3.3.3

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 3.3.3.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Task 8

Configure R1 to perform the NHRP mapping dynamically. R1 should NOT be configured with a static NHRP mapping at all. If this configuration is performed successfully, all the routers in this topology should have reachability to every route advertised by the other routers.

On R1:

```
R1(config)#Int Tunnel 1
R1(config-if)#No ip nhrp map multicast 200.1.3.3
R1(config-if)#No ip nhrp map multicast 200.1.2.2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#No ip nhrp map 123.1.1.3 200.1.3.3
R1(config-if)#No ip nhrp map 123.1.1.2 200.1.2.2
```

Let's add the following command. The following command allows Multicast traffic to be mapped dynamically to all NHCs, therefore, instead of mapping multicast for each spoke's NBMA-IP statically, this command will perform those mapping dynamically:

```
R1(config-if)#IP NHRP map multicast dynamic

R1#Clear ip nhrp
```

On R2 and R3:

```
Rx(config)#Int tunnel 1
Rx(config-if)#IP NHRP NHS 123.1.1.1
```

The spokes (NHCs) need to register this mapping with the NHS, therefore, the NHCs MUST know the Tunnel IP address of the NHS in order to perform the registration.

To verify the configuration:

On R1:

```
R1#Show ip nhrp detail

123.1.1.2/32 via 123.1.1.2
    Tunnel1 created 00:05:51, expire 01:54:19
    Type: dynamic, Flags: unique registered
    NBMA address: 200.1.2.2
123.1.1.3/32 via 123.1.1.3
    Tunnel1 created 00:06:26, expire 01:54:20
    Type: dynamic, Flags: unique registered
    NBMA address: 200.1.3.3

R1#Show ip bgp | B Net
     Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*  2.2.2.0/24       123.1.1.2                              0 300 200 i
*>                  123.1.1.2                0             0 200 i
*  3.3.3.0/24       123.1.1.3                              0 200 300 i
*>                  123.1.1.3                0             0 300 i
```

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2#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

R2#Ping 3.3.3.3

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

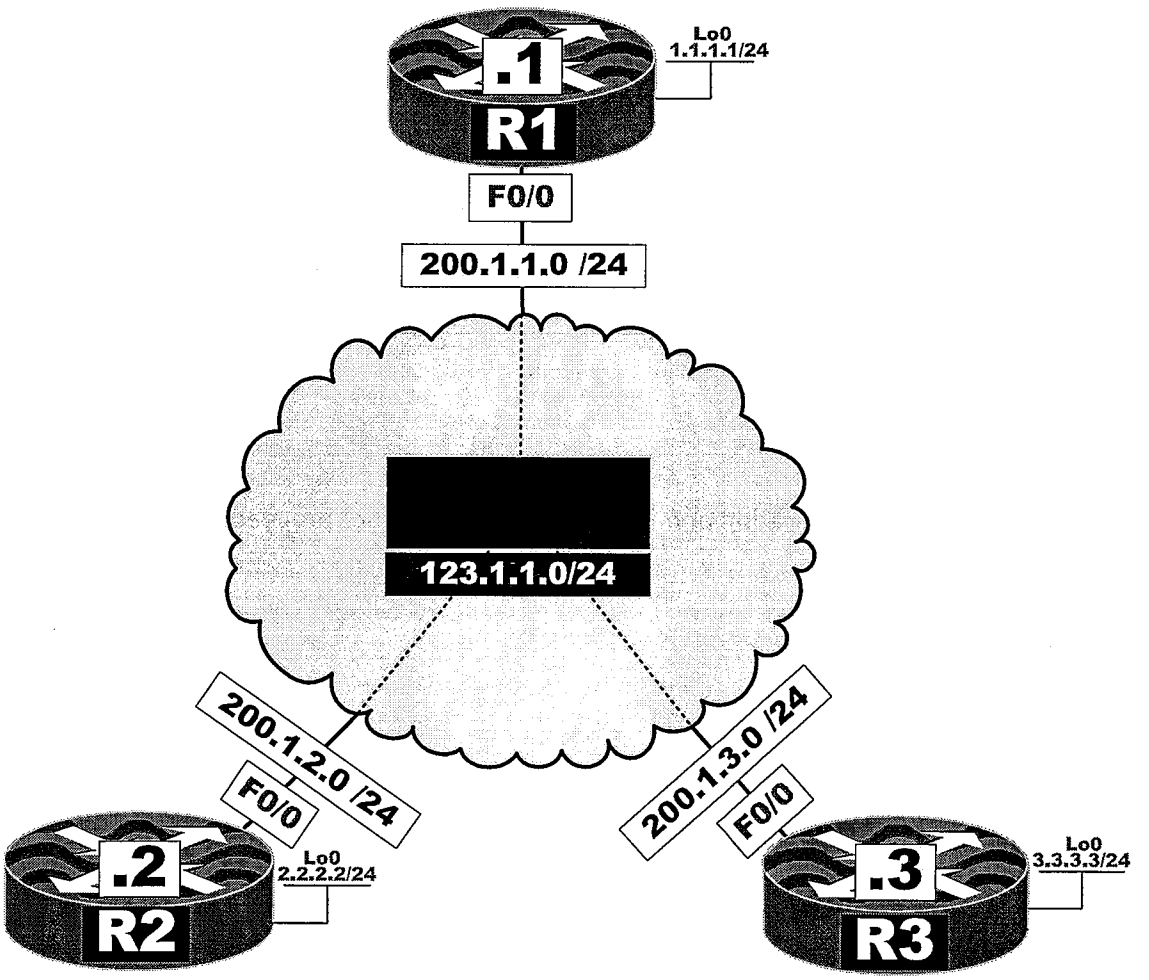
Task 9

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 6 – Running Routing Protocols on DMVPN Phase #2 and Phase #3



Task 1

SW1 represents the Internet; configure a static default route on each router pointing to the appropriate interface on SW1. If this configuration is performed correctly, these

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

routers should be able to ping and have reachability to the F0/0 interfaces of all routers in this topology. The switch interface to which the routers are connected to should have a ".10" in the host portion of the IP address for that subnet.

Let's configure SW1's interfaces for these routers. Since in this lab SW1 represents the Internet, the IP addresses in the following configuration should be configured as the default gateway on the routers.

On SW1:

```
SW1(config)#Int range f0/1-3
SW1(config-if-range)#No switchport

SW1(config)#Int F0/1
SW1(config-if)#ip address 200.1.1.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#ip address 200.1.2.10 255.255.255.0
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#ip address 200.1.3.10 255.255.255.0
SW1(config-if)#No shut
```

Let's NOT forget to enable "IP routing" or else the switch will not be able to route from one subnet to another.

```
SW1(config)#IP routing
```

Let's configure the routers:

On R1:

```
R1(config)#int f0/0
R1(config-if)#ip addr 200.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 200.1.1.10
```

On R2:

```
R2(config)#Int f0/0
R2(config-if)#ip addr 200.1.2.2 255.255.255.0
R2(config-if)#No shut
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#ip route 0.0.0.0 0.0.0.0 200.1.2.10
```

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 200.1.3.3 255.255.255.0
R3(config-if)#No shut

R3(config)#ip route 0.0.0.0 0.0.0.0 200.1.3.10
```

To verify the configuration:

On R1:

```
R1#Ping 200.1.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Ping 200.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R2:

```
R2#Ping 200.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.1, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 200.1.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.3.3, timeout is 2 seconds:
!!!!!
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

Task 2

Configure mGRE tunnel interface on the routers in the above diagram. These routers should use their F0/0 interface as the source of the tunnel. R1 should be configured as NHS, and R2 and R3 should be configured as NHCs. R1 should NOT be configured with any NHRP mappings. The NHCs should NOT traverse the NHS to reach networks advertised by other NHCs.

DMVPN Phase 2:

The main problem with DMVPN Phase 1 is that the spoke-to-spoke communication can ONLY occur through the hub router.

DMVPN Phase 2 allows spoke-to-spoke communication.

But how?

In Phase #2, the SPOKES are configured with the "Tunnel mode gre multipoint" tunnel interface configuration command and NOT GRE, this means that the spokes should NOT be configured with the "Tunnel destination" command.

More details.....

With NHRP, the NBMA network is initially laid out as a hub and spoke network, with the NHS as the hub and NHCs as the spokes. The NHCs are configured with static mapping to reach the NHS, and this mapping is used by the NHCs to send registration requests to the NHS. This allows the NHS to dynamically learn the mapping information for the spokes (NHCs).
Once the NBMA hub and spoke network is dynamically built, NHRP resolution requests and responses are used between the NHCs and the NHS to dynamically discover the spoke-to-spoke (NHC to NHC) mapping so the spokes can bypass the hub and connect to each other directly. Of course the layer two connectivity must be there to allow spoke-to-spoke communication without traversing the hub.
This configuration tremendously reduces the amount of configuration on the hub when a new spoke is added to the network.

On R1:

```
R1(config)#Int Tunnel 1
R1(config-if)#IP addr 123.1.1.1 255.255.255.0
R1(config-if)#Tunnel source 200.1.1.1
R1(config-if)#Tunnel mode gre multipoint
R1(config-if)#IP NHRP network-id 111
R1(config-if)#IP NHRP map multicast dynamic
```

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#Int Tunnel 1
R2(config-if)#IP address 123.1.1.2 255.255.255.0
R2(config-if)#Tunnel source 200.1.2.2
R2(config-if)#Tunnel mode gre multipoint
R2(config-if)#IP NHRP network-id 222
R2(config-if)#IP NHRP NHS 123.1.1.1
R2(config-if)#IP NHRP map 123.1.1.1 200.1.1.1
R2(config-if)#IP NHRP map multicast 200.1.1.1
```

On R3:

```
R3(config)#Int Tunnel 1
R3(config-if)#IP addr 123.1.1.3 255.255.255.0
R3(config-if)#Tunnel source 200.1.3.3
R3(config-if)#Tunnel mode gre multipoint
R3(config-if)#IP NHRP Network-id 333
R3(config-if)#IP NHRP NHS 123.1.1.1
R3(config-if)#IP NHRP map 123.1.1.1 200.1.1.1
R3(config-if)#IP NHRP map multicast 200.1.1.1
```

To verify the configuration:

On R1:

```
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 = 1/3/4 ms

R1#Ping 123.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R1#Show ip nhrp detail

123.1.1.2/32 via 123.1.1.2
    Tunnel1 created 00:02:46, expire 01:59:01
    Type: dynamic, Flags: unique registered
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
    NBMA address: 200.1.2.2
123.1.1.3/32 via 123.1.1.3
    Tunnel1 created 00:00:50, expire 01:59:09
    Type: dynamic, Flags: unique registered
    NBMA address: 200.1.3.3
```

We can see that the mapping of Tunnel IP address to the NBMA IP address was done dynamically, this means that the spoke routers (NHCs) must have sent a registration request to the NHS. Let's verify the NHCs:

On R2:

```
R2#Ping 123.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R2#Ping 123.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/8 ms

R2#Show ip nhrp detail

123.1.1.1/32 via 123.1.1.1
   Tunnel1 created 00:10:50, never expire
   Type: static, Flags: used
   NBMA address: 200.1.1.1
123.1.1.2/32 via 123.1.1.2
   Tunnel1 created 00:00:06, expire 01:59:53
   Type: dynamic, Flags: router unique local
   NBMA address: 200.1.2.2
     (no-socket)
  Requester: 123.1.1.3 Request ID: 1
123.1.1.3/32 via 123.1.1.3
   Tunnel1 created 00:00:06, expire 01:59:53
   Type: dynamic, Flags: router used
   NBMA address: 200.1.3.3
```

We can see that the local router (R2), has established two tunnels dynamically so it can reach them directly. Let's verify this configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Traceroute 123.1.1.3

Type escape sequence to abort.
Tracing the route to 123.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 4 msec *  0 msec
```

NOTE: Because of the dynamic mapping, R2 can reach R3 directly and it does not have to traverse through R1 to reach R3.

On R3:

```
R3#Ping 123.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R3#Ping 123.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 123.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms

R3#Show ip nhrp detail

123.1.1.1/32 via 123.1.1.1
   Tunnel1 created 00:15:30, never expire
   Type: static, Flags: used
   NBMA address: 200.1.1.1
123.1.1.2/32 via 123.1.1.2
   Tunnel1 created 00:08:56, expire 01:51:03
   Type: dynamic, Flags: router
   NBMA address: 200.1.2.2
123.1.1.3/32 via 123.1.1.3
   Tunnel1 created 00:08:56, expire 00:03:00
   Type: dynamic, Flags: temporary
   NBMA address: 200.1.1.1
  Requester: 123.1.1.2 Request ID: 2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 3

Configure all three routers to advertise their loopback 0 interfaces using RIPv2. If this configuration is performed correctly, every router should have reachability to every loopback0 interface. The spoke routers (R2 and R3) should NOT traverse through R1 to reach networks advertised by other spokes.

On R1:

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 123.0.0.0
R1(config-router)#Network 1.0.0.0
```

On R2:

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 123.0.0.0
R2(config-router)#Network 2.0.0.0
```

On R3:

```
R3(config)#Router rip
R3(config-router)#No au
R3(config-router)#Ver 2
R3(config-router)#Network 123.0.0.0
R3(config-router)#Network 3.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is 200.1.1.10 to network 0.0.0.0

      2.0.0.0/24 is subnetted, 1 subnets
R        2.2.2.0 [120/1] via 123.1.1.2, 00:00:02, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/1] via 123.1.1.3, 00:00:01, Tunnel1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:06, Tunnel1
```

On R3:

```
R3#Show ip route rip | B Gate
Gateway of last resort is 200.1.3.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:04, Tunnel1
```

The hub router can see the networks that the spokes are advertising, but the spokes can only see the network that the hub is advertising. Why don't they see each other's network?

The answer is "Split-horizon". The rule for Split-horizon states that "I will NOT advertise a route out of the interface through which I learned the route".

Let's disable the "IP Split-horizon" and verify the result:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#No ip split-horizon

R1#Clear ip route *
```

To verify the configuration:

On R2:

```
R2#Show ip route rip | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:22, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/2] via 123.1.1.3, 00:00:22, Tunnel1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

This is perfect, it worked, but the next-hop IP address for network 3.3.3.0/24 is pointing to 123.1.1.3 which is R3. Does this mean that R2 goes directly to R3 to reach network 3.0.0.0/8?

Lets verify using the traceroute:

```
R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 0 msec *  0 msec

R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 0 msec *  0 msec

R2#Traceroute 3.3.3.3 source 2.2.2.2

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)

  1 123.1.1.3 4 msec *  0 msec
```

We can see that the spoke routers are no longer traversing through the hub router (R1) to reach networks that the other spokes are advertising. This is a major benefit in Phase-2 over Phase-1.

On R3:

```
R3#Show ip route rip | B Gate
Gateway of last resort is 200.1.3.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
R        1.1.1.0 [120/1] via 123.1.1.1, 00:00:14, Tunnel1
      2.0.0.0/24 is subnetted, 1 subnets
R        2.2.2.0 [120/2] via 123.1.1.2, 00:00:14, Tunnel1

R3#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  1 123.1.1.2 4 msec *  0 msec

R3#Traceroute 2.2.2.2 source 3.3.3.3 numeric

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec
```

We can see that RIPv2 works well without any modification, the spokes can reach each other's advertised networks without traversing through the hub router (R1). The ONLY issue was the IP split-horizon that needed to be disabled so the spoke routers could see networks advertised by other spokes.

Task 4

Remove the RIP routing protocol and configure EIGRP AS 100, if this configuration is successful, all the routers should have reachability to every network advertised in EIGRP routing domain. The spoke routers should reach each other's networks by going directly to each other without traversing through R1.

On R1:

On All Routers:

```
Rx(config)#No router rip
```

On R1:

```
R1(config)#Int Tunnel 1
R1(config-if)#IP split-horizon
```

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 123.1.1.1 0.0.0.0
R1(config-router)#Network 1.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp 100
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config-router)#Network 123.1.1.2 0.0.0.0
R2(config-router)#Network 2.2.2.2 0.0.0.0
```

On R3:

```
R3(config)#Router eigrp 100
R3(config-router)#Network 123.1.1.3 0.0.0.0
R3(config-router)#Network 3.3.3.3 0.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is 200.1.1.10 to network 0.0.0.0

      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/27008000] via 123.1.1.2, 00:01:19, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
D        3.3.3.0 [90/27008000] via 123.1.1.3, 00:00:34, Tunnel1
```

On R2:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:02:25, Tunnel1
```

On R3:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:02:07, Tunnel1
```

Once again we see the same problem that we had with RIPv2, let's disable the split-horizon for EIGRP in AS 100 and verify:

On R1:

```
R1(config)#Int tunnel 1
R1(config-if)#No ip split-horizon eigrp 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R2:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:04:46, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
D        3.3.3.0 [90/28288000] via 123.1.1.1, 00:00:38, Tunnel1
```

On R3:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:04:20, Tunnel1
      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/28288000] via 123.1.1.1, 00:00:56, Tunnel1
```

We can see that the spokes are going through the hub to reach the networks advertised by the other spokes. Let's verify:

On R2:

```
R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 0 msec 4 msec
  2 123.1.1.3 4 msec *  0 msec

R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 0 msec 4 msec
  2 123.1.1.3 4 msec *  0 msec
```

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 4 msec 4 msec 0 msec
  2 123.1.1.2 0 msec *  0 msec

R3#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 4 msec 4 msec
  2 123.1.1.2 4 msec *  0 msec
```

The reason the spoke routers are traversing through the hub router has nothing to do with the way we have configured NHRP, this has to do with EIGRP and the way EIGRP is handling the hext-hop IP address.

By default, EIGRP will change the next-hop IP address to its own interface's IP address when it receives a network from a given neighbor and advertises that same network/s to another neighbor, this behavior can be changed such that EIGRP does NOT change the next-hop IP address.

Let's configure this and verify:

On R1:

```
R1(config)#Int Tunnel 1
R1(config-if)#No ip next-hop-self Eigrp 100
```

To verify the configuration:

On R2:

```
R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 4 msec *  0 msec

R2#Traceroute 3.3.3.3 source 2.2.2.2

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 4 msec *  0 msec
```

On R2:

```
R3#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec

R3#Traceroute 2.2.2.2 source 3.3.3.3

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec
```

Task 5

Remove the EIGRP routing protocol and configure OSPF in Area 0. If this configuration is successful, all the routers in this topology should have reachability to every network advertised in this area. The spoke routers should reach each other's networks by going directly to each other without traversing through R1. The loopback interfaces should be advertised with their correct mask.

On All Routers:

```
Rx(config)#No router eigrp 100
```

Remember from Lab-1 that the tunnel interfaces regardless of their mode (GRE or mGRE) will default to OSPF Point-to-point network type, and the hub router will constantly establish an adjacency with R2 and then it will tear down that adjacency and establish an adjacency with R3, and this cycle will repeat itself over and over again, such that you won't get a chance to enter any command/s on R1. If this happens in the lab because you forgot to change OSPF's network type on the hub router, you can shutdown the F0/0 interfaces on both R2 and R3 and this will stop the console messages.

In this scenario, the spoke routers should go directly to other spoke routers, therefore, OSPF's network

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

type can be manipulated to achieve this goal. In this case the network type is changed to "Broadcast". Let's briefly cover the other OSPF network types:

Point-to-point on all three routers:
We know that point-to-point on all three routers will NOT work, this was the default and we had problems.

Broadcast on all three routers:
Since the next-hop IP address does NOT change on OSPF broadcast network types, this network type will work, we have to remember to configure an OSPF priority of zero on the spokes and allow Multicast traffic to go through, this is done with the "IP NHRP map multicast" command.

Non-broadcast:
This network type should also work, we have to remember to set the OSPF priority to zero on the spoke routers (R2 and R3), and configure the "Neighbor" command on the Hub router. With this network type we don't need to enable Multicast traffic through, becasue once the neighbor commands are configured, OSPF will use Unicast.

Point-to-multipoint:
This network type will also work and we don't need to worry about DR/BDR but once again we need to remember to allow multicast traffic through.

Let's configure the OSPF network type on the Tunnel interfaces on the hub and the spoke routers to be "Broadcast" but remember that the hub router should be the DR, so besides changing the network type, we should also set the OSPF priority on the Tunnel interfaces of the spokes to zero so they do not participate in the DR/BDR election.

On R1:

```
R1(config)#Int Tunnel 1
R1(config-if)#IP OSPF network broadcast

R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Network 123.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 1.1.1.1 0.0.0.0 area 0

R1(config)#Int Lo0
R1(config-if)#IP ospf network point-to-point
```

On R2:

```
R2(config)#Int Tunnel 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-if)#IP OSPF network broadcast
R2(config-if)#IP OSPF priority 0

R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2

R2(config-router)#Network 123.1.1.2 0.0.0.0 area 0
R2(config-router)#Network 2.2.2.2 0.0.0.0 area 0

R2(config)#Int Lo0
R2(config-if)#IP ospf network point-to-point
```

On R3:

```
R3(config)#Int Tunnel 1
R3(config-if)#IP OSPF network broadcast
R3(config-if)#IP OSPF priority 0

R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Network 123.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 3.3.3.3 0.0.0.0 area 0

R3(config)#Int Lo0

R3(config-if)#IP ospf network point-to-point
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
O        1.1.1.0 [110/1001] via 123.1.1.1, 00:04:19, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
O        3.3.3.0 [110/1001] via 123.1.1.3, 00:04:19, Tunnel1

R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 0 msec *  0 msec
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2#Traceroute 3.3.3.3 Source 2.2.2.2

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 4 msec *  0 msec
```

On R3:

```
R3#Show ip route ospf | B Gate
Gateway of last resort is 200.1.3.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
O        1.1.1.0 [110/1001] via 123.1.1.1, 00:05:33, Tunnel1
      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/1001] via 123.1.1.2, 00:05:23, Tunnel1

R3#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec

R3#Traceroute 2.2.2.2 Source 3.3.3.3

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec
```

Task 6

Remove OSPF and configure BGP instead, if this configuration is performed successfully the routers should have reachability to every route advertised in this topology. R1 should be configured in AS 100, R2 should be configured in AS 200, and R3 should be configured in AS 300. The spoke routers should reach each other's networks by going directly to each other without traversing through R1.

On All Routers:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Rx(config)#No Router ospf 1

Rx(config)#Int tunnel 1
Rx(config-if)#No ip ospf network broadcast
```

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 123.1.1.2 remote-as 200
R1(config-router)#Neighbor 123.1.1.3 remote-as 300
R1(config-router)#Network 1.1.1.0 mask 255.255.255.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Neighbor 123.1.1.1 remote-as 100
R2(config-router)#Neighbor 123.1.1.3 remote-as 300
R2(config-router)#Neighbor 123.1.1.3 ebgp-multihop 2
R2(config-router)#Network 2.2.2.0 mask 255.255.255.0
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Neighbor 123.1.1.1 remote-as 100
R3(config-router)#Neighbor 123.1.1.2 remote-as 200
R3(config-router)#Neighbor 123.1.1.2 ebgp-multihop 2
R3(config-router)#Network 3.3.3.0 mask 255.255.255.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
*  2.2.2.0/24       123.1.1.2                              0 300 200 i
*>                  123.1.1.2                0             0 200 i
*  3.3.3.0/24       123.1.1.3                              0 200 300 i
*>                  123.1.1.3                0             0 300 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

We can see multiple paths for networks 2.2.2.0/24 and 3.3.3.0/24, this is because R2 has an adjacency with R3 and it advertises its loopback 0 to both neighbors and both neighbors advertise the same network to R1, but R1 takes the path by going directly to R2 for network 2.2.2.0/24 and R3 for network 3.3.3.0/24. Why? Well, you should look at the AS-Path attribute, remember the shortest as-path wins when the "Weight", "Local-Preference", are equal and the third item in the bestpath list is "Network, Aggregation and Redistribution" and after this the next attribute is the "AS-Path" attribute, since R1 has a shorter AS-Path through R2 for network 2.2.2.0/24, it chooses R2 as the best path, the same applied for network 3.3.3.0/24.

On R2:

```
R2#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.0/24       123.1.1.1                              0 300 100 i
*>                  123.1.1.1                0             0 100 i
*> 2.2.2.0/24       0.0.0.0                  0         32768 i
*  3.3.3.0/24       123.1.1.3                              0 100 300 i
*>                  123.1.1.3                0             0 300 i
```

On R3:

```
R3#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path

*  1.1.1.0/24       123.1.1.1                              0 200 100 i
*>                  123.1.1.1                0             0 100 i
*> 2.2.2.0/24       123.1.1.2                0             0 200 i
*                   123.1.1.2                              0 100 200 i
*> 3.3.3.0/24       0.0.0.0                  0         32768 i
```

To verify and test the configuration:

On R3:

```
R3#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

R3#Ping 1.1.1.1 source lo0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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 3.3.3.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

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 = 1/2/4 ms

R3#Ping 2.2.2.2 source 3.3.3.3
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 3.3.3.3
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R3#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec

R3#Traceroute 2.2.2.2 source 3.3.3.3

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.2 4 msec *  0 msec
```

Task 7

Re-configure the routers for NHRP Phase 3. You should configure mGRE tunnel interfaces on the routers in the above diagram. These routers should use their F0/0 interface as the source of the tunnel. R1 should be configured as NHS, and R2 and R3 should be configured as NHCs. R1 should NOT be configured with any NHRP mappings. The NHCs should NOT traverse the NHS to reach networks advertised by other NHCs.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

You should test NHRP Phase 3 using EIGRP, DO NOT use PBR or the "No ip next-hop Eigrp" command to accomplish this task.

NHRP Phase 2 and 3 are identical with one major difference. The IP next-hop EIGRP was introduced in IOS 12.3, and the "IP NHRP Shortcut" was introduced in 12.4(6)T. In NHRP Phase 2, spoke-to-spoke was achieved through NHRP registration requests and replies but once EIGRP was configured the next-hop had to be changed and this was done through manipiulation of EIGRP using the "No ip next-hop EIGRP" command.

With the "IP NHRP redirect" configured on the hub and the "IP NHRP Shortcut" command on the spoke/s, we no longer need to manipulate EIGRP to achieve spoke-to-spoke communication for the networks that the spokes advertise. Let's configure Phase 3 and verify its behavior:

Let's check the existing configuration on the hub router before we configure the "IP NHRP redirect":

On R1:

```
R1#Show run int tunnel 1 | B interface

interface Tunnel1
 ip address 123.1.1.1 255.255.255.0
 no ip redirects
 ip nhrp map multicast dynamic
 ip nhrp network-id 111
 tunnel source 200.1.1.1
 tunnel mode gre multipoint

R1(config)#Int Tunnel 1
R1(config-if)#IP NHRP redirect
```

The "IP NHRP Redirect" message is an indiction that the currect path to the destination is not optimal and that the receiver of the message should find a better path to the destination.

Let's configure the spokes:

On R2 and R3:

```
R2(config)#Int tunnel 1
R2(config-if)#IP NHRP Shortcut
```

The "IP NHRP Shortcut" message overrides the routing table only if it receives an "IP NHRP Redirect" message. Let's remove BGP and configure EIGRP and test this feature:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#No router bgp 100

R1(config)#Router eigrp 100
R1(config-router)#Network 123.1.1.1 0.0.0.0
R1(config-router)#Network 1.1.1.1 0.0.0.0
```

We know that the IP split-horizon for EIGRP needs to be disable for the spokes to see each other's routes:

```
R1(config)#Int tunnel 1
R1(config-if)#No ip split-horizon eigrp 100
```

On R2:

```
R2(config)#No router bgp 200

R2(config)#Router eigrp 100
R2(config-router)#Network 123.1.1.2 0.0.0.0
R2(config-router)#Network 2.2.2.2 0.0.0.0
```

On R3:

```
R3(config)#No router bgp 300

R3(config)#Router eigrp 100
R3(config-router)#Network 123.1.1.3 0.0.0.0
R3(config-router)#Network 3.3.3.3 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | B Gate
Gateway of last resort is 200.1.2.10 to network 0.0.0.0

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/27008000] via 123.1.1.1, 00:03:00, Tunnel1
      3.0.0.0/24 is subnetted, 1 subnets
D        3.3.3.0 [90/28288000] via 123.1.1.1, 00:01:06, Tunnel1
```

NOTE: The next-hop to reach network 3.3.3.0/24 is the hub router, but because the "IP NHRP Redirect" is configured on the hub and the "IP NHRP Shortcut" is configured on the spokes, they override the entries in the routing table.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To test the configuration:

On R2:

```
R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.1 0 msec 0 msec 4 msec
  2 123.1.1.3 4 msec *  0 msec

R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 123.1.1.3 4 msec *  0 msec
```

We can see that the first traceroute went through the hub router, and the second and subsequent traffic went directly to R3. The questions is why?

Well, when the first traceroute was sent out the local router did not have a mapping for R3 so it had no choice but to send a resolution request to the NHS (R1) and once it received the response from R1, it cached that information and the subsequent traffic used the cached entry for R3, so the question is how do you want to provide spoke-to-spoke reachability? Do you want to achieve it through the routing protocol or NHRP? If the answer is through routing protocol, then Phase 2 is your choice, but if you need to resolve the problem through NHRP, then Phase 3 should be your choice.

to see the cached entries:

```
R2#Show ip nhrp detail

3.3.3.0/24 via 123.1.1.3
   Tunnel1 created 00:07:03, expire 01:52:56
   Type: dynamic, Flags: router used
   NBMA address: 200.1.3.3
123.1.1.1/32 via 123.1.1.1
   Tunnel1 created 03:18:03, never expire
   Type: static, Flags: used
   NBMA address: 200.1.1.1
123.1.1.3/32 via 123.1.1.3
   Tunnel1 created 01:06:08, expire 00:53:51
   Type: dynamic, Flags: router
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
NBMA address: 200.1.3.3
```

Task 8

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

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Advanced
CCIE Routing & Switching
v5.0

www.MicronicsTraining.com

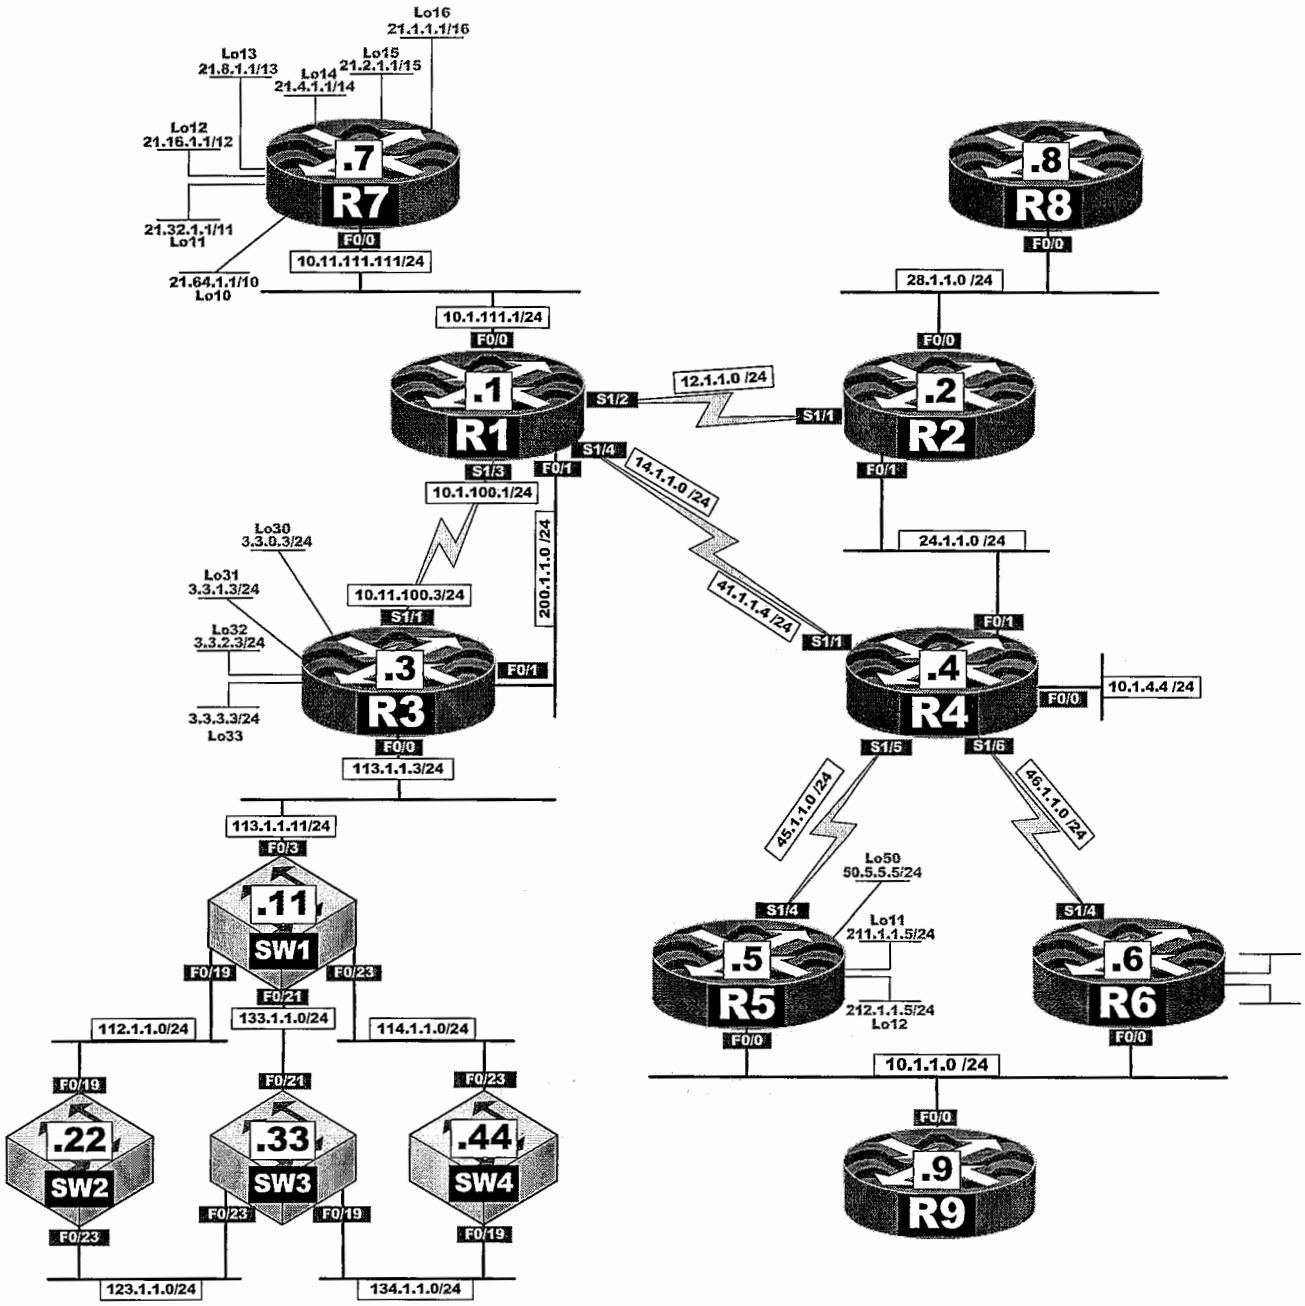
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

RIPv2

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Configuring RIPv2



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "RIPv2"→ "Lab-1".

Task 1

Configure RIPv2 on all routers/switches and advertise their directly connected interface/s in this routing domain. These devices should NOT have a classful nature. If this configuration is successful, the routers/switches should successfully exchange routes.

Use the following policy when configuring this task:

1. R7 and R1 must use two static routes for NLRI.
2. R1 and R4 can NOT use a static route/s, R1 can NOT use any solution that's used in the previous policy or the next ones to provide reachability.
3. R1 and R3 should NOT use static route/s, run an extra layer 2 protocol on the interface, secondary IP addresses, SPAN, RSPAN, and/or Tunneling of any type or disable the sanity check in RIP to accomplish this task for their connection through Serial interfaces.
4. You can NOT use PBR to resolve any of the conditions in this task.

On R1:

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
```

Note RIP process is NOT started unless a Network is advertised:

```
R1(config-router)#do Show run | s router rip
R1#
```

Let's advertise Network 10.0.0.0:

```
R1(config-router)#Netw 10.0.0.0
R1(config-router)#Netw 12.0.0.0
R1(config-router)#Netw 14.0.0.0
R1(config-router)#Netw 200.1.1.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R7:

Let's see the IP addresses configured on R7:

```
R7#Show ip int brief | Exc unass

Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/0            10.11.111.111   YES manual up                    up
Loopback10                 21.64.1.1       YES manual up                    up
Loopback11                 21.32.1.1       YES manual up                    up
Loopback12                 21.16.1.1       YES manual up                    up
Loopback13                 21.8.1.1        YES manual up                    up
Loopback14                 21.4.1.1        YES manual up                    up
Loopback15                 21.2.1.1        YES manual up                    up
Loopback16                 21.1.1.1        YES manual up                    up

R7(config)#Router rip
R7(config-router)#No au
R7(config-router)#Ver 2
R7(config-router)#Network 10.0.0.0
R7(config-router)#Network 21.0.0.0
```

To verify the configuration:

On R1:

```
R1#Sh ip rou rip | B Gate
Gateway of last resort is not set
R1#
```

The reason R1 and R7 are NOT exchanging routes is because they are on different networks, the following verifies this point:

On R1:

```
R1#Sh run int F0/0 | Inc ip address

 ip address 10.1.111.1 255.255.255.0
```

On R7:

```
R7#Show run int g0/0 | Inc ip address

 ip address 10.11.111.111 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

By default, sanity check enabled, Sanity check is performed by every RIP router before exchanging updates to ensure that the source IP address of incoming routing updates is on the same IP network as the interface through which the update was received. This sanity check can be disabled by configuring the "No validate-update-source" router configuration command; this command must be configured on both routers.

On R1 and R7:

```
(config)#router rip
(config-router)#No validate-update-source
```

To verify the configuration:

On R1:

```
R1#Clear ip rou *

R1#Show ip route rip | B Gate
Gateway of last resort is not set

      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/1] via 10.11.111.111, 00:00:16
R        21.2.0.0/15 [120/1] via 10.11.111.111, 00:00:16
R        21.4.0.0/14 [120/1] via 10.11.111.111, 00:00:16
R        21.8.0.0/13 [120/1] via 10.11.111.111, 00:00:16
R        21.16.0.0/12 [120/1] via 10.11.111.111, 00:00:16
R        21.32.0.0/11 [120/1] via 10.11.111.111, 00:00:16
R        21.64.0.0/10 [120/1] via 10.11.111.111, 00:00:16
```

Note even though the routes are in the routing table, the local router does NOT have NLRI to those routes, because the next-hop IP address is not reachable, since ONLY two static routes are allowed, we can configure one on R1 and the second one on R2, these static routes are for the IP address of the next-hop router:

```
R1(config)#ip route 10.11.111.111 255.255.255.255 F0/0
```

On R7:

```
R7#Show ip route rip | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R        10.1.100.0/24 [120/1] via 10.1.111.1, 00:00:21
      12.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        12.1.1.0 [120/1] via 10.1.111.1, 00:00:21
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/1] via 10.1.111.1, 00:00:21
R     200.1.1.0/24 [120/1] via 10.1.111.1, 00:00:21

R7(config)#ip route 10.1.111.1 255.255.255.255 G0/0
```

To verify the configuration:

On R1:

```
R1#Ping 21.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 21.16.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Configuring R2 & R3:

On R2:

```
R2(config)#router rip
R2(config-router)#No au
R2(config-router)#ver 2
R2(config-router)#Network 12.0.0.0
R2(config-router)#Netw 24.0.0.0
R2(config-router)#Netw 28.0.0.0
```

To verify the configuration:

The following routes are advertised by R1, some are directly connected routes and some are routes that R1 received from R7:

On R2:

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

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R        10.1.100.0/24 [120/1] via 12.1.1.1, 00:00:04, Serial1/1
R        10.1.111.0/24 [120/1] via 12.1.1.1, 00:00:04, Serial1/1
R        10.11.111.111/32 [120/1] via 12.1.1.1, 00:00:04, Serial1/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/1] via 12.1.1.1, 00:00:04, Serial1/1
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R        21.2.0.0/15 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R        21.4.0.0/14 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R        21.8.0.0/13 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R        21.16.0.0/12 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R        21.32.0.0/11 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R        21.64.0.0/10 [120/2] via 12.1.1.1, 00:00:04, Serial1/1
R     200.1.1.0/24 [120/1] via 12.1.1.1, 00:00:04, Serial1/1
```

On R3:

```
R3(config)#router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Netw 10.0.0.0
R3(config-router)#Netw 113.0.0.0
R3(config-router)#Netw 200.1.1.0
```

To verify the configuration:

On R3:

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

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R        10.1.100.0/24 [120/1] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        10.1.111.0/24 [120/1] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        10.11.111.111/32 [120/1] via 200.1.1.1, 00:00:00, FastEthernet0/1
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 200.1.1.1, 00:00:00, FastEthernet0/1
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/1] via 200.1.1.1, 00:00:00, FastEthernet0/1
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        21.2.0.0/15 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        21.4.0.0/14 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        21.8.0.0/13 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        21.16.0.0/12 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
R        21.32.0.0/11 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        21.64.0.0/10 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
      24.0.0.0/24 is subnetted, 1 subnets
R        24.1.1.0 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
      28.0.0.0/24 is subnetted, 1 subnets
R        28.1.1.0 [120/2] via 200.1.1.1, 00:00:00, FastEthernet0/1
```

R3 is NOT receiving any routes through its S1/1 interface, this is because of the sanity check, in this case the use of static routes, running another layer 2 protocol (This also mean PPP) or disabling the sanity check is prohibited, therefore, our ONLY choice here is to assign the IP addresses of the S1/1 interfaces on R3 and the IP address of the S1/3 interface on R1 to a loopback interface and use the "IP unnumbered Loopback" interface configuration command.

When a router receives routing update/s, it installs the source IP address of the update/s as the next hop in its routing table. Normally, the next hop is a directly-connected router. When the "IP unnumbered" command is configured, the S1/1 interface "borrows" its IP address from the configured Loopback interface, which is in a different network and routes learned through the IP unnumbered interface have the interface as the next hop instead of the source IP address of the routing update. Thus avoiding an invalid next hop address problem.

On R3:

```
R3(config-if)#int s1/1
R3(config-if)#No ip addr
R3(config-if)#ip unnumbered lo300

R3(config-if)#int lo300
R3(config-if)#ip addr 10.11.100.3 255.255.255.0
```

On R1:

```
R1(config)#int s1/3
R1(config-if)#No ip addr
R1(config-if)#ip unnumbered lo100

R1(config-if)#int lo100
R1(config-if)#ip addr 10.1.100.1 255.255.255.0
```

You may have to wait a little for RIP to converge, unless the routing table is cleared using the "Clear ip route *" command.

To verify the configuration:

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R        10.1.100.0/24 [120/1] via 200.1.1.1, 00:00:23, FastEthernet0/1
                       [120/1] via 10.1.100.1, 00:00:13, Serial1/1
R        10.1.111.0/24 [120/1] via 200.1.1.1, 00:00:23, FastEthernet0/1
                       [120/1] via 10.1.100.1, 00:00:13, Serial1/1
R        10.11.111.111/32 [120/1] via 200.1.1.1, 00:00:23, FastEthernet0/1
                          [120/1] via 10.1.100.1, 00:00:13, Serial1/1
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 200.1.1.1, 00:00:23, FastEthernet0/1
                  [120/1] via 10.1.100.1, 00:00:13, Serial1/1
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/1] via 200.1.1.1, 00:00:23, FastEthernet0/1
                  [120/1] via 10.1.100.1, 00:00:13, Serial1/1
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                     [120/2] via 10.1.100.1, 00:00:13, Serial1/1
R        21.2.0.0/15 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                     [120/2] via 10.1.100.1, 00:00:13, Serial1/1
R        21.4.0.0/14 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                     [120/2] via 10.1.100.1, 00:00:13, Serial1/1
R        21.8.0.0/13 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                     [120/2] via 10.1.100.1, 00:00:13, Serial1/1
R        21.16.0.0/12 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                      [120/2] via 10.1.100.1, 00:00:13, Serial1/1
R        21.32.0.0/11 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                      [120/2] via 10.1.100.1, 00:00:13, Serial1/1
R        21.64.0.0/10 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                      [120/2] via 10.1.100.1, 00:00:13, Serial1/1
      24.0.0.0/24 is subnetted, 1 subnets
R        24.1.1.0 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                  [120/2] via 10.1.100.1, 00:00:13, Serial1/1
      28.0.0.0/24 is subnetted, 1 subnets
R        28.1.1.0 [120/2] via 200.1.1.1, 00:00:23, FastEthernet0/1
                  [120/2] via 10.1.100.1, 00:00:13, Serial1/1
```

Note routes are being exchanged through S1/1 interface of R3; with the IP unnumbered configuration, the sanity check no longer needs to be disabled for this connection and the NLRI is provided because the interface is the next hop and NOT the IP address.

To test the configuration:

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Ping 10.11.111.111

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.11.111.111, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 40/43/44 ms

R3#Traceroute 10.11.111.111 Numeric

Type escape sequence to abort.
Tracing the route to 10.11.111.111
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.100.1 8 msec
    200.1.1.1 0 msec
    10.1.100.1 8 msec
  2 10.11.111.111 0 msec *  0 msec
```

Configuring R4 & R1:

Note once again the IP address of their Serial interfaces connecting them to each other are from different networks:

On R4:

```
R4#Sh run int S1/1 | inc ip address

ip address 41.1.1.4 255.255.255.0
```

On R1:

```
R1#Show run int S1/4 | Inc ip address

ip address 14.1.1.1 255.255.255.0
```

Once again the sanity check problem is detected between R1 and R4, but remember the task stated that the solution from the previous policy or the next policy can NOT be used. Therefore, to fix this problem using a different solution, PPP should be configured on the serial interfaces, with PPP configured, host routes are created and injected into the routing table, the host routes will point to neighboring router's interface facing the local router, let's configure and verify:

On R1:

```
R1(config)#Int S1/4
R1(config-if)#Encapsulation PPP
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R4:

```
R4(config)#Int S1/1
R4(config-if)#Encapsulation PPP
```

To verify the host route that was injected by configuring PPP:

On R1:

If you don't see the host route, you should "Shut" and "No Shut" the interface.

```
R1#Show ip route | Inc C

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
C        10.1.100.0/24 is directly connected, Loopback100
C        10.1.111.0/24 is directly connected, FastEthernet0/0
C        12.1.1.0/24 is directly connected, Serial1/2
C        14.1.1.0/24 is directly connected, Serial1/4
C        41.1.1.4 is directly connected, Serial1/4
C        200.1.1.0/24 is directly connected, FastEthernet0/1

R1#Ping 41.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 41.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

On R4:

```
R4#Sh ip rou | i C

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
C       24.1.1.0 is directly connected, FastEthernet0/1
C       10.1.4.0 is directly connected, FastEthernet0/0
C       41.1.1.0 is directly connected, Serial1/1
C       46.1.1.0 is directly connected, Serial1/6
C       14.1.1.1 is directly connected, Serial1/1
C       45.1.1.0 is directly connected, Serial1/5

R4#Ping 41.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 41.1.1.4, 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 = 56/57/60 ms
```

Configuring R4:

On R4:

```
R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#Netw 46.0.0.0
R4(config-router)#Netw 45.0.0.0
R4(config-router)#Netw 41.0.0.0
R4(config-router)#Netw 24.0.0.0
R4(config-router)#Netw 10.0.0.0
R4(config-router)#No validate-update-source
```

To verify the configuration:

On R4:

```
R4#Clear ip route *

R4#Show ip route | B Gate
Gateway of last resort is not set

     3.0.0.0/24 is subnetted, 4 subnets
R       3.3.0.0 [120/2] via 14.1.1.1, 00:00:15
R       3.3.1.0 [120/2] via 14.1.1.1, 00:00:15
R       3.3.2.0 [120/2] via 14.1.1.1, 00:00:15
R       3.3.3.0 [120/2] via 14.1.1.1, 00:00:15
     21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R       21.16.0.0/12 [120/2] via 14.1.1.1, 00:00:15
R       21.4.0.0/14 [120/2] via 14.1.1.1, 00:00:15
R       21.1.0.0/16 [120/2] via 14.1.1.1, 00:00:15
R       21.2.0.0/15 [120/2] via 14.1.1.1, 00:00:15
R       21.8.0.0/13 [120/2] via 14.1.1.1, 00:00:15
R       21.32.0.0/11 [120/2] via 14.1.1.1, 00:00:15
R       21.64.0.0/10 [120/2] via 14.1.1.1, 00:00:15
R    200.1.1.0/24 [120/1] via 14.1.1.1, 00:00:15
     113.0.0.0/24 is subnetted, 1 subnets
R       113.1.1.0 [120/2] via 14.1.1.1, 00:00:15
     24.0.0.0/24 is subnetted, 1 subnets
C       24.1.1.0 is directly connected, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
R       10.11.111.111/32 [120/1] via 14.1.1.1, 00:00:16
C       10.1.4.0/24 is directly connected, FastEthernet0/0
R       10.1.111.0/24 [120/1] via 14.1.1.1, 00:00:17
R       10.11.100.0/24 [120/2] via 14.1.1.1, 00:00:17
R       10.1.100.0/24 [120/1] via 14.1.1.1, 00:00:17
     41.0.0.0/24 is subnetted, 1 subnets
C       41.1.1.0 is directly connected, Serial1/1
     12.0.0.0/24 is subnetted, 1 subnets
R       12.1.1.0 [120/1] via 24.1.1.2, 00:00:26, FastEthernet0/1
                 [120/1] via 14.1.1.1, 00:00:17
     46.0.0.0/24 is subnetted, 1 subnets
C       46.1.1.0 is directly connected, Serial1/6
     28.0.0.0/24 is subnetted, 1 subnets
R       28.1.1.0 [120/1] via 24.1.1.2, 00:00:26, FastEthernet0/1
     14.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
R       14.1.1.0/24 [120/2] via 24.1.1.2, 00:00:26, FastEthernet0/1
C       14.1.1.1/32 is directly connected, Serial1/1
     45.0.0.0/24 is subnetted, 1 subnets
C       45.1.1.0 is directly connected, Serial1/5
```

To test the configuration:

On R4:

```
R4#Ping 21.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 21.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

Configuring R5:

On R5:

```
R5(config)#router rip
R5(config-router)#No au
R5(config-router)#ver 2
R5(config-router)#Netw 50.0.0.0
R5(config-router)#Netw 212.1.1.0
R5(config-router)#Netw 211.1.1.0
R5(config-router)#Netw 45.0.0.0
R5(config-router)#Netw 10.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R5:

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

      3.0.0.0/24 is subnetted, 4 subnets
R        3.3.0.0 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        3.3.1.0 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        3.3.2.0 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        3.3.3.0 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C        10.1.1.0/24 is directly connected, FastEthernet0/0
L        10.1.1.5/32 is directly connected, FastEthernet0/0
R        10.1.4.0/24 [120/1] via 45.1.1.4, 00:00:15, Serial1/4
R        10.1.100.0/24 [120/2] via 45.1.1.4, 00:00:15, Serial1/4
R        10.1.111.0/24 [120/2] via 45.1.1.4, 00:00:15, Serial1/4
R        10.11.100.0/24 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        10.11.111.111/32 [120/2] via 45.1.1.4, 00:00:15, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/2] via 45.1.1.4, 00:00:15, Serial1/4
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        21.2.0.0/15 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        21.4.0.0/14 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        21.8.0.0/13 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        21.16.0.0/12 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        21.32.0.0/11 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R        21.64.0.0/10 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
      24.0.0.0/24 is subnetted, 1 subnets
R        24.1.1.0 [120/1] via 45.1.1.4, 00:00:15, Serial1/4
      28.0.0.0/24 is subnetted, 1 subnets
R        28.1.1.0 [120/2] via 45.1.1.4, 00:00:15, Serial1/4
      41.0.0.0/24 is subnetted, 1 subnets
R        41.1.1.0 [120/1] via 45.1.1.4, 00:00:15, Serial1/4
      45.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        45.1.1.0/24 is directly connected, Serial1/4
L        45.1.1.5/32 is directly connected, Serial1/4
      46.0.0.0/24 is subnetted, 1 subnets
R        46.1.1.0 [120/1] via 45.1.1.4, 00:00:15, Serial1/4
      50.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        50.5.5.0/24 is directly connected, Loopback50
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
L        50.5.5.5/32 is directly connected, Loopback50
      113.0.0.0/24 is subnetted, 1 subnets
R        113.1.1.0 [120/3] via 45.1.1.4, 00:00:15, Serial1/4
R     200.1.1.0/24 [120/2] via 45.1.1.4, 00:00:15, Serial1/4
      211.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        211.1.1.0/24 is directly connected, Loopback11
L        211.1.1.5/32 is directly connected, Loopback11
      212.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        212.1.1.0/24 is directly connected, Loopback12
L        212.1.1.5/32 is directly connected, Loopback12
```

To test the configuration:

On R5:

```
R5#Ping 21.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 21.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/116/117 ms
```

Configuring R6:

On R6:

```
R6(config)#router rip
R6(config-router)#No au
R6(config-router)#ver 2
R6(config-router)#Netw 212.1.1.0
R6(config-router)#Netw 211.1.1.0
R6(config-router)#Netw 46.0.0.0
R6(config-router)#Netw 10.0.0.0
```

To verify the configuration:

On R6:

```
R6#Show ip route | B Gate
Gateway of last resort is not set


      3.0.0.0/24 is subnetted, 4 subnets
R        3.3.0.0 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        3.3.1.0 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        3.3.2.0 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        3.3.3.0 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C        10.1.1.0/24 is directly connected, FastEthernet0/0
L        10.1.1.6/32 is directly connected, FastEthernet0/0
R        10.1.4.0/24 [120/1] via 46.1.1.4, 00:00:08, Serial1/4
R        10.1.100.0/24 [120/2] via 46.1.1.4, 00:00:08, Serial1/4
R        10.1.111.0/24 [120/2] via 46.1.1.4, 00:00:08, Serial1/4
R        10.11.100.0/24 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        10.11.111.111/32 [120/2] via 46.1.1.4, 00:00:08, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/2] via 46.1.1.4, 00:00:08, Serial1/4
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        21.2.0.0/15 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        21.4.0.0/14 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        21.8.0.0/13 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        21.16.0.0/12 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        21.32.0.0/11 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R        21.64.0.0/10 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
      24.0.0.0/24 is subnetted, 1 subnets
R        24.1.1.0 [120/1] via 46.1.1.4, 00:00:08, Serial1/4
      28.0.0.0/24 is subnetted, 1 subnets
R        28.1.1.0 [120/2] via 46.1.1.4, 00:00:08, Serial1/4
      41.0.0.0/24 is subnetted, 1 subnets
R        41.1.1.0 [120/1] via 46.1.1.4, 00:00:08, Serial1/4
      45.0.0.0/24 is subnetted, 1 subnets
R        45.1.1.0 [120/1] via 46.1.1.4, 00:00:08, Serial1/4
                  [120/1] via 10.1.1.5, 00:00:17, FastEthernet0/0
      46.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        46.1.1.0/24 is directly connected, Serial1/4
L        46.1.1.6/32 is directly connected, Serial1/4
      50.0.0.0/24 is subnetted, 1 subnets
R        50.5.5.0 [120/1] via 10.1.1.5, 00:00:17, FastEthernet0/0
      113.0.0.0/24 is subnetted, 1 subnets
R        113.1.1.0 [120/3] via 46.1.1.4, 00:00:08, Serial1/4
R     200.1.1.0/24 [120/2] via 46.1.1.4, 00:00:08, Serial1/4
      211.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        211.1.1.0/24 is directly connected, Loopback11
L        211.1.1.6/32 is directly connected, Loopback11
      212.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        212.1.1.0/24 is directly connected, Loopback12
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
L        212.1.1.6/32 is directly connected, Loopback12
```

To test the configuration:

On R6:

```
R6#Ping 21.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 21.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/116/117 ms
```

Configuring R9:

On R9:

```
R9(config)#router rip
R9(config-router)#No au
R9(config-router)#ver 2
R9(config-router)#Network 10.0.0.0
```

To verify the configuration:

On R9:

It may take up to 30 seconds for the routes to be populated:

```
R9#Sh ip rou | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 4 subnets
R        3.3.0.0 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R        3.3.1.0 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R        3.3.2.0 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R        3.3.3.0 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
C        10.1.1.0/24 is directly connected, FastEthernet0/0
L        10.1.1.9/32 is directly connected, FastEthernet0/0
R        10.1.4.0/24 [120/2] via 10.1.1.6, 00:00:09, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                    [120/2] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       10.1.100.0/24 [120/3] via 10.1.1.6, 00:00:09, FastEthernet0/0
                      [120/3] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       10.1.111.0/24 [120/3] via 10.1.1.6, 00:00:09, FastEthernet0/0
                      [120/3] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       10.11.100.0/24 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                       [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       10.11.111.111/32 [120/3] via 10.1.1.6, 00:00:09, FastEthernet0/0
                         [120/3] via 10.1.1.5, 00:00:18, FastEthernet0/0
     12.0.0.0/24 is subnetted, 1 subnets
R       12.1.1.0 [120/3] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/3] via 10.1.1.5, 00:00:18, FastEthernet0/0
     14.0.0.0/24 is subnetted, 1 subnets
R       14.1.1.0 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
     21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R       21.1.0.0/16 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                    [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       21.2.0.0/15 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                    [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       21.4.0.0/14 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                    [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       21.8.0.0/13 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                    [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       21.16.0.0/12 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                     [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       21.32.0.0/11 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                     [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R       21.64.0.0/10 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                     [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
     24.0.0.0/24 is subnetted, 1 subnets
R       24.1.1.0 [120/2] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/2] via 10.1.1.5, 00:00:18, FastEthernet0/0
     28.0.0.0/24 is subnetted, 1 subnets
R       28.1.1.0 [120/3] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/3] via 10.1.1.5, 00:00:18, FastEthernet0/0
     41.0.0.0/24 is subnetted, 1 subnets
R       41.1.1.0 [120/2] via 10.1.1.6, 00:00:09, FastEthernet0/0
                 [120/2] via 10.1.1.5, 00:00:18, FastEthernet0/0
     45.0.0.0/24 is subnetted, 1 subnets
R       45.1.1.0 [120/1] via 10.1.1.5, 00:00:18, FastEthernet0/0
     46.0.0.0/24 is subnetted, 1 subnets
R       46.1.1.0 [120/1] via 10.1.1.6, 00:00:09, FastEthernet0/0
     50.0.0.0/24 is subnetted, 1 subnets
R       50.5.5.0 [120/1] via 10.1.1.5, 00:00:18, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      113.0.0.0/24 is subnetted, 1 subnets
R        113.1.1.0 [120/4] via 10.1.1.6, 00:00:09, FastEthernet0/0
                   [120/4] via 10.1.1.5, 00:00:18, FastEthernet0/0
R     200.1.1.0/24 [120/3] via 10.1.1.6, 00:00:09, FastEthernet0/0
                   [120/3] via 10.1.1.5, 00:00:18, FastEthernet0/0
R     211.1.1.0/24 [120/1] via 10.1.1.6, 00:00:09, FastEthernet0/0
                   [120/1] via 10.1.1.5, 00:00:18, FastEthernet0/0
R     212.1.1.0/24 [120/1] via 10.1.1.6, 00:00:09, FastEthernet0/0
                   [120/1] via 10.1.1.5, 00:00:18, FastEthernet0/0
```

To test the configuration:

On R9:

```
R9#Ping 21.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 21.16.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 116/116/117 ms
```

Configuring R8:

On R8:

```
R8(config)#router rip
R8(config-router)#No au
R8(config-router)#ver 2
R8(config-router)#Network 28.0.0.0
```

To verify the configuration:

On R8:

```
R8#Show ip route | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 4 subnets
R        3.3.0.0 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        3.3.1.0 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        3.3.2.0 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        3.3.3.0 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R        10.1.1.0/24 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        10.1.4.0/24 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        10.1.100.0/24 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        10.1.111.0/24 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        10.11.100.0/24 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        10.11.111.111/32 [120/2] via 28.1.1.2, 00:00:02,
GigabitEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        21.2.0.0/15 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        21.4.0.0/14 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        21.8.0.0/13 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        21.16.0.0/12 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        21.32.0.0/11 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R        21.64.0.0/10 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      24.0.0.0/24 is subnetted, 1 subnets
R        24.1.1.0 [120/1] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      28.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        28.1.1.0/24 is directly connected, GigabitEthernet0/0
L        28.1.1.8/32 is directly connected, GigabitEthernet0/0
      41.0.0.0/24 is subnetted, 1 subnets
R        41.1.1.0 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
R        45.1.1.0 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      46.0.0.0/24 is subnetted, 1 subnets
R        46.1.1.0 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      50.0.0.0/24 is subnetted, 1 subnets
R        50.5.5.0 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
      113.0.0.0/24 is subnetted, 1 subnets
R        113.1.1.0 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R     200.1.1.0/24 [120/2] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R     211.1.1.0/24 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
R     212.1.1.0/24 [120/3] via 28.1.1.2, 00:00:02, GigabitEthernet0/0
```

To test the configuration:

On R8:

```
R8#Ping 21.16.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 21.16.1.1, 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 = 56/58/60 ms
```

Configuring SW1:

The initial configuration file enabled ip routing by entering the "ip routing" global configuration command, if this command is NOT configured, the switches will NOT have the capability to perform routing.

On SW1:

```
SW1(config)#router rip
SW1(config-router)#No au
SW1(config-router)#ver 2
SW1(config-router)#Netw 113.0.0.0
SW1(config-router)#Netw 112.0.0.0
SW1(config-router)#Netw 133.1.0.0
SW1(config-router)#Netw 114.0.0.0
```

To test the configuration:

On SW1:

```
SW1#Show ip route rip

     50.0.0.0/24 is subnetted, 1 subnets
R       50.5.5.0 [120/4] via 113.1.1.3, 00:00:12, FastEthernet0/3
     3.0.0.0/24 is subnetted, 4 subnets
R       3.3.0.0 [120/1] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       3.3.1.0 [120/1] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       3.3.2.0 [120/1] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       3.3.3.0 [120/1] via 113.1.1.3, 00:00:12, FastEthernet0/3
     21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R       21.16.0.0/12 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       21.4.0.0/14 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       21.1.0.0/16 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       21.2.0.0/15 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       21.8.0.0/13 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       21.32.0.0/11 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R       21.64.0.0/10 [120/3] via 113.1.1.3, 00:00:12, FastEthernet0/3
R    200.1.1.0/24 [120/1] via 113.1.1.3, 00:00:12, FastEthernet0/3
R    212.1.1.0/24 [120/4] via 113.1.1.3, 00:00:13, FastEthernet0/3
     24.0.0.0/24 is subnetted, 1 subnets
R       24.1.1.0 [120/3] via 113.1.1.3, 00:00:13, FastEthernet0/3
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R       10.11.111.111/32 [120/2] via 113.1.1.3, 00:00:13, FastEthernet0/3
R       10.1.1.0/24 [120/4] via 113.1.1.3, 00:00:13, FastEthernet0/3
R       10.1.4.0/24 [120/3] via 113.1.1.3, 00:00:13, FastEthernet0/3
R       10.1.111.0/24 [120/2] via 113.1.1.3, 00:00:14, FastEthernet0/3
R       10.11.100.0/24 [120/1] via 113.1.1.3, 00:00:14, FastEthernet0/3
R       10.1.100.0/24 [120/2] via 113.1.1.3, 00:00:14, FastEthernet0/3
     41.0.0.0/24 is subnetted, 1 subnets
R       41.1.1.0 [120/4] via 113.1.1.3, 00:00:14, FastEthernet0/3
     12.0.0.0/24 is subnetted, 1 subnets
R       12.1.1.0 [120/2] via 113.1.1.3, 00:00:14, FastEthernet0/3
     46.0.0.0/24 is subnetted, 1 subnets
R       46.1.1.0 [120/3] via 113.1.1.3, 00:00:15, FastEthernet0/3
     28.0.0.0/24 is subnetted, 1 subnets
R       28.1.1.0 [120/3] via 113.1.1.3, 00:00:15, FastEthernet0/3
     14.0.0.0/24 is subnetted, 1 subnets
R       14.1.1.0 [120/2] via 113.1.1.3, 00:00:15, FastEthernet0/3
R    211.1.1.0/24 [120/4] via 113.1.1.3, 00:00:15, FastEthernet0/3
     45.0.0.0/24 is subnetted, 1 subnets
R       45.1.1.0 [120/3] via 113.1.1.3, 00:00:15, FastEthernet0/3
```

To test the configuration:

On SW1:

```
SW1#Ping 10.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 142/142/143 ms
```

Configuring SW2:

On SW2:

```
SW2(config)#router rip
SW2(config-router)#No au
SW2(config-router)#ver 2
SW2(config-router)#Netw 112.0.0.0
SW2(config-router)#Netw 123.0.0.0
```

To verify the configuration:

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2#Show ip route rip

     50.0.0.0/24 is subnetted, 1 subnets
R       50.5.5.0 [120/5] via 112.1.1.11, 00:00:13, FastEthernet0/19
     3.0.0.0/24 is subnetted, 4 subnets
R       3.3.0.0 [120/2] via 112.1.1.11, 00:00:13, FastEthernet0/19
R       3.3.1.0 [120/2] via 112.1.1.11, 00:00:13, FastEthernet0/19
R       3.3.2.0 [120/2] via 112.1.1.11, 00:00:13, FastEthernet0/19
R       3.3.3.0 [120/2] via 112.1.1.11, 00:00:13, FastEthernet0/19
     21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R       21.16.0.0/12 [120/4] via 112.1.1.11, 00:00:13, FastEthernet0/19
R       21.4.0.0/14 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
R       21.1.0.0/16 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
R       21.2.0.0/15 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
R       21.8.0.0/13 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
R       21.32.0.0/11 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
R       21.64.0.0/10 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
     114.0.0.0/24 is subnetted, 1 subnets
R       114.1.1.0 [120/1] via 112.1.1.11, 00:00:14, FastEthernet0/19
R    200.1.1.0/24 [120/2] via 112.1.1.11, 00:00:14, FastEthernet0/19
     113.0.0.0/24 is subnetted, 1 subnets
R       113.1.1.0 [120/1] via 112.1.1.11, 00:00:14, FastEthernet0/19
R    212.1.1.0/24 [120/5] via 112.1.1.11, 00:00:14, FastEthernet0/19
     24.0.0.0/24 is subnetted, 1 subnets
R       24.1.1.0 [120/4] via 112.1.1.11, 00:00:14, FastEthernet0/19
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R       10.11.111.111/32 [120/3] via 112.1.1.11, 00:00:16,
FastEthernet0/19
R       10.1.1.0/24 [120/5] via 112.1.1.11, 00:00:16, FastEthernet0/19
R       10.1.4.0/24 [120/4] via 112.1.1.11, 00:00:16, FastEthernet0/19
R       10.1.111.0/24 [120/3] via 112.1.1.11, 00:00:16, FastEthernet0/19
R       10.11.100.0/24 [120/2] via 112.1.1.11, 00:00:16, FastEthernet0/19
R       10.1.100.0/24 [120/3] via 112.1.1.11, 00:00:16, FastEthernet0/19
     41.0.0.0/24 is subnetted, 1 subnets
R       41.1.1.0 [120/5] via 112.1.1.11, 00:00:17, FastEthernet0/19
     12.0.0.0/24 is subnetted, 1 subnets
R       12.1.1.0 [120/3] via 112.1.1.11, 00:00:17, FastEthernet0/19
     46.0.0.0/24 is subnetted, 1 subnets
R       46.1.1.0 [120/4] via 112.1.1.11, 00:00:17, FastEthernet0/19
     133.1.0.0/24 is subnetted, 1 subnets
R       133.1.1.0 [120/1] via 112.1.1.11, 00:00:17, FastEthernet0/19
     28.0.0.0/24 is subnetted, 1 subnets
R       28.1.1.0 [120/4] via 112.1.1.11, 00:00:17, FastEthernet0/19
     14.0.0.0/24 is subnetted, 1 subnets
R       14.1.1.0 [120/3] via 112.1.1.11, 00:00:17, FastEthernet0/19
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R    211.1.1.0/24 [120/5] via 112.1.1.11, 00:00:17, FastEthernet0/19
     45.0.0.0/24 is subnetted, 1 subnets
R       45.1.1.0 [120/4] via 112.1.1.11, 00:00:17, FastEthernet0/19
```

To test the configuration:

On SW2:

```
SW2#Ping 10.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 151/157/160 ms
```

Configuring SW3:

On SW3:

```
SW3(config)#router rip
SW3(config-router)#No au
SW3(config-router)#ver 2
SW3(config-router)#Netw 123.0.0.0
SW3(config-router)#Netw 133.1.0.0
SW3(config-router)#Netw 134.1.0.0
```

To verify the configuration:

On SW3:

```
SW3#Sh ip route rip

     50.0.0.0/24 is subnetted, 1 subnets
R       50.5.5.0 [120/5] via 133.1.1.11, 00:00:23, FastEthernet0/21
     3.0.0.0/24 is subnetted, 4 subnets
R       3.3.0.0 [120/2] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       3.3.1.0 [120/2] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       3.3.2.0 [120/2] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       3.3.3.0 [120/2] via 133.1.1.11, 00:00:23, FastEthernet0/21
     21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R       21.16.0.0/12 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       21.4.0.0/14 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       21.1.0.0/16 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       21.2.0.0/15 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R       21.8.0.0/13 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       21.32.0.0/11 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
R       21.64.0.0/10 [120/4] via 133.1.1.11, 00:00:23, FastEthernet0/21
     114.0.0.0/24 is subnetted, 1 subnets
R       114.1.1.0 [120/1] via 133.1.1.11, 00:00:24, FastEthernet0/21
R    200.1.1.0/24 [120/2] via 133.1.1.11, 00:00:24, FastEthernet0/21
     113.0.0.0/24 is subnetted, 1 subnets
R       113.1.1.0 [120/1] via 133.1.1.11, 00:00:24, FastEthernet0/21
     112.0.0.0/24 is subnetted, 1 subnets
R       112.1.1.0 [120/1] via 133.1.1.11, 00:00:24, FastEthernet0/21
                  [120/1] via 123.1.1.22, 00:00:25, FastEthernet0/23
R    212.1.1.0/24 [120/5] via 133.1.1.11, 00:00:26, FastEthernet0/21
     24.0.0.0/24 is subnetted, 1 subnets
R       24.1.1.0 [120/4] via 133.1.1.11, 00:00:26, FastEthernet0/21
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R       10.11.111.111/32 [120/3] via 133.1.1.11, 00:00:26,
FastEthernet0/21
R       10.1.1.0/24 [120/5] via 133.1.1.11, 00:00:26, FastEthernet0/21
R       10.1.4.0/24 [120/4] via 133.1.1.11, 00:00:26, FastEthernet0/21
R       10.1.111.0/24 [120/3] via 133.1.1.11, 00:00:26, FastEthernet0/21
R       10.11.100.0/24 [120/2] via 133.1.1.11, 00:00:26, FastEthernet0/21
R       10.1.100.0/24 [120/3] via 133.1.1.11, 00:00:26, FastEthernet0/21
     41.0.0.0/24 is subnetted, 1 subnets
R       41.1.1.0 [120/5] via 133.1.1.11, 00:00:26, FastEthernet0/21
     12.0.0.0/24 is subnetted, 1 subnets
R       12.1.1.0 [120/3] via 133.1.1.11, 00:00:26, FastEthernet0/21
     46.0.0.0/24 is subnetted, 1 subnets
R       46.1.1.0 [120/4] via 133.1.1.11, 00:00:26, FastEthernet0/21
     28.0.0.0/24 is subnetted, 1 subnets
R       28.1.1.0 [120/4] via 133.1.1.11, 00:00:27, FastEthernet0/21
     14.0.0.0/24 is subnetted, 1 subnets
R       14.1.1.0 [120/3] via 133.1.1.11, 00:00:27, FastEthernet0/21
R    211.1.1.0/24 [120/5] via 133.1.1.11, 00:00:27, FastEthernet0/21
     45.0.0.0/24 is subnetted, 1 subnets
R       45.1.1.0 [120/4] via 133.1.1.11, 00:00:27, FastEthernet0/21
```

To test the configuration:

On SW3:

```
SW3#Ping 10.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.9, 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 = 156/156/160 ms
```

Configuring SW4:

On SW4:

```
SW4(config)#router rip
SW4(config-router)#No au
SW4(config-router)#ver 2
SW4(config-router)#Netw 134.1.0.0
SW4(config-router)#Netw 114.0.0.0
```

To verify the configuration:

On SW4:

```
SW4#Sh ip route rip

     50.0.0.0/24 is subnetted, 1 subnets
R       50.5.5.0 [120/5] via 114.1.1.1, 00:00:12, FastEthernet0/23
     3.0.0.0/24 is subnetted, 4 subnets
R       3.3.0.0 [120/2] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       3.3.1.0 [120/2] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       3.3.2.0 [120/2] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       3.3.3.0 [120/2] via 114.1.1.1, 00:00:12, FastEthernet0/23
     21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R       21.16.0.0/12 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       21.4.0.0/14 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       21.1.0.0/16 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       21.2.0.0/15 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       21.8.0.0/13 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       21.32.0.0/11 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R       21.64.0.0/10 [120/4] via 114.1.1.1, 00:00:12, FastEthernet0/23
R    200.1.1.0/24 [120/2] via 114.1.1.1, 00:00:12, FastEthernet0/23
     113.0.0.0/24 is subnetted, 1 subnets
R       113.1.1.0 [120/1] via 114.1.1.1, 00:00:13, FastEthernet0/23
     112.0.0.0/24 is subnetted, 1 subnets
R       112.1.1.0 [120/1] via 114.1.1.1, 00:00:13, FastEthernet0/23
R    212.1.1.0/24 [120/5] via 114.1.1.1, 00:00:13, FastEthernet0/23
     24.0.0.0/24 is subnetted, 1 subnets
R       24.1.1.0 [120/4] via 114.1.1.1, 00:00:13, FastEthernet0/23
     10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R       10.11.111.111/32 [120/3] via 114.1.1.1, 00:00:14, FastEthernet0/23
R       10.1.1.0/24 [120/5] via 114.1.1.1, 00:00:14, FastEthernet0/23
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R       10.1.4.0/24 [120/4] via 114.1.1.1, 00:00:14, FastEthernet0/23
R       10.1.111.0/24 [120/3] via 114.1.1.1, 00:00:14, FastEthernet0/23
R       10.11.100.0/24 [120/2] via 114.1.1.1, 00:00:14, FastEthernet0/23
R       10.1.100.0/24 [120/3] via 114.1.1.1, 00:00:14, FastEthernet0/23
     41.0.0.0/24 is subnetted, 1 subnets
R       41.1.1.0 [120/5] via 114.1.1.1, 00:00:15, FastEthernet0/23
     123.0.0.0/24 is subnetted, 1 subnets
R       123.1.1.0 [120/1] via 134.1.1.33, 00:00:13, FastEthernet0/19
     12.0.0.0/24 is subnetted, 1 subnets
R       12.1.1.0 [120/3] via 114.1.1.1, 00:00:15, FastEthernet0/23
     46.0.0.0/24 is subnetted, 1 subnets
R       46.1.1.0 [120/4] via 114.1.1.1, 00:00:15, FastEthernet0/23
     133.1.0.0/24 is subnetted, 1 subnets
R       133.1.1.0 [120/1] via 134.1.1.33, 00:00:13, FastEthernet0/19
                  [120/1] via 114.1.1.1, 00:00:16, FastEthernet0/23
     28.0.0.0/24 is subnetted, 1 subnets
R       28.1.1.0 [120/4] via 114.1.1.1, 00:00:16, FastEthernet0/23
     14.0.0.0/24 is subnetted, 1 subnets
R       14.1.1.0 [120/3] via 114.1.1.1, 00:00:16, FastEthernet0/23
R    211.1.1.0/24 [120/5] via 114.1.1.1, 00:00:16, FastEthernet0/23
     45.0.0.0/24 is subnetted, 1 subnets
R       45.1.1.0 [120/4] via 114.1.1.1, 00:00:17, FastEthernet0/23
```

To test the configuration:

On SW4:

```
SW4#Ping 10.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 44/44/44 ms
```

Another way to check and verify which router/s your local router is exchanging routes with is to use the "Show ip protocols" command:

On R1:

```
R1#Show ip protocol | B Gate

    Gateway         Distance      Last Update

    10.11.111.111        120      00:00:13
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    12.1.1.2              120       00:00:24
    41.1.1.4              120       00:00:11
    10.11.100.3           120       01:37:28
    200.1.1.3             120       00:00:01
  Distance: (default is 120)
```

On SW1:

```
SW1#Show ip protocol | B Gate

    Gateway         Distance      Last Update
    114.1.1.44           120      00:00:23
    112.1.1.22           120      00:00:16
    113.1.1.3            120      00:00:25
    133.1.1.33           120      00:00:11

  Distance: (default is 120)
```

On SW2:

```
SW2#Show ip protocol | B Gate

    Gateway         Distance      Last Update
    123.1.1.33           120      00:00:22
    112.1.1.11           120      00:00:11
  Distance: (default is 120)
```

On SW3:

```
SW3#Show ip protocol | B Gate

    Gateway         Distance      Last Update
    123.1.1.22           120      00:00:15
    133.1.1.11           120      00:00:07
    134.1.1.44           120      00:00:00
  Distance: (default is 120)
```

On SW4:

```
SW4#Show ip protocol | B Gate

    Gateway         Distance      Last Update
    114.1.1.1            120      00:00:10
    134.1.1.33           120      00:00:19
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Distance: (default is 120)
```

Task 2

Configure R4 such that it sends RIPv2 updates out of its F0/0 interface to a Broadcast destination, whereas, R1 should be configured to send unicast updates to its neighboring router R7.

RIPv2 sends updates to a multicast destination of 224.0.0.9, this behavior can be changed such that the updates are sent to a broadcast destination, starting with 12.1(5)T the "IP RIP v2-broadcast" interface command will change the destination of the updates from Multicast to broadcast.

On R4:

```
R4(config)#int f0/0
R4(config-if)#ip rip v2-broadcast
```

To verify the configuration:

We can use conditional debugging, if a regular "Debug ip rip" is configured, you may not be able to get access to the console.

```
R4#Debug condition interface F0/0
Condition 1 set

R4#Debug ip pack det
IP packet debugging is on (detailed)

IP: s=10.1.4.4 (local), d=255.255.255.255 (FastEthernet0/0), len 532,
sending broad/multicast
      UDP src=520, dst=520
  IP: s=10.1.4.4 (local), d=255.255.255.255 (FastEthernet0/0), len 532,
sending full packet
      UDP src=520, dst=520
  IP: s=10.1.4.4 (local), d=255.255.255.255 (FastEthernet0/0), len 192,
sending broad/multicast
      UDP src=520, dst=520
  IP: s=10.1.4.4 (local), d=255.255.255.255 (FastEthernet0/0), len 192,
sending full packet
      UDP src=520, dst=520
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

To configure R1 so it sends Unicast updates to R7:

On R1:

```
R1(config)#router rip
R1(config-router)#passive-interface F0/0
R1(config-router)#Neighbor 10.11.111.111
```

Note if the passive-interface is not configured, R1 will send two updates out of its f0/0 interface, one to a multicast destination of 224.0.0.9 and the second to a unicast destination of 10.11.111.111.

To verify the configuration:

On R1:

```
R1#Debug condition interface F0/0
Condition 1 set

R1#Debug ip packet det
IP packet debugging is on (detailed)

IP: s=10.1.111.1 (local), d=10.11.111.111 (FastEthernet0/0), len 532,
sending
      UDP src=520, dst=520
IP: s=10.1.111.1 (local), d=10.11.111.111 (FastEthernet0/0), len 532,
sending full packet
      UDP src=520, dst=520

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

Task 3

Configure the following loopback interfaces on R3; starting with Network 30.3.0.0 /24 this router should ONLY advertise every other network.

Int lo0 = 30.3.0.1 /24	Int lo6 = 30.3.6.1 /24
Int lo1 = 30.3.1.1 /24	Int lo7 = 30.3.7.1 /24

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Int lo2 = 30.3.2.1 /24	Int lo8 = 30.3.8.1 /24
Int lo3 = 30.3.3.1 /24	Int lo9 = 30.3.9.1 /24
Int lo4 = 30.3.4.1 /24	Int lo10 = 30.3.10.1 /24
Int lo5 = 30.3.5.1 /24	

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "RIPv2"→ "Task-3".

On R3:

Note the following access-list denies the odd numbered subnets of 30.3.0.0 and permitting the rest, which includes the even numbered networks. If this was done by permitting the even numbered subnets of 30.3.0.0 network ONLY, an extra access-list statement should have been configured to allow the existing networks; you should NOT forget the future network.

```
R3(config)#access-list 30 deny 30.3.1.0 0.0.254.0
R3(config)#access-list 30 permit any

R3(config)#router rip
R3(config-router)#netw 30.0.0.0
R3(config-router)#distribute-list 30 out S1/1
R3(config-router)#distribute-list 30 out F0/0
R3(config-router)#distribute-list 30 out F0/1
```

On All Routers:

You may have to do the following "Clear ip route *" multiple times to see the result, remember RIP's convergence is VERY slow:

```
Rx#Clear ip route *
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | Inc _30.3.

R       30.3.0.0 [120/1] via 200.1.1.3, 00:00:22, FastEthernet0/1
R       30.3.2.0 [120/1] via 200.1.1.3, 00:00:22, FastEthernet0/1
R       30.3.4.0 [120/1] via 200.1.1.3, 00:00:22, FastEthernet0/1
R       30.3.6.0 [120/1] via 200.1.1.3, 00:00:22, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        30.3.8.0 [120/1] via 200.1.1.3, 00:00:22, FastEthernet0/1
R        30.3.10.0 [120/1] via 200.1.1.3, 00:00:22, FastEthernet0/1
```

On SW1:

```
SW1#Sh ip route rip | Inc _30.3.

R        30.3.8.0 [120/1] via 113.1.1.3, 00:00:11, FastEthernet0/3
R        30.3.10.0 [120/1] via 113.1.1.3, 00:00:11, FastEthernet0/3
R        30.3.4.0 [120/1] via 113.1.1.3, 00:00:11, FastEthernet0/3
R        30.3.6.0 [120/1] via 113.1.1.3, 00:00:11, FastEthernet0/3
R        30.3.0.0 [120/1] via 113.1.1.3, 00:00:11, FastEthernet0/3
R        30.3.2.0 [120/1] via 113.1.1.3, 00:00:11, FastEthernet0/3
```

Task 4

Configure the following loopback interfaces on R6; starting with Network 60.6.0.0 /24 this router should ONLY advertise every 8th Network.

Int lo0 = 60.6.0.1 /24	Int lo6 = 60.6.6.1 /24
Int lo1 = 60.6.1.1 /24	Int lo7 = 60.6.7.1 /24
Int lo2 = 60.6.2.1 /24	Int lo8 = 60.6.8.1 /24
Int lo3 = 60.6.3.1 /24	Int lo9 = 60.6.9.1 /24
Int lo4 = 60.6.4.1 /24	Int lo10 = 60.6.10.1 /24
Int lo5 = 60.6.5.1 /24	

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "RIPv2"→ "Task-4".

On R6:

If "248" which is the third octet in the inverse mask of the access-list is converted to binary, we should have:

Decimal	128	64	32	16	**8**	4	2	1
Binary	**1**	**1**	**1**	**1**	**1**	0	0	0

Since decimal value of 8 is the last position used to represent the decimal value of 248 in binary, the

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

subsequent networks will be increments of 8, if 240 was used the networks will be increments of 16, if 224 was used the networks will be increments of 32 and so forth.

```
R6(config)#access-list 60 permit 60.6.0.0 0.0.248.0
R6(config)#access-list 60 deny 60.6.0.0 0.0.255.255
R6(config)#access-list 60 permit any

R6(config)#router rip
R6(config-router)#distribute-list 60 out S1/4
R6(config-router)#distribute-list 60 out f0/0
R6(config-router)#Network 60.0.0.0

R6#Clear ip route *
```

To verify the configuration:

On R9:

```
R9#Sh ip route rip | Inc 60.6.

R        60.6.8.0 [120/1] via 10.1.56.6, 00:00:25, FastEthernet0/0
R        60.6.0.0 [120/1] via 10.1.56.6, 00:00:25, FastEthernet0/0
```

On R2:

```
R2#Show ip route rip | Inc 60.6.

R        60.6.8.0 [120/2] via 10.1.24.4, 00:00:13, FastEthernet0/1
R        60.6.0.0 [120/2] via 10.1.24.4, 00:00:13, FastEthernet0/1
```

On R1:

```
R1#Show ip route rip | Inc 60.6.

R        60.6.8.0 [120/2] via 10.11.14.4, 00:00:16
R        60.6.0.0 [120/2] via 10.11.14.4, 00:00:16
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 5

Configure the following loopback interfaces on R4; this router should be configured such that it ONLY advertises the even numbered hosts of odd numbered networks of these loopback interfaces plus all other networks.

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "RIPv2"→ "Task-5".

Int lo1	Int lo6
ip addr 40.4.1.1 255.255.255.0 ip addr 40.4.1.2 255.255.255.255 sec ip addr 40.4.1.3 255.255.255.255 sec ip addr 40.4.1.4 255.255.255.255 sec ip addr 40.4.1.5 255.255.255.255 sec ip addr 40.4.1.6 255.255.255.255 sec ip addr 40.4.1.7 255.255.255.255 sec ip addr 40.4.1.8 255.255.255.255 sec ip addr 40.4.1.9 255.255.255.255 sec ip addr 40.4.1.10 255.255.255.255 sec	ip addr 40.4.6.1 255.255.255.0 ip addr 40.4.6.2 255.255.255.255 sec ip addr 40.4.6.3 255.255.255.255 sec ip addr 40.4.6.4 255.255.255.255 sec ip addr 40.4.6.5 255.255.255.255 sec ip addr 40.4.6.6 255.255.255.255 sec ip addr 40.4.6.7 255.255.255.255 sec ip addr 40.4.6.8 255.255.255.255 sec ip addr 40.4.6.9 255.255.255.255 sec ip addr 40.4.6.10 255.255.255.255 sec
Int lo2	**Int lo7**
ip addr 40.4.2.1 255.255.255.0 ip addr 40.4.2.2 255.255.255.255 sec ip addr 40.4.2.3 255.255.255.255 sec ip addr 40.4.2.4 255.255.255.255 sec ip addr 40.4.2.5 255.255.255.255 sec ip addr 40.4.2.6 255.255.255.255 sec ip addr 40.4.2.7 255.255.255.255 sec ip addr 40.4.2.8 255.255.255.255 sec ip addr 40.4.2.9 255.255.255.255 sec ip addr 40.4.2.10 255.255.255.255 sec	ip addr 40.4.7.1 255.255.255.0 ip addr 40.4.7.2 255.255.255.255 sec ip addr 40.4.7.3 255.255.255.255 sec ip addr 40.4.7.4 255.255.255.255 sec ip addr 40.4.7.5 255.255.255.255 sec ip addr 40.4.7.6 255.255.255.255 sec ip addr 40.4.7.7 255.255.255.255 sec ip addr 40.4.7.8 255.255.255.255 sec ip addr 40.4.7.9 255.255.255.255 sec ip addr 40.4.7.10 255.255.255.255 sec
Int lo3	**Int lo8**
ip addr 40.4.3.1 255.255.255.0 ip addr 40.4.3.2 255.255.255.255 sec ip addr 40.4.3.3 255.255.255.255 sec	ip addr 40.4.8.1 255.255.255.0 ip addr 40.4.8.2 255.255.255.255 sec ip addr 40.4.8.3 255.255.255.255 sec

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

ip addr 40.4.3.4 255.255.255.255 sec ip addr 40.4.3.5 255.255.255.255 sec ip addr 40.4.3.6 255.255.255.255 sec ip addr 40.4.3.7 255.255.255.255 sec ip addr 40.4.3.8 255.255.255.255 sec ip addr 40.4.3.9 255.255.255.255 sec ip addr 40.4.3.10 255.255.255.255 sec	ip addr 40.4.8.4 255.255.255.255 sec ip addr 40.4.8.5 255.255.255.255 sec ip addr 40.4.8.6 255.255.255.255 sec ip addr 40.4.8.7 255.255.255.255 sec ip addr 40.4.8.8 255.255.255.255 sec ip addr 40.4.8.9 255.255.255.255 sec ip addr 40.4.8.10 255.255.255.255 sec
Int lo4 ip addr 40.4.4.1 255.255.255.0 ip addr 40.4.4.2 255.255.255.255 sec ip addr 40.4.4.3 255.255.255.255 sec ip addr 40.4.4.4 255.255.255.255 sec ip addr 40.4.4.5 255.255.255.255 sec ip addr 40.4.4.6 255.255.255.255 sec ip addr 40.4.4.7 255.255.255.255 sec ip addr 40.4.4.8 255.255.255.255 sec ip addr 40.4.4.9 255.255.255.255 sec ip addr 40.4.4.10 255.255.255.255 sec	**Int lo9** ip addr 40.4.9.1 255.255.255.0 ip addr 40.4.9.2 255.255.255.255 sec ip addr 40.4.9.3 255.255.255.255 sec ip addr 40.4.9.4 255.255.255.255 sec ip addr 40.4.9.5 255.255.255.255 sec ip addr 40.4.9.6 255.255.255.255 sec ip addr 40.4.9.7 255.255.255.255 sec ip addr 40.4.9.8 255.255.255.255 sec ip addr 40.4.9.9 255.255.255.255 sec ip addr 40.4.9.10 255.255.255.255 sec
Int lo5 ip addr 40.4.5.1 255.255.255.0 ip addr 40.4.5.2 255.255.255.255 sec ip addr 40.4.5.3 255.255.255.255 sec ip addr 40.4.5.4 255.255.255.255 sec ip addr 40.4.5.5 255.255.255.255 sec ip addr 40.4.5.6 255.255.255.255 sec ip addr 40.4.5.7 255.255.255.255 sec ip addr 40.4.5.8 255.255.255.255 sec ip addr 40.4.5.9 255.255.255.255 sec ip addr 40.4.5.10 255.255.255.255 sec	**Int lo10** ip addr 40.4.10.1 255.255.255.0 ip addr 40.4.10.2 255.255.255.255 sec ip addr 40.4.10.3 255.255.255.255 sec ip addr 40.4.10.4 255.255.255.255 sec ip addr 40.4.10.5 255.255.255.255 sec ip addr 40.4.10.6 255.255.255.255 sec ip addr 40.4.10.7 255.255.255.255 sec ip addr 40.4.10.8 255.255.255.255 sec ip addr 40.4.10.9 255.255.255.255 sec ip addr 40.4.10.10 255.255.255.255 sec

On R4:

```
R4(config)#access-list 40 permit 40.4.1.0 0.0.254.254
R4(config)#access-list 40 deny 40.4.0.0 0.0.255.255
R4(config)#access-list 40 permit any

R4(config)#router rip
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config-router)#netw 40.0.0.0
R4(config-router)#distribute-list 40 out s1/5
R4(config-router)#distribute-list 40 out s1/6
R4(config-router)#distribute-list 40 out f0/0
R4(config-router)#distribute-list 40 out s1/1
R4(config-router)#distribute-list 40 out f0/1

R4#Clear ip route *
```

To verify the configuration:

On R1:

```
R1#Clear ip route *

R1#Show ip route rip | inc 40.4.

R        40.4.1.0/24 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.1.2/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.1.4/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.1.6/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.1.8/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.1.10/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.3.0/24 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.3.2/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.3.4/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.3.6/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.3.8/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.3.10/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.5.0/24 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.5.2/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.5.4/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.5.6/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.5.8/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.5.10/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.7.0/24 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.7.2/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.7.4/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.7.6/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.7.8/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.7.10/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.9.0/24 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.9.2/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.9.4/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.9.6/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R        40.4.9.8/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
R        40.4.9.10/32 [120/2] via 12.1.1.2, 00:00:18, Serial1/2
```

Task 6

Configure the routers in this routing domain based on the following timers:

- Periodic updates every 30 seconds
- They should declare a route invalid after 1.5 minutes
- They should suppress routing information regarding a better path for 1.5 minutes
- They should flush routes 30 seconds after they are invalid
- They should postpone their periodic updates by 100 milliseconds

On All Routers:

```
(config)#router rip
(config-router)#timers basic 30 90 90 120 100
```

To verify the configuration:

On Any Router:

```
R7#Show ip prot | Inc Send|Inv

Sending updates every 30 seconds, next due in 27 seconds
  Invalid after 90 seconds, hold down 90, flushed after 100
    Interface              Send  Recv  Triggered RIP  Key-chain
```

Task 7

Since R1 is a very fast router, it should be configured such that it adds an interpacket delay of 50 milliseconds between the update packets.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#router rip
R1(config-router)#output-delay 50
```

Task 8

Configure R2 to set the number of received but NOT yet processed RIP update packets in the RIP input queue to 100.

On R2:

```
R2(config)#router rip

R2(config-router)#input-queue 100
```

Task 9

Configure all routers to suppress a flash update when a topology change occurs 10 seconds before a regular update.

On All routers:

```
(config)#router rip
(config-router)#flash-update-threshold 10
```

To verify the configuration:

On Any Router:

```
#Show ip protocol | Inc Flash

  Flash update is suppressed when next update due within 10 seconds
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 10

R1 and R4 should be configured to send event driven updates and NOT periodic updates.

On R1:

```
R1(config)#int s1/4
R1(config-if)#ip rip triggered
```

On R4:

```
R4(config)#int s1/1
R4(config-if)#ip rip triggered
```

To verify the configuration:

On R4:

```
R4#Show ip protocols | Inc Trigg|Serial1/1

    Serial1/1 filtered by 40, default is 40
    Interface              Send  Recv  Triggered RIP  Key-chain
    Serial1/1              2     2
```

On R1:

```
R1#Show ip protocols | Inc Trigg|Serial1/4

    Interface              Send  Recv  Triggered RIP  Key-chain
    Serial1/4              2     2           Yes
```

Task 11

R1 and R2 should be configured to authenticate their routing updates through their Serial connection; these two routers should use an unencrypted key of "**ccie**" for this purpose.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Key chain R1-R2
R1(config-keychain)#key 1
R1(config-keychain-key)#key-string ccie

R1(config-keychain-key)#int S1/2
R1(config-subif)#ip rip authentication key-chain R1-R2
```

On R2:

```
R2(config)#key chain R2-R1
R2(config-keychain)#key 1
R2(config-keychain-key)#key-string ccie

R2(config-keychain-key)#int S1/1
R2(config-subif)#ip rip authentication key-chain R2-R1
```

To verify the configuration:

On R1:

```
R1#Show ip protocol | Inc Serial1/2|Key

    Interface              Send  Recv  Triggered RIP  Key-chain
    Serial1/2              2     2                   R1-R2
```

On R2:

```
R2#Show ip protocol | Inc Serial1/1|Key

    Interface              Send  Recv  Triggered RIP  Key-chain
    Serial1/1              2     2                   R2-R1
```

Task 12

Configure R5, R6 and R9 to use authentication using the strongest authentication method available to RIPv2, these routers should use "**Micronic?**" as their password.

On R5:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5(config)#key chain R5-R6-R9
R5(config-keychain)#key 1
R5(config-keychain-key)#key-string Micronic?

R5(config-keychain-key)#int f0/0
R5(config-if)#ip rip authentication key-chain R5-R6-R9
R5(config-if)#ip rip authentication mode md5
```

On R6:

```
R6(config)#key chain R5-R6-R9
R6(config-keychain)#key 1
R6(config-keychain-key)#key-string Micronic?

R6(config-keychain-key)#int f0/0
R6(config-if)#ip rip authentication key-chain R5-R6-R9
R6(config-if)#ip rip authentication mode md5
```

On R9:

```
R9(config)#key chain R5-R6-R9
R9(config-keychain)#key 1
R9(config-keychain-key)#key-string Micronic?

R9(config-keychain-key)#int f0/0
R9(config-if)#ip rip authentication key-chain R5-R6-R9
R9(config-if)#ip rip authentication mode md5
```

To verify the configuration:

On R5:

```
R5#Show key chain R5-R6-R9 | Inc key 1

    key 1 -- text "Micronic?"
```

On R6:

```
R6#Show key chain R5-R6-R9 | Inc key 1

    key 1 -- text "Micronic?"
```

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9#Show key chain R5-R6-R9 | Inc key 1

    key 1 -- text "Micronic?"
```

Task 13

Configure authentication on the link between R3 and SW1 such that the output of the "debug ip rip" on these two devices resembles the following, these two routers should successfully exchange updates:

On SW1:

RIP: received packet with text authentication R3
RIP: received v2 update from 10.1.113.3 on FastEthernet0/3

On R3:

RIP: received packet with text authentication SW1
RIP: received v2 update from 10.1.113.1 on FastEthernet0/0

Typically commands and configuration in Cisco IOS are NOT order specific; RIPv2's authentication is an exception to this rule, meaning that it is order specific. The key chain must be configured before it is applied, or else authentication will NOT work.

On R3:

```
R3(config)#key chain TST
R3(config-keychain)#key 1
R3(config-keychain-key)#key-string R3
R3(config-keychain-key)#send-lifetime 10:00:00 jan 1 2009 in

R3(config-keychain-key)#key 2
R3(config-keychain-key)#key-string SW1
R3(config-keychain-key)#accept-lifetime 10:00:00 jan 1 2009 inf

R3(config-keychain-key)#int F0/0
R3(config-if)#ip rip authentication key TST
```

On SW1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1(config)#key chain TST
SW1(config-keychain)#key 2
SW1(config-keychain-key)#key-string R3
SW1(config-keychain-key)#accept-lifetime 10:00:00 jan 1 2009 in

SW1(config-keychain-key)#key 1
SW1(config-keychain-key)#key-string SW1
SW1(config-keychain-key)#send-lifetime 10:00:00 jan 1 2009 inf

SW1(config-keychain-key)#int f0/3
SW1(config-if)#ip rip authentication key TST
```

On Both Devices:

The following sets the clock on both devices:

```
#Clock set 09:55:00 jan 1 2009
```

SHUTDOWN THE OTHER INTERFACES ON THESE TWO ROUTERS BEFORE PERFORMING THE DEBUG, OR ELSE YOU WON'T BE ABLE TO CONSOLE INTO ANY OF THE ROUTERS

On SW1:

```
SW1#Debug ip rip

RIP: received packet with text authentication R3
RIP: received v2 update from 10.1.113.3 on FastEthernet0/3
```

On R3:

```
R3#Debug ip rip

RIP: received packet with text authentication SW1
RIP: received v2 update from 10.1.113.1 on FastEthernet0/0
```

Task 14

R1 should be configured to accept existing and future routes that have a prefix-length of 10 to 14, these routes should be received from R7 and R7 ONLY. DO NOT use an access-list/s or a neighbor command to accomplish this task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To resolve this task two prefix-lists are configured, the first one references all the routes with prefix-length of 10 to 14, and the second one references R7's IP address; A distribute-list is configured to reference the prefix-list that identifies the desired prefixes and the gateway option is used to reference R7.

On R1:

```
R1#Sh ip route rip | Inc 10.11.111.111

R       21.16.0.0/12 [120/1] via 10.11.111.111, 00:00:10
R       21.4.0.0/14 [120/1] via 10.11.111.111, 00:00:10
R       21.1.0.0/16 [120/1] via 10.11.111.111, 00:00:10
R       21.2.0.0/15 [120/1] via 10.11.111.111, 00:00:10
R       21.8.0.0/13 [120/1] via 10.11.111.111, 00:00:10
R       21.32.0.0/11 [120/1] via 10.11.111.111, 00:00:10
R       21.64.0.0/10 [120/1] via 10.11.111.111, 00:00:10

R1(config)#ip prefix-list NET permit 0.0.0.0/0 GE 10 LE 14

R1(config)#ip prefix-list R7 permit 10.11.111.111/32

R1(config)#router rip
R1(config-router)#distribute-list prefix NET gateway R7 in F0/0
```

To verify the configuration:

On R1:

```
R1#Clear ip rou *

R1#Show ip route rip | Inc 10.11.111.111

R       21.16.0.0/12 [120/1] via 10.11.111.111, 00:00:09
R       21.4.0.0/14 [120/1] via 10.11.111.111, 00:00:09
R       21.8.0.0/13 [120/1] via 10.11.111.111, 00:00:09
R       21.32.0.0/11 [120/1] via 10.11.111.111, 00:00:09
R       21.64.0.0/10 [120/1] via 10.11.111.111, 00:00:09
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 15

Configure R7 to inject a default route ONLY if its Lo15 and Lo16 are both UP, this router should **NOT** inject a default route if any of these two interfaces are down.

On R7:

Step #1

A boolean logic is configured identifying 2 objects, object 1 and object 2:

```
R7(config)#track 3 list boolean and
R7(config-track)#object 1
R7(config-track)#object 2
```

Step #2

The two objects are defined:

```
R7(config-track)#track 1 ip route 21.1.0.0 255.255.0.0 reach
R7(config-track)#track 2 ip route 21.2.0.0 255.254.0.0 reach
```

Step #3

A fake route is configured, this prefix DOES NOT exist in the network, it's configured so it could be tracked and matched, you can pick any prefix:

```
R7(config)#ip route 211.211.211.211 255.255.255.255 null0 track 3
```

Step #4

The fake prefix is referenced by the following prefix-list:

```
R7(config)#ip prefix-list NET permit 211.211.211.211/32
```

Step #5

The following route-map references the prefix-list from the previous step:

```
R7(config)#route-map TST permit 10
R7(config-route-map)#match ip addr prefix NET
```

Step #6

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The default-information originate is configured to reference the route-map:

```
R7(config-route-map)#router rip
R7(config-router)#default-information originate route-map TST
```

In this configuration, if the state of the boolean is UP, meaning that the two networks referenced by the two objects are in this router's routing table, then, the fake route will NOT be removed, which means that the route-map will be able to reference the fake route, and as a result of that a default route is injected, however, if the state of the boolean is down, the fake route is removed, which means that the route-map can NOT match the fake route, and as a result of that, the default route is NOT injected:

To verify the configuration:

On R7:

```
R7#Sh track

Track 1
  IP route 21.1.0.0 255.255.0.0 reachability
  Reachability is Up (connected)

    1 change, last change 00:15:22
  First-hop interface is Loopback16

Track 2
  IP route 21.2.0.0 255.254.0.0 reachability
  Reachability is Up (connected)
    1 changes, last change 00:02:37
  First-hop interface is Loopback15

Track 3
  List boolean and
  Boolean AND is Up
    1 changes, last change 00:02:36
    object 1 Up
    object 2 Up
  Tracked by:
    STATIC-IP-ROUTING 0
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Before the default route can be verified on R1, the following prefix-list must be added to the existing prefix-list that was configured in the previous task, remember in the previous task R1 was configured to ONLY allow prefixes that had a prefix-length of 10 – 14 coming from R7, if the following prefix-list is NOT added, the default route will be filtered and it will NOT be injected:

```
R1(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0

R1#Show ip route rip | Inc 0.0.0.0/0

R*   0.0.0.0/0 [120/1] via 10.11.111.111, 00:00:05
```

To test the configuration:

On R7:

```
R7(config)#int lo15
R7(config-if)#shut
```

Note you should receive the following console messages:

```
%LINK-5-CHANGED: Interface Loopback15, changed state to
administratively down

%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback15, changed
state to down
%TRACK-6-STATE: 2 ip route 21.2.0.0/15 reachability Up -> Down
%TRACK-6-STATE: 3 list boolean and Up -> Down
```

To verify the state of the Boolean:

```
R7#Show track

Track 1
  IP route 21.1.0.0 255.255.0.0 reachability
  Reachability is Up (connected)
    1 change, last change 00:09:32
  First-hop interface is Loopback16

Track 2
  IP route 21.2.0.0 255.254.0.0 reachability
  Reachability is Down (no route)
    2 changes, last change 00:02:17
  First-hop interface is unknown
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Track 3
  List boolean and
  Boolean AND is Down
    3 changes, last change 00:02:17
    object 1 Up
    object 2 Down

  Tracked by:
    STATIC-IP-ROUTING 0
```

To verify the removal of the default route:

On R1:

```
R1#Show ip route rip | Inc 0.0.0.0/0
R1#
```

Note the default route is removed. To re-inject the default route:

On R7:

```
R7(config)#int lo15
R7(config-if)#No shut
```

Once again you should see the following console messages:

```
%LINK-3-UPDOWN: Interface Loopback15, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback15, changed
state to up
%TRACK-6-STATE: 2 ip route 21.2.0.0/15 reachability Down -> Up
%TRACK-6-STATE: 3 list boolean and Down -> Up
```

To verify the result:

On R1:

```
R1#Show ip route rip | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/1] via 10.11.111.111, 00:00:18
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 18

Configure Lo0 and Lo1 interfaces on R8, assign 112.11.1.1 /24 and 112.12.2.2 /24 IP addresses to these two loopback interfaces respectively. R8 should be configured to inject a default route as long as one of these 2 interfaces is up. If both interfaces are down, then, the default route should be removed. **The condition stated in this task should be fully tested and verified.**

On R8:

```
R8(config)#int lo0
R8(config-if)#ip addr 112.11.1.1 255.255.255.0

R8(config-if)#int lo1
R8(config-if)#ip addr 112.12.2.2 255.255.255.0


R8(config)#access-list 1 permit 112.11.1.0 0.0.0.255
R8(config)#access-list 1 permit 112.12.2.0 0.0.0.255

R8(config)#route-map TST permit 10
R8(config-route-map)#match ip addr 1

R8(config-route-map)#router rip

R8(config-router)#default-information originate route-map TST
```

To verify the configuration:

On R2:

```
R2#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 12.1.1.1, 00:00:08, Serial1/1
```

Note the default route is injected.

To test the configuration:

On R8:

```
R8(config)#int lo0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R8(config-if)#Shut
```

On R2:

The default route is injected because Lo1 is up.

```
R2#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 12.1.1.1, 00:00:15, Serial1/1
```

On R8:

```
R8(config-if)#int lo1
R8(config-if)#Shut
```

On R2:

Note the following default route is injected by R7 from the previous task and NOT R8: Remember that convergence in RIPv2 is slow.

```
R2#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 12.1.1.1, 00:00:17, Serial1/1
```

To enable the loopback interfaces on R8:

On R8:

```
R8(config-if)#int lo0
R8(config-if)#No Shut

R8(config-if)#Int lo1
R8(config-if)#No Shut
```

On R2:

To speed up convergence:

```
R2#Clear ip route *
```

You should see the following default route being injected by R4:

```
R2#Sh ip route rip | Inc 0.0.0.0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R*    0.0.0.0/0 [120/3] via 24.1.1.4, 00:00:17, FastEthernet0/1
```

But the next Show ip route command on R2 shows that R2 is getting the default route from R1:

```
R2#Sh ip rou rip | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 12.1.1.1, 00:00:16, Serial1/1
```

Note R2 is taking the default route injected by R7 through R1, even though it has a higher hop count. This is because R8 is no longer advertising the default route, the question is why?
This is because when a router receives a given route through an interface, it will never advertise that same route out of that interface (Split-horizon). R2 was advertising the default route that it received from R1 to R8 and because of the split-horizon rule, R8 does NOT advertise its own default route.

In order to complete this task so all the tests pass, R8 should be configured to deny the default route coming in through its F0/0 interface, so it can advertise the default route that its configured to advertise out of that same interface (F0/0):

On R8:

```
R8#Show ip route rip | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/3] via 28.1.1.2, 00:00:28, GigabitEthernet0/0
```

To filter the default route:

```
R8(config)#ip prefix-list Default seq 5 deny 0.0.0.0/0
R8(config)#ip prefix-list Default seq 10 permit 0.0.0.0/0 LE 32

R8(config)#router rip
R8(config-router)#distribute-list prefix Default in g0/0
```

To test the configuration:

On R2:

```
R2#Clear ip route *

R2#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/3] via 24.1.1.4, 00:00:10, FastEthernet0/1
```

To test the final configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R8:

```
R8(config-if)#int lo0
R8(config-if)#No shut

R8(config-if)#int lo1
R8(config-if)#No shut
```

On R2:

```
R2#Clear ip route *

R2#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 12.1.1.1, 00:00:06, Serial1/1
```

Task 19

Configure R4 to filter any default routes injected by either R7 or R8.

On R4:

```
R4(config)#ip prefix-list TST seq 5 deny 0.0.0.0/0
R4(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R4(config)#router rip
R4(config-router)#distribute-list prefix TST in F0/1
R4(config-router)#distribute-list prefix TST in F0/0
```

To verify the configuration:

```
R4#Clear ip route *

R4#Show ip route rip | Inc 0.0.0.0/0
R4#
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 20

Configure Lo0 and Lo1 interfaces on R9 and assign 113.1.1.3 /24 and 113.2.2.3 /24 IP addresses to these two interfaces respectively.
R9 should be configured to inject a default route ONLY **if Lo0 is UP** and **Lo1 is in DOWN state**, this task should be fully tested and verified. DO NOT use a fake route to accomplish this task.

On R9:

```
R9(config)#int lo0
R9(config-if)#ip addr 113.1.1.3 255.255.255.0

R9(config-if)#int lo1
R9(config-if)#ip addr 113.2.2.3 255.255.255.0

R9(config)#track 3 list boolean and
R9(config-track)#object 1
R9(config-track)#object 2 NOT

R9(config-track)#track 1 ip route 113.1.1.0 255.255.255.0 reach

*Mar  4 07:02:56.488: %TRACKING-5-STATE: 3 list boolean and Down->Up

R9(config-track)#track 2 ip route 113.2.2.0 255.255.255.0 reach

*Mar  4 07:03:11.513: %TRACKING-5-STATE: 3 list boolean and Up->Down

R9(config)#ip route 0.0.0.0 0.0.0.0 null0 track 3

R9(config)#ip prefix-list TST permit 0.0.0.0/0

R9(config)#route-map TST permit 10
R9(config-route-map)#match ip addr prefix TST

R9(config-route-map)#router rip
R9(config-router)#default-information originate route-map TST
```

To verify the configuration:

On R9:

```
R9(config-if)#int lo1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9(config-if)#Shut
```

You should see the following console message:

```
%TRACKING-5-STATE: 2 ip route 113.2.2.0/24 reachability Up->Down

%TRACKING-5-STATE: 3 list boolean and Down->Up
%LINK-5-CHANGED: Interface Loopback1, changed state to
administratively down
%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback1, changed
state to down
```

To verify the configuration:

On R5:

```
R5#Show ip route rip | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/1] via 10.1.1.9, 00:00:16, FastEthernet0/0
```

On R6:

```
R6#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/1] via 10.1.1.9, 00:00:08, FastEthernet0/0
```

Task 21

Configure R5 and R6 to filter the default route received from R9 but allow the rest of the prefixes in their routing table.

On R5 and R6:

```
Rx(config)#ip prefix-list Default seq 5 deny 0.0.0.0/0
Rx(config)#ip prefix-list Default seq 10 permit 0.0.0.0/0 LE 32

Rx(config)#router rip
Rx(config-router)#distribute-list prefix Default in F0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5#Clear ip route *
```

To test the configuration:

On R5:

```
R5#Show ip route rip | Inc 0.0.0.0/0
R5#
```

On R6:

```
R6#Sh ip route rip | inc 0.0.0.0/0
R6#
```

Task 22

R9 should be configured such that it always prefers to reach network 200.1.1.0 /24 through R6, R9 should use R5 ONLY when R6 is down. DO NOT use offset-list to accomplish this task.

On R9:

```
R9#Sh ip route 200.1.1.0

Routing entry for 200.1.1.0/24
  Known via "rip", distance 120, metric 4

  Redistributing via rip
  Last update from 10.1.1.6 on FastEthernet0/0, 00:00:00 ago
  Routing Descriptor Blocks:
    10.1.1.6, from 10.1.1.6, 00:00:00 ago, via FastEthernet0/0
      Route metric is 4, traffic share count is 1
  * 10.1.1.5, from 10.1.1.5, 00:00:16 ago, via FastEthernet0/0
      Route metric is 4, traffic share count is 1
```

To configure the task:

```
R9(config)#router rip
R9(config-router)#distance 121 10.1.1.5 0.0.0.0 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9(config)#access-list 1 permit 200.1.1.0 0.0.0.255

R9#clear ip route *
```

To verify the configuration:

On R9:

```
R9#Show ip route 200.1.1.0

Routing entry for 200.1.1.0/24
  Known via "rip", distance 120, metric 4
  Redistributing via rip
  Last update from 10.1.1.6 on FastEthernet0/0, 00:00:09 ago
  Routing Descriptor Blocks:
  * 10.1.1.6, from 10.1.1.6, 00:00:09 ago, via FastEthernet0/0
      Route metric is 4, traffic share count is 1
```

To test the configuration:

On R6:

```
R6(config)#Int F0/0
R6(config-if)#Shut
```

On R9:

```
R9#Clear ip route *

R9#Show ip route 200.1.1.0

Routing entry for 200.1.1.0/24
  Known via "rip", distance 121, metric 4
  Redistributing via rip
  Last update from 10.1.1.5 on FastEthernet0/0, 00:00:01 ago
  Routing Descriptor Blocks:
  * 10.1.1.5, from 10.1.1.5, 00:00:01 ago, via FastEthernet0/0
      Route metric is 4, traffic share count is 1
```

Note R9 is choosing R5 because R6 is down, to see the administrative distance:

```
R9#Show ip route rip | Inc 200.1.1.0

R     200.1.1.0/24 [121/4] via 10.1.1.5, 00:00:24, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R6:

```
R6(config)#Int F0/0
R6(config-if)#No Shut
```

Task 23

R9 should be configured to filter network 50.5.5.0 /24. DO NOT use an access-list to accomplish this task.

On R9:

```
R9(config)#ip prefix-list NET-50 deny 50.5.5.0/24
R9(config)#ip prefix-list NET-50 permit 0.0.0.0/0 le 32

R9(config)#router rip
R9(config-router)#distribute-list prefix NET-50 in F0/0
```

To verify the configuration:

On R9:

```
R9#Clear ip route *


R9#Sh ip route rip | Inc 50.5.5.0
R9#
```

Task 24

Configure R4 such that it injects a default route into RIPv2 routing domain; this default route should be given ONLY to R5 and NOT R6. DO NOT configure R6 to accomplish this task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R4:

```
R4(config)#router rip
R4(config-router)#default-information originate route-map TST

R4(config)#route-map TST permit 10
R4(config-route-map)#Set interface S1/5

R4#Clear ip route *
```

To test the configuration:

On R5:

```
R5#Sh ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/1] via 45.1.1.4, 00:00:14, Serial1/4
```

On R6:

```
R6#Sh ip route rip | inc 0.0.0.0/0
R6#
```

Task 25

R3 should be configured to summarize its Loopback Lo0 – Lo3, 30 – 33 and advertise two summary routes into the RIP routing domain.

On R3:

```
R3(config)#Int S1/1
R3(config-subif)#ip summary-address rip 3.3.0.0 255.255.252.0
R3(config-subif)#ip summary-address rip 30.3.0.0 255.255.252.0

R3(config-subif)#Int f0/1
R3(config-if)#ip summary-address rip 3.3.0.0 255.255.252.0
R3(config-if)#ip summary-address rip 30.3.0.0 255.255.252.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config-if)#int f0/0
R3(config-if)#ip summary-address rip 3.3.0.0 255.255.252.0
R3(config-if)#ip summary-address rip 30.3.0.0 255.255.252.0
```

Task 26

Configure Lo200 with an IP address of 120.2.2.2 /24 on SW2; this switch should advertise this Network in RIPv2 routing domain.
SW1 should be configured such that this network is never advertised to any router beyond R9.

On SW2:

```
SW2(config)#int lo200
SW2(config-if)#ip addr 120.2.2.2 255.255.255.0

SW2(config-if)#router rip
SW2(config-router)#netw 120.0.0.0
```

To verify the configuration:

On SW1:

```
SW1#Show ip route rip | Inc 120.2.2.0

R        120.2.2.0 [120/1] via 112.1.1.22, 00:00:01, FastEthernet0/19
```

To configure SW1:

```
SW1(config)#access-list 1 permit 120.2.2.0 0.0.0.255

SW1(config)#router rip
SW1(config-router)#offset-list 1 out 8 F0/3

SW1#Clear ip route *
```

To verify the configuration:

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9#Sh ip route rip | Inc 120.2.2.0

R        120.2.2.0 [120/15] via 10.1.1.6, 00:00:04, FastEthernet0/0
```

Note R9 will advertise this route with a hop count of 16 (Adding one to what was advertised to it), and the receiving RIP speaker will reject the route because of its Hop count exceeding the maximum allowable hop count.

Task 27

Configure F0/1 interface of R9 with an IP address of 203.1.1.3 /24 and advertise this interface in RIPv2; configure the appropriate device/s such that R9's next-hop IP address for subnets of 21.0.0.0 network and the default route (advertised by R7) is 10.11.111.111 (R7's F0/0 interface).
Do not use a global configuration mode on any router, GRE, IPnIP, LDP, or MPLS to accomplish this task.
R9 may not have reachability to any of the IP addresses resulting from this task.

On R9:

```
R9#Sh ip rou rip | Inc _21.

      21.0.0.0/8 is variably subnetted, 5 subnets, 5 masks
R        21.4.0.0/14 [120/4] via 10.1.1.6, 00:00:21, FastEthernet0/0
R        21.8.0.0/13 [120/4] via 10.1.1.6, 00:00:21, FastEthernet0/0
R        21.16.0.0/12 [120/4] via 10.1.1.6, 00:00:21, FastEthernet0/0
R        21.32.0.0/11 [120/4] via 10.1.1.6, 00:00:21, FastEthernet0/0
R        21.64.0.0/10 [120/4] via 10.1.1.6, 00:00:21, FastEthernet0/0
R     211.1.1.0/24 [120/1] via 10.1.1.6, 00:00:21, FastEthernet0/0

R     212.1.1.0/24 [120/1] via 10.1.1.6, 00:00:21, FastEthernet0/0
```

In order to configure this task, SPAN/RSPAN is used on the switches; basically SPAN/RSPAN takes the traffic from one or more interface/s and forwards it to another interface on the same switch (SPAN) or another switch (RSPAN). Since the two routers in this topology are connected to two different switches, RSPAN is used to accomplish this task.

On SW1:

```
SW1(config)#vlan 99
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1(config-vlan)#remote-span

SW1(config)#monitor session 1 source interface f0/7 both
SW1(config)#monitor session 1 destination remote vlan 99

SW1(config)#int f0/22
SW1(config-if)#swi tru enc dot
SW1(config-if)#swi mode trunk
SW1(config-if)#No shut
```

On SW3:

```
SW3(config)#vtp domain RIPv2
Changing VTP domain name from NULL to RIPv2

SW3(config)#int f0/22
SW3(config-if)#swi tru enc dot
SW3(config-if)#swi mode trunk
SW3(config-if)#No shut
```

If you don't see the following VLAN on SW3, you should create a VLAN on one of the switches (SW1 or SW3) so the VTP can synchronise the VLAN DB.

```
SW3#Show vlan remote-span

Remote SPAN VLANs
------------------------------------------------------------------------------
99

SW3(config)#monitor session 1 source remote vlan 99
SW3(config)#monitor session 1 destination interface f0/9
```

On R9:

```
R9(config)#int F0/1
R9(config-if)#ip addr 203.1.1.3 255.255.255.0
R9(config-if)#No Shut

R9(config-if)#router rip
R9(config-router)#netw 203.1.1.0
R9(config-router)#No validate-update-source
```

To verify the configuration:

```
R9#Sh ip rou rip | Inc _21.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      21.0.0.0/8 is variably subnetted, 7 subnets, 7 masks
R        21.1.0.0/16 [120/1] via 10.11.111.111, 00:00:00
R        21.2.0.0/15 [120/1] via 10.11.111.111, 00:00:00
R        21.4.0.0/14 [120/1] via 10.11.111.111, 00:00:00
R        21.8.0.0/13 [120/1] via 10.11.111.111, 00:00:00
R        21.16.0.0/12 [120/1] via 10.11.111.111, 00:00:00
R        21.32.0.0/11 [120/1] via 10.11.111.111, 00:00:00
R        21.64.0.0/10 [120/1] via 10.11.111.111, 00:00:00
R     211.1.1.0/24 [120/1] via 10.1.1.6, 00:00:01, FastEthernet0/0
R     212.1.1.0/24 [120/1] via 10.1.1.6, 00:00:01, FastEthernet0/0
```

Task 28

Erase the startup configuration of the routers, startup configuration and the VLAN.dat of the switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Advanced
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

EIGRP:

Lab-1:
EIGRP Named Configuration

Lab-2:
EIGRP Stub

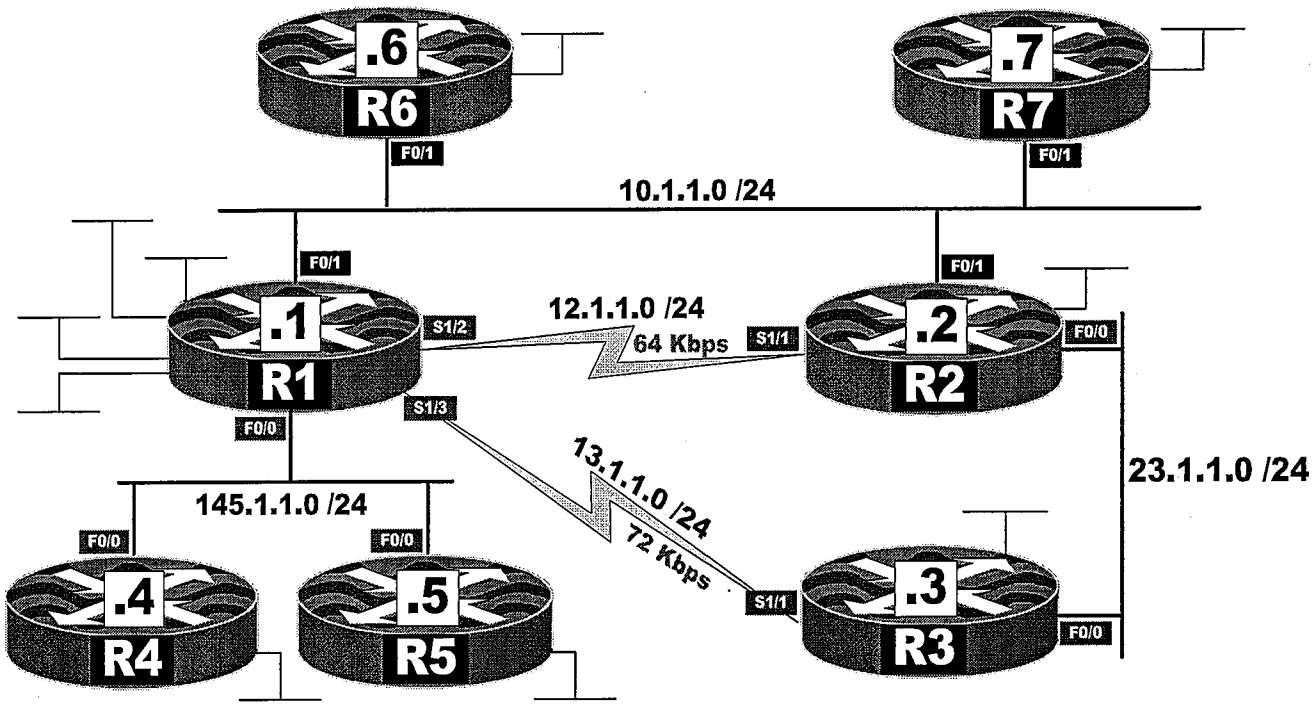
Lab-3:
Advanced EIGRP Stub Configuration

Lab-4:
EIGRP Filtering

Lab-5:
Configuring EIGRP Mock Lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 1 - EIGRP Named Configuration



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "EIGRP"→ "Lab-1".

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Configure EIGRP on R1, R2 and R3 based on the following policy:

Router	Interface	AS Number
R1	F0/1	200
	F0/0	100
	S1/2	100
	S1/3	100
	Looback 0 - 3	100
R2	F0/1	200
	F0/0	100
	S1/1	100
	Loopback0	100
R3	F0/0	100
	S1/1	100
	Loopbck0	100

R1 should be configured to use Unicast to establish a neighbor adjacency with R2

R1 should use Multicast to establish a neighbor adjacency with R3.

R1 – R3 should use EIGRP named configuration to accomplish this task.

Configuring EIGRP with a virtual instance name creates an EIGRP named configuration. If you need to configure another EIGRP instance in another ASN, you must use another EIGRP virtual instance with a different name and specify the ASN in the address-family command.

The "Auto-summary" is disabled by default, and if you enter the "No auto-summary" command it will not be displayed in the output of the "show run | s router eigrp" command, unless it is specifically enabled, in which case it will be displayed.

Using the EIGRP named configuration the "Network" statement is configured under the "Address-family" configuration mode and NOT directly under the router EIGRP configuration mode.

The named configuration supports both IPv4 and IPv6, and VRF instance/s using different address-family types in a single EIGRP instance.

The named configuration provides a single place for all EIGRP commands and it can be displayed using a single show command:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#Network 12.1.1.1 0.0.0.0
R1(config-router-af)#Network 13.1.1.1 0.0.0.0
R1(config-router-af)#Network 145.1.1.1 0.0.0.0
R1(config-router-af)#Network 1.1.0.1 0.0.0.0
R1(config-router-af)#Network 1.1.1.1 0.0.0.0
R1(config-router-af)#Network 1.1.2.1 0.0.0.0
R1(config-router-af)#Network 1.1.3.1 0.0.0.0
```

We can see that the Neighbor command is configured under the Address-family and NOT directly under the router eigrp configuration mode:

```
R1(config-router-af)#Neighbor 12.1.1.2 S1/2

R1(config)#Router eigrp tst200
R1(config-router)#Address-family ipv4 unicast autonomous-system 200
R1(config-router-af)#Network 10.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp tst
R2(config-router)#Address-family ipv4 unicast autonomous-system 100
R2(config-router-af)#Network 12.1.1.2 0.0.0.0
R2(config-router-af)#Network 23.1.1.2 0.0.0.0
R2(config-router-af)#Network 2.2.2.2 0.0.0.0
R2(config-router-af)#Neighbor 12.1.1.1 S1/1
```

You should see the following console message stating that a neighbor adjacency has been established between the local router and R1 (12.1.1.1):

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 12.1.1.1 (Serial1/1) is up:
new adjacency

R2(config)#Router eigrp tst200
R2(config-router)#Address-family ipv4 unicast autonomous-system 200
R2(config-router-af)#Network 10.1.1.2 0.0.0.0
```

You should see the following console message stating that the local router has established an adjacency with R1 (10.1.1.1):

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 200: Neighbor 10.1.1.1 (FastEthernet0/1) is

up: new adjacency
```

On R3:

```
R3(config)#Router eigrp tst
R3(config-router)#Address-family ipv4 unicast autonomous-system 100
R3(config-router-af)#Network 23.1.1.3 0.0.0.0
R3(config-router-af)#Network 13.1.1.3 0.0.0.0
R3(config-router-af)#Network 3.3.3.3 0.0.0.0
```

To verify the configuration:

On R3:

```
R3#Show ip route eigrp 100 | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 4 subnets
D        1.1.0.0 [90/36195328] via 13.1.1.1, 00:08:42, Serial1/1
D        1.1.1.0 [90/36195328] via 13.1.1.1, 00:08:42, Serial1/1
D        1.1.2.0 [90/36195328] via 13.1.1.1, 00:08:42, Serial1/1
D        1.1.3.0 [90/36195328] via 13.1.1.1, 00:08:42, Serial1/1
D     2.0.0.0/8 [90/156160] via 23.1.1.2, 00:08:42, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/40514560] via 23.1.1.2, 00:08:42, FastEthernet0/0
```

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     2.0.0.0/8 [90/36197888] via 13.1.1.3, 00:09:23, Serial1/3
D     3.0.0.0/8 [90/36195328] via 13.1.1.3, 00:09:18, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/36069888] via 13.1.1.3, 00:09:23, Serial1/3
```

We can see that EIGRP has the capability to establish Unicast or Multicast or in this case both simultaneously.

Task 2

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Configure EIGRP on R4 and R5 in AS 100.

On R4:

```
R4(config)#Router eigrp tst
R4(config-router)#Address-family ipv4 unicast autonomous-system 100
R4(config-router-af)#Network 145.1.1.4 0.0.0.0
R4(config-router-af)#Network 4.4.4.4 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 145.1.1.1 (FastEthernet0/0) is
up: new adjacency
```

On R5:

```
R5(config)#Router eigrp tst
R5(config-router)#Address-family ipv4 unicast autonomous-system 100
R5(config-router-af)#Network 5.5.5.5 0.0.0.0
R5(config-router-af)#Network 145.1.1.5 0.0.0.0
```

You should see the following console messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 145.1.1.4 (FastEthernet0/0) is
up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 145.1.1.1 (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

      1.0.0.0/24 is subnetted, 4 subnets
D        1.1.0.0 [90/156160] via 145.1.1.1, 00:01:04, FastEthernet0/0
D        1.1.1.0 [90/156160] via 145.1.1.1, 00:01:04, FastEthernet0/0
D        1.1.2.0 [90/156160] via 145.1.1.1, 00:01:04, FastEthernet0/0
D        1.1.3.0 [90/156160] via 145.1.1.1, 00:01:04, FastEthernet0/0
D     2.0.0.0/8 [90/36200448] via 145.1.1.1, 00:01:04, FastEthernet0/0
D     3.0.0.0/8 [90/36197888] via 145.1.1.1, 00:01:04, FastEthernet0/0
D     4.0.0.0/8 [90/156160] via 145.1.1.4, 00:01:04, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/40514560] via 145.1.1.1, 00:01:04, FastEthernet0/0

      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/36069888] via 145.1.1.1, 00:01:04, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/36072448] via 145.1.1.1, 00:01:04, FastEthernet0/0
```

Task 3

Configure R1, R4 and R5 to use unicast to establish their neighbor adjacency.

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#Neighbor 145.1.1.4 F0/0
```

NOTE: Once the above "Neighbor" command is configured the following console message states that both neighbors (R4 and R5) lost their adjacencies to the local router:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 145.1.1.5 (FastEthernet0/0) is
down: Static peer configured

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 145.1.1.4 (FastEthernet0/0) is
down: Static peer configured
```

But why didn't R5 re-establish its adjacency?
In task 1, we saw that EIGRP established a Unicast adjacency with R2 and a Multicast adjacency with R3, but in this case once the "neighbor" command was configured all adjacencies on that segment went down.
So what is different about R1, R4 and R5 versus R1, R2 and R3?
R1, R2, and R3 were configured in a point-to-point manner, whereas, R1, R4 and R5 are on an Ethernet segment. Therefore, on MAB (Multi-access Broadcast) segment EIGRP can NOT Unicast to one or some neighbors and Multicast to others, it's either one or the other.
Let's configure R4 and R5:

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#Neighbor 145.1.1.5 F0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R4:

```
R4(config)#Router eigrp tst
R4(config-router)#Address-family ipv4 unicast autonomous-system 100
R4(config-router-af)#Neighbor 145.1.1.1 F0/0
R4(config-router-af)#Neighbor 145.1.1.5 F0/0
```

On R5:

```
R5(config)#Router eigrp tst
R5(config-router)#Address-family ipv4 unicast autonomous-system 100
R5(config-router-af)#Neighbor 145.1.1.1 F0/0
R5(config-router-af)#Neighbor 145.1.1.4 F0/0
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 4 subnets
D        1.1.0.0 [90/156160] via 145.1.1.1, 00:07:46, FastEthernet0/0
D        1.1.1.0 [90/156160] via 145.1.1.1, 00:07:46, FastEthernet0/0
D        1.1.2.0 [90/156160] via 145.1.1.1, 00:07:46, FastEthernet0/0
D        1.1.3.0 [90/156160] via 145.1.1.1, 00:07:46, FastEthernet0/0
D     2.0.0.0/8 [90/36200448] via 145.1.1.1, 00:07:46, FastEthernet0/0
D     3.0.0.0/8 [90/36197888] via 145.1.1.1, 00:07:46, FastEthernet0/0
D     5.0.0.0/8 [90/156160] via 145.1.1.5, 00:00:17, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/40514560] via 145.1.1.1, 00:07:46, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/36069888] via 145.1.1.1, 00:07:46, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/36072448] via 145.1.1.1, 00:07:46, FastEthernet0/0
```

Task 4

Configure R6 in EIGRP AS 200, this router should run EIGRP AS 200 on its F0/1 and its Loopback0 interface, you should use an EIGRP named configuration to accomplish this task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6(config)#Router eigrp tst200
R6(config-router)#Address-family ipv4 unicast autonomous-system 200
R6(config-router-af)#Network 6.6.6.6 0.0.0.0
R6(config-router-af)#Network 10.1.1.6 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 200: Neighbor 10.1.1.2 (FastEthernet0/1) is
up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 200: Neighbor 10.1.1.1 (FastEthernet0/1) is
up: new adjacency
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp 200 | B Gate
Gateway of last resort is not set

D      6.0.0.0/8 [90/156160] via 10.1.1.6, 00:03:03, FastEthernet0/1
```

Task 5

Configure OSPF area 0 on R6's F0/1 and R7's F0/1 and Loopback0 interfaces. You should use a router-id of your choice.

On R6:

```
R6(config)#Router OSPF 1
R6(config-router)#Router-id 0.0.0.6
R6(config-router)#Network 10.1.1.6 0.0.0.0 area 0
```

On R7:

```
R7(config)#Router OSPF 1
R7(config-router)#Router-id 0.0.0.7
R7(config-router)#Network 10.1.1.7 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R7(config-router)#Network 7.7.7.7 0.0.0.0 area 0
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.6 on FastEthernet0/1 from LOADING to
FULL, Loading Done
```

To verify the configuration:

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 10.1.1.7, 00:00:49, FastEthernet0/1
```

Task 6

Configure R6 to redistribute OSPF into EIGRP such that R1 and R2 go directly to R7 to reach network 7.0.0.0/8 network.

In the address-family configuration mode you will see the "Topology base", the route redistribution is configured in this configuration mode. The following commands are configured in this mode:

```
R6(config)#Router eigrp tst
R6(config-router)#Address-family ipv4 unicast autonomous-system 200
R6(config-router-af)#Topology base
R6(config-router-af-topology)#?
Address Family Topology configuration commands:
  auto-summary          Enable automatic network number summarization
  default               Set a command to its defaults
  default-information   Control distribution of default information
  default-metric        Set metric of redistributed routes
  distance              Define an administrative distance
  distribute-list       Filter entries in eigrp updates
  eigrp                 EIGRP specific commands
  exit-af-topology      Exit from Address Family Topology configuration mode
  maximum-paths         Forward packets over multiple paths
  metric                Modify metrics and parameters for advertisement
  no                    Negate a command or set its defaults
  offset-list           Add or subtract offset from EIGRP metrics
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  redistribute            Redistribute IPv4 routes from another routing protocol
  snmp                    Modify snmp parameters

  summary-metric          Specify summary to apply metric/filtering
  timers                  Adjust topology specific timers
  traffic-share           How to compute traffic share over alternate paths
  variance                Control load balancing variance
```

On R6:

```
R6(config)#Router eigrp tst200
R6(config-router)#Address-family ipv4 unicast autonomous-system 200
R6(config-router-af)#Topology base
R6(config-router-af-topology)#Redistribute OSPF 1 metric 1 1 1 1 1
```

To verify the configuration:

On R1:

```
R1#Show ip route 7.7.7.7

Routing entry for 7.7.7.7/32
  Known via "eigrp 200", distance 170, metric 2560002816, type external
  Redistributing via eigrp 200
  Last update from 10.1.1.6 on FastEthernet0/1, 00:11:13 ago
  Routing Descriptor Blocks:
  * 10.1.1.6, from 10.1.1.6, 00:11:13 ago, via FastEthernet0/1
      Route metric is 2560002816, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

We can see that the next-hop is pointing to R6 and NOT R7. To resolve this task we can use the third party next-hop, in classic EIGRP configuration this feature was configured directly under the interface, in the EIGRP named configuration we have the "af-interface" mode directly under the "address-family" configuration mode. Let's see our options in the "af-interface" mode:

On R6:

```
R6(config)#Router eigrp tst200
R6(config-router)#Address-family ipv4 unicast autonomous-system 200
R6(config-router-af)#af-interface F0/1
R6(config-router-af-interface)#?
Address Family Interfaces configuration commands:
  authentication          authentication subcommands
  bandwidth-percent       Set percentage of bandwidth percentage limit
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
  bfd                    Enable Bidirectional Forwarding Detection
  dampening-change       Percent interface metric must change to cause update

  dampening-interval     Time in seconds to check interface metrics
  default                Set a command to its defaults
  exit-af-interface      Exit from Address Family Interface configuration mode
  hello-interval         Configures hello interval
  hold-time              Configures hold time
  next-hop-self          Configures EIGRP next-hop-self
  no                     Negate a command or set its defaults
  passive-interface      Suppress address updates on an interface
  shutdown               Disable Address-Family on interface
  split-horizon          Perform split horizon
  summary-address        Perform address summarization
```

We can see that most if not all the interface configuration commands for EIGRP are listed in this mode. Let's configure the third party next-hop:

```
R6(config-router-af-interface)#No next-hop-self
```

To verify the configuration:

On R1:

```
R1#Show ip route 7.7.7.7

Routing entry for 7.7.7.7/32
  Known via "eigrp 200", distance 170, metric 2560002816, type external
  Redistributing via eigrp 200
  Last update from 10.1.1.7 on FastEthernet0/1, 00:01:09 ago
  Routing Descriptor Blocks:
  * 10.1.1.7, from 10.1.1.6, 00:01:09 ago, via FastEthernet0/1
      Route metric is 2560002816, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

Task 7

Configure the hello interval of all routers in AS 200 to be twice as much as the default.

Let's see the default value for the "Hello interval" on R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip eigrp 200 interface detail


EIGRP-IPv4 VR(tst200) Address-Family Interfaces for AS(200)
                      Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/1              2       0/0        6       0/1           50           0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Un/reliable mcasts: 0/4  Un/reliable ucasts: 6/8
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 3  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is not set
```

Now that we know the default values, let's change the Hello interval to 10 seconds:

On R1, R2 and R6:

```
Rx(config)#Router eigrp tst200
Rx(config-router)#Address-family ipv4 unicast autonomous-system 200
Rx(config-router-af)#af-interface F0/1
Rx(config-router-af-interface)#hello-interval 10
```

To verify the configuration:

On R1:

```
R1#Show ip eigrp 200 interface detail

EIGRP-IPv4 VR(tst200) Address-Family Interfaces for AS(200)
                      Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/1              2       0/0        6       0/1           50           0
  Hello-interval is 10, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Un/reliable mcasts: 0/4  Un/reliable ucasts: 6/8
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 3  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is not set
```

NOTE: When the hello interval is changed the hold-time is NOT changed automatically.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 8

R4 should be configured such that in the worst case scenario uses 10 percent of the bandwidth for its EIGRP updates. This policy should apply to the existing and future interfaces.

By default, EIGRP will ONLY consume 50 percent of its interface's bandwidth, this can be changed using the "Bandwidth percent" command in interface configuration mode or in the "af-interface" configuration mode under the address-family. Let's configure this feature in the "af-interface" configuration mode:

On R4:

```
R4(config)#Router eigrp tst
R4(config-router)#Address-family ipv4 unicast autonomous-system 100
```

The following configuration applies the percentage to all (existing and future) interfaces on this router:

```
R4(config-router-af)#af-interface default
R4(config-router-af-interface)#bandwidth-percent 10
```

Task 9

Configure R1 to summarize its loopback interfaces and advertise a single summary.

NOTE: The "af-interface default" can NOT be used to summarize and the specific interface/s must be specified:

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#af-interface s1/2
R1(config-router-af-interface)#summary-address 1.1.0.0 255.255.252.0
R1(config-router-af)#af-interface s1/3
R1(config-router-af-interface)#summary-address 1.1.0.0 255.255.252.0
```

To verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

      1.0.0.0/22 is subnetted, 1 subnets
D        1.1.0.0 [90/36197888] via 23.1.1.3, 00:02:24, FastEthernet0/0
D     3.0.0.0/8 [90/156160] via 23.1.1.3, 02:19:11, FastEthernet0/0
D     4.0.0.0/8 [90/36200448] via 23.1.1.3, 01:47:13, FastEthernet0/0
D     5.0.0.0/8 [90/36200448] via 23.1.1.3, 01:46:18, FastEthernet0/0
D     6.0.0.0/8 [90/156160] via 10.1.1.6, 00:46:08, FastEthernet0/1
      7.0.0.0/32 is subnetted, 1 subnets
D EX     7.7.7.7 [170/2560002816] via 10.1.1.7, 00:46:08, FastEthernet0/1
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/36069888] via 23.1.1.3, 02:19:17, FastEthernet0/0
      145.1.0.0/24 is subnetted, 1 subnets
D        145.1.1.0 [90/36072448] via 23.1.1.3, 02:04:01, FastEthernet0/0
```

On R3:

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

      1.0.0.0/22 is subnetted, 1 subnets
D        1.1.0.0 [90/36195328] via 13.1.1.1, 00:03:23, Serial1/1
D     2.0.0.0/8 [90/156160] via 23.1.1.2, 02:20:14, FastEthernet0/0
D     4.0.0.0/8 [90/36197888] via 13.1.1.1, 01:48:13, Serial1/1
D     5.0.0.0/8 [90/36197888] via 13.1.1.1, 01:47:18, Serial1/1
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/40514560] via 23.1.1.2, 02:20:14, FastEthernet0/0
      145.1.0.0/24 is subnetted, 1 subnets
D        145.1.1.0 [90/36069888] via 13.1.1.1, 02:05:01, Serial1/1
```

We can see that R2 traverses through R3 to get to the loopback interfaces or the summary for the loopback interfaces, whereas, R3 goes directly to R1 to reach the summary route; this is because the bandwidth of R2 is set to 64Kbps whereas the bandwidth of R3 is set to 72Kbps. To test this:

On R2:

```
R2#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 23.1.1.3 0 msec 0 msec 0 msec
  2 13.1.1.1 16 msec *  12 msec
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

```
R3#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 13.1.1.1 12 msec *  12 msec
```

Task 10

Configure R1 to assign a fixed metric to its summary route such that R2 and R3 both use R1 as the next-hop to reach the more specific routes of the summary regardless of the bandwidth command on the serial links to R2 and R3.

Before we configure this task, let's see the cost of the summary route on R2 and R3:

On R2:

```
R2#Show ip route 1.1.0.0 | Inc metric

  Known via "eigrp 100", distance 90, metric 36197888, type internal
      Route metric is 36197888, traffic share count is 1
```

On R3:

```
R3#Show ip route 1.1.0.0 | Inc metric

  Known via "eigrp 100", distance 90, metric 36195328, type internal
      Route metric is 36195328, traffic share count is 1
```

The "Traceroutes" performed at the end of the previous task revealed that R2 goes through R3 to reach the more specific routes of the summary. Let's assign a FIXED cost to the summary so that both R2 and R3 use R1 as the next-hop to reach the components of the summary route:

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-router-af)#topology base
```

The following command assigns a fixed metric for the EIGRP summary route:

```
R1(config-router-af-topology)#summary-metric 1.1.0.0/22 1 1 1 1 1
```

To verify the configuration:

On R2:

```
R2#Show ip route 1.1.0.0 | Inc metric

  Known via "eigrp 100", distance 90, metric 2560512256, type internal
      Route metric is 2560512256, traffic share count is 1
```

On R3:

```
R3#Show ip route 1.1.0.0 | Inc metric

  Known via "eigrp 100", distance 90, metric 2560512256, type internal
      Route metric is 2560512256, traffic share count is 1
```

To test the configuration:

On R2:

```
R2#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.1 12 msec *  12 msec
```

On R3:

```
R3#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 13.1.1.1 16 msec *  12 msec
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 11

Configure R1 to limit the number of received prefixes from R5 to 10, R1 should be configured to receive a warning message once 50 percent of this threshold is reached, and a warning message for every additional route that exceeds the threshold.

To test and configure this feature, let's configure 10 additional loopback interfaces on R5 and then advertise one at a time and verify the result:

On R5:

```
R5(config-if)#Int Lo1
R5(config-if)#ip addr 51.1.1.5 255.0.0.0

R5(config-if)#Int Lo2
R5(config-if)#ip addr 52.1.1.5 255.0.0.0

R5(config-if)#Int Lo3
R5(config-if)#ip addr 53.1.1.5 255.0.0.0

R5(config-if)#Int Lo4
R5(config-if)#ip addr 54.1.1.5 255.0.0.0

R5(config-if)#Int Lo5
R5(config-if)#ip addr 55.1.1.5 255.0.0.0

R5(config-if)#Int Lo6
R5(config-if)#ip addr 56.1.1.5 255.0.0.0

R5(config-if)#Int Lo7
R5(config-if)#ip addr 57.1.1.5 255.0.0.0

R5(config-if)#Int Lo8
R5(config-if)#ip addr 58.1.1.5 255.0.0.0

R5(config-if)#Int Lo9
R5(config-if)#ip addr 59.1.1.5 255.0.0.0

R5(config-if)#Int Lo10
R5(config-if)#ip addr 60.1.1.5 255.0.0.0
```

Let's configure R1 for this feature:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#Neighbor 145.1.1.5 maximum-prefix 10 50 warning-only
```

To test and verify the configuration:

On R5:

```
R5(config)#Router eigrp tst
R5(config-router)#Address-family ipv4 unicast autonomous-system 100
R5(config-router-af)#Network 51.1.1.5 0.0.0.0
R5(config-router-af)#Network 52.1.1.5 0.0.0.0
R5(config-router-af)#Network 53.1.1.5 0.0.0.0
```

So far R1 has not received any warning messages because it is only receiving four networks from R5, let's advertise one more and verify the result:

```
R5(config-router-af)#Network 54.1.1.5 0.0.0.0
```

Let's check the console messages on R1:

On R1:

Sure enough the first message was received:

```
%DUAL-4-PFXLIMITTHR: EIGRP-IPv4 100: Neighbor threshold prefix level(5)
reached.
```

Let's advertise another network on R5 and verify the result:

On R5:

```
R5(config-router-af)#Network 55.1.1.5 0.0.0.0
```

No console messages on R1 yet, let's advertise another network on R5:

```
R5(config-router-af)#Network 56.1.1.5 0.0.0.0
```

No console messages on R1 yet, let's advertise another three networks on R5 to reach the threshold of 10 networks:

```
R5(config-router-af)#Network 57.1.1.5 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5(config-router-af)#Network 58.1.1.5 0.0.0.0
R5(config-router-af)#Network 59.1.1.5 0.0.0.0
```

On R1:

```
%DUAL-3-PFXLIMIT: EIGRP-IPv4 100: Neighbor prefix limit reached(10).
```

Let's advertise another network on R5 and verify the result:

On R5:

```
R5(config-router-af)#Network 60.1.1.5 0.0.0.0
```

On R1:

We can see that for every additional network above the set threshold a console message is generated on R1's console:

```
%DUAL-3-PFXLIMIT: EIGRP-IPv4 100: Neighbor prefix limit reached(10).
```

Task 12

Configure R1 to limit the number of prefixes received from R4 to five, R1 should be configured to tear down the adjacency if R4 exceeds the specified threshold.

Let's configure five more loopback interfaces on R4 so we can test this feature:

On R4:

```
R4(config)#Int Lo1
R4(config-if)#IP addr 41.1.1.4 255.0.0.0

R4(config)#Int Lo2
R4(config-if)#IP addr 42.1.1.4 255.0.0.0

R4(config)#Int Lo3
R4(config-if)#IP addr 43.1.1.4 255.0.0.0

R4(config)#Int Lo4
R4(config-if)#IP addr 44.1.1.4 255.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#Int Lo5
R4(config-if)#IP addr 45.1.1.4 255.0.0.0
```

Let's configure R1 for this feature:

On R1:

```
R1(config)#Router eigrp tst
R1(config-router)#Address-family ipv4 unicast autonomous-system 100
R1(config-router-af)#Neighbor 145.1.1.4 maximum-prefix 5
```

Let's go back to R1 and advertise an extra Network and verify the results:

On R4:

```
R4(config)#Router eigrp tst
R4(config-router)#Address-family ipv4 unicast autonomous-system 100
R4(config-router-af)#Network 41.1.1.4 0.0.0.0
```

So far R1 has not received any console messages because it is only receiving two networks from R4, let's advertise three more and verify the result:

```
R4(config-router-af)#Network 42.1.1.4 0.0.0.0
R4(config-router-af)#Network 43.1.1.4 0.0.0.0
```

Sure enough we can see the following console message on R1, this is because 80 percent of the configured threshold is reached:

```
%DUAL-4-PFXLIMITTHR: EIGRP-IPv4 100: Neighbor threshold prefix level(3)
reached.
```

Let's advertise one more and verify the console messages on R1:

```
R4(config-router-af)#Network 44.1.1.4 0.0.0.0
```

We should see the following two console messages on R1 stating that the maximum threshold has been reached:

On R1:

```
%DUAL-3-PFXLIMIT: EIGRP-IPv4 100: Neighbor prefix limit reached(5).

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 145.1.1.4 (FastEthernet0/0)
is down: prefix-limit exceeded
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

We can see that the neighbor adjacency is torn down, let's verify the neighbor table of R4:

```
R4#Show ip eigrp neighbor

EIGRP-IPv4 VR(tst) Address-Family Neighbors for AS(100)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   145.1.1.5               Fa0/0             12 00:14:24    1   200  0  23

R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      5.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      51.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      52.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      53.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      54.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      55.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      56.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      57.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      58.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      59.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
D      60.0.0.0/8 [90/156160] via 145.1.1.5, 00:13:57, FastEthernet0/0
```

We can see that R4 only exchanges routes with R5 and NOT R1.

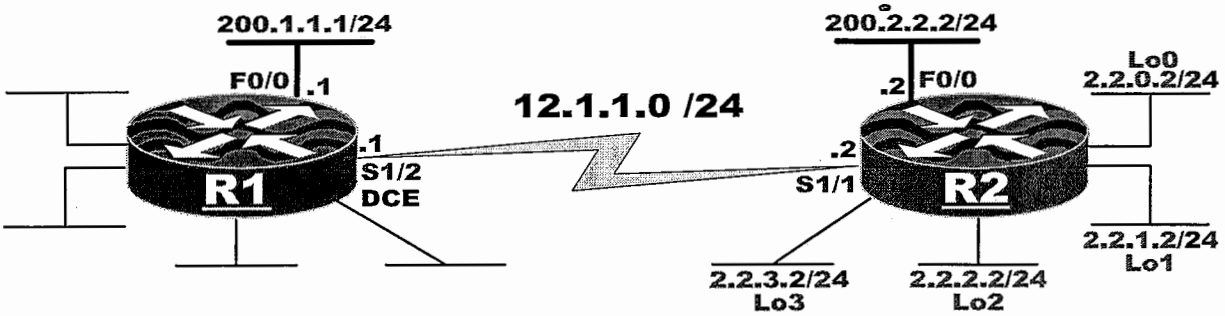
Task 13

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 2 - EIGRP Stub



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "EIGRP"→ "Lab-2".

Task 1

Configure EIGRP AS 100 on the serial and all loopback interfaces of these two routers. R1 should configure EIGRP using the classic method and R2 should use the new EIGRP named configuration to accomplish this task. **DO NOT run EIGRP on the F0/0 interfaces of these two routers.**

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 1.1.0.1 0.0.0.0
R1(config-router)#Network 1.1.1.1 0.0.0.0
R1(config-router)#Network 1.1.2.1 0.0.0.0
R1(config-router)#Network 1.1.3.1 0.0.0.0
R1(config-router)#Network 12.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp tst
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config-router)#Address-family ipv4 unicast autonomous-system 100
R2(config-router-af)#Network 2.2.0.2 0.0.0.0
R2(config-router-af)#Network 2.2.1.2 0.0.0.0
R2(config-router-af)#Network 2.2.2.2 0.0.0.0
R2(config-router-af)#Network 2.2.3.2 0.0.0.0
R2(config-router-af)#Network 12.1.1.2 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 12.1.1.1 (Serial1/1) is up:
new adjacency
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 4 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:03:18, Serial1/2
D        2.2.1.0 [90/20640000] via 12.1.1.2, 00:03:18, Serial1/2
D        2.2.2.0 [90/20640000] via 12.1.1.2, 00:03:18, Serial1/2
D        2.2.3.0 [90/20640000] via 12.1.1.2, 00:03:18, Serial1/2
```

On R2:

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

      1.0.0.0/24 is subnetted, 4 subnets
D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:03:42, Serial1/1
D        1.1.1.0 [90/20640000] via 12.1.1.1, 00:03:42, Serial1/1
D        1.1.2.0 [90/20640000] via 12.1.1.1, 00:03:42, Serial1/1
D        1.1.3.0 [90/20640000] via 12.1.1.1, 00:03:42, Serial1/1
```

Task 2

Configure R1 and R2 to summarize their Loopback interfaces in EIGRP.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#Int S1/2
R1(config-if)#IP summary-address eigrp 100 1.1.0.0 255.255.252.0
```

On R2:

```
R2(config)#Router eigrp tst
R2(config-router)#Address-family ipv4 unicast autonomous-system 100
R2(config-router-af)#af-interface S1/1
R2(config-router-af-interface)#summary-address 2.2.0.0 255.255.252.0
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 00:01:33, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:00:34, Serial1/2
```

On R2:

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

      1.0.0.0/22 is subnetted, 1 subnets
D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:02:21, Serial1/1
      2.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        2.2.0.0/22 is a summary, 00:01:22, Null0
```

Task 3

Configure the following static routes on R1 and R2 and redistribute them into EIGRP.

Router	Static route
R1	Create a static route to network 11.0.0.0 /8 via F0/0
R2	**Create a static route to network 22.0.0.0 /8 via F0/0**

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1(config)#IP route 11.0.0.0 255.0.0.0 F0/0

R1(config)#Router eigrp 100
R1(config-router)#Redistribute static
```

On R2:

```
R2(config)#IP route 22.0.0.0 255.0.0.0 F0/0

R2(config)#Router eigrp tst
R2(config-router)#Address-family ipv4 unicast auto 100
R2(config-router-af)#Topology Base
R2(config-router-af-topology)#Redistribute static
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 00:16:40, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:15:41, Serial1/2
D EX  22.0.0.0/8 [170/20514560] via 12.1.1.2, 00:01:37, Serial1/2
```

On R2:

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

      1.0.0.0/22 is subnetted, 1 subnets
D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:17:13, Serial1/1
      2.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        2.2.0.0/22 is a summary, 00:16:14, Null0
D EX  11.0.0.0/8 [170/20514560] via 12.1.1.1, 00:05:18, Serial1/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 4

Advertise the FastEthernet interfaces of these two routers in RIPv2 and disable auto-summarization. You should redistribute RIP into EIGRP, use any metric for redistributed routes.

On R1:

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 200.1.1.0

R1(config)#Router eigrp 100
R1(config-router)#Redistribute rip metric 1 1 1 1 1
```

On R2:

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 200.2.2.0

R2(config)#Router eigrp tst
R2(config-router)#Address-family ipv4 unicast auto 100
R2(config-router-af)#Topology Base
R2(config-router-af-topology)#Redistribute rip metric 1 1 1 1 1
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 00:31:50, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:30:51, Serial1/2
D EX  22.0.0.0/8 [170/20514560] via 12.1.1.2, 00:16:47, Serial1/2
D EX  200.2.2.0/24 [170/2560512256] via 12.1.1.2, 00:00:47, Serial1/2
```

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

      1.0.0.0/22 is subnetted, 1 subnets
D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:32:20, Serial1/1
      2.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        2.2.0.0/22 is a summary, 00:31:21, Null0
D EX  11.0.0.0/8 [170/20514560] via 12.1.1.1, 00:20:25, Serial1/1
D EX  200.1.1.0/24 [170/2560512256] via 12.1.1.1, 00:02:03, Serial1/1
```

Task 5

Configure Eigrp Stub routing on R1 using the "**Eigrp stub connected**"; test this option and verify the routes in the routing tables of both routers.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Eigrp stub connected
```

You should see the following console message stating that the neighbor adjacency was reset:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 12.1.1.2 (Serial1/2) is down:
peer info changed

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 12.1.1.2 (Serial1/2) is up:
new adjacency
```

The "EIGRP Stub" routing can also be configured in EIGRP named configuration in the "Address-family" configuration mode, but in this case it is ONLY configured on R1. The following shows how EIGRP Stub routing can be configured in EIGRP named configuration:

```
R2(config)#Router eigrp tst
R2(config-router)#Address-family ipv4 unicast auto 100
R2(config-router-af)#Eigrp stub ?
  connected      Do advertise connected routes
  leak-map       Allow dynamic prefixes based on the leak-map
  receive-only   Set receive only neighbor
  redistributed  Do advertise redistributed routes
  static         Do advertise static routes
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  summary          Do advertise summary routes
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 12.1.1.1

D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:08:13, Serial1/1
D        1.1.1.0 [90/20640000] via 12.1.1.1, 00:08:13, Serial1/1
D        1.1.2.0 [90/20640000] via 12.1.1.1, 00:08:13, Serial1/1
D        1.1.3.0 [90/20640000] via 12.1.1.1, 00:08:13, Serial1/1
```

NOTE: Only the directly connected networks that R1 used a network command to advertise were advertised to R2.

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 01:09:13, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:09:11, Serial1/2
D EX  22.0.0.0/8 [170/20514560] via 12.1.1.2, 00:09:11, Serial1/2
D EX  200.2.2.0/24 [170/2560512256] via 12.1.1.2, 00:09:11, Serial1/2
```

NOTE: The routing table of R1 was NOT affected at all.

Task 6

Remove the "Eigrp stub connected" option configured in the previous task and re-configure Eigrp Stub routing on R1 using the "**Eigrp stub summary**"; test this option and verify the routes in the routing tables of both routers.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#No eigrp stub connected
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-router)#Eigrp stub summary
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 12.1.1.1

D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:00:30, Serial1/1
```

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 01:12:59, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:00:57, Serial1/2
D EX  22.0.0.0/8 [170/20514560] via 12.1.1.2, 00:00:57, Serial1/2
D EX  200.2.2.0/24 [170/2560512256] via 12.1.1.2, 00:00:57, Serial1/2
```

By looking at the routing table of these two routers, we can clearly see that ONLY the summarized route/s were advertised to R2, but once again the routing table of R1 was NOT affected at all.

Task 7

Remove the "Eigrp stub summary" option configured in the previous task and re-configure Eigrp Stub routing on R1 using the "**Eigrp stub Static**"; test this option and verify the routes in the routing tables of both routers.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#No eigrp stub summary
R1(config-router)#Eigrp stub Static
```

To verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Show ip route eigrp | Inc 12.1.1.1

D EX  11.0.0.0/8 [170/20514560] via 12.1.1.1, 00:00:57, Serial1/1
```

On R1:

```
R1#Sh ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 01:16:33, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:01:41, Serial1/2
D EX  22.0.0.0/8 [170/20514560] via 12.1.1.2, 00:01:41, Serial1/2
D EX  200.2.2.0/24 [170/2560512256] via 12.1.1.2, 00:01:41, Serial1/2
```

The output of the above show command reveals that R1's routing table was NOT affected at all but R2 received the static routes that R1 redistributed into its routing table. NOTE: The routing protocol/s that are redistributed into EIGRP are NOT advertised and ONLY the static routes are advertised.

Task 8

Remove the "Eigrp stub Static" option configured in the previous task and re-configure Eigrp Stub routing on R1 using the "**Eigrp stub redistributed**"; test this option and verify the routes in the routing tables of both routers.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#No eigrp stub static
R1(config-router)#Eigrp stub redistributed
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 12.1.1.1

D EX  11.0.0.0/8 [170/20514560] via 12.1.1.1, 00:00:26, Serial1/1
D EX  200.1.1.0/24 [170/2560512256] via 12.1.1.1, 00:00:26, Serial1/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

NOTE: All the redistributed routes are advertised to R2, this means that the static route/s and the routing protocol/s that were redistributed into Eigrp were advertised with this option.

Task 9

Remove the "Eigrp stub redistributed" option configured in the previous task and re-configure Eigrp Stub routing on R1 using the "**Eigrp stub receive-only**"; test this option and verify the routes in the routing tables of both routers.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#No eigrp stub redistributed
R1(config-router)#Eigrp stub receive-only
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 12.1.1.1
R2#
```

We can clearly see that nothing was advertised to R2, with this option configured on R1, R1 receives all the routes from R2, but does NOT advertise any routes to R2.

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        1.1.0.0/22 is a summary, 01:31:37, Null0
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.0.0 [90/20640000] via 12.1.1.2, 00:01:27, Serial1/2
D EX  22.0.0.0/8 [170/20514560] via 12.1.1.2, 00:01:27, Serial1/2
D EX  200.2.2.0/24 [170/2560512256] via 12.1.1.2, 00:01:27, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 10

Remove the "Eigrp stub receive-only" option configured in the previous task and re-configure Eigrp Stub routing on R1 using the "**Eigrp stub**"; test this option and verify the routes in the routing tables of both routers.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#No eigrp stub receive-only
R1(config-router)#Eigrp stub
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 12.1.1.1

D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:01:54, Serial1/1
```

NOTE: All the directly connected and the summary route/s are advertised to R2, the reason we only see a single summary route is because when summarization is performed the specific routes that are included in the summary are suppressed, to prove this fact let's configure Loopback100 interface on R1 and assign an IP address of 100.1.1.1/24 and advertise this route in Eigrp AS 100. If this is performed successfully, R2 should see the summary plus the 100.1.1.0/24 routes:

On R1:

```
R1(config)#Int Lo100
R1(config-if)#Ip addr 100.1.1.1 255.255.255.0

R1(config)#Router eigrp 100
R1(config-router)#Netw 100.1.1.1 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 12.1.1.1

D        1.1.0.0 [90/20640000] via 12.1.1.1, 00:04:45, Serial1/1
D        100.1.1.0 [90/20640000] via 12.1.1.1, 00:00:07, Serial1/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

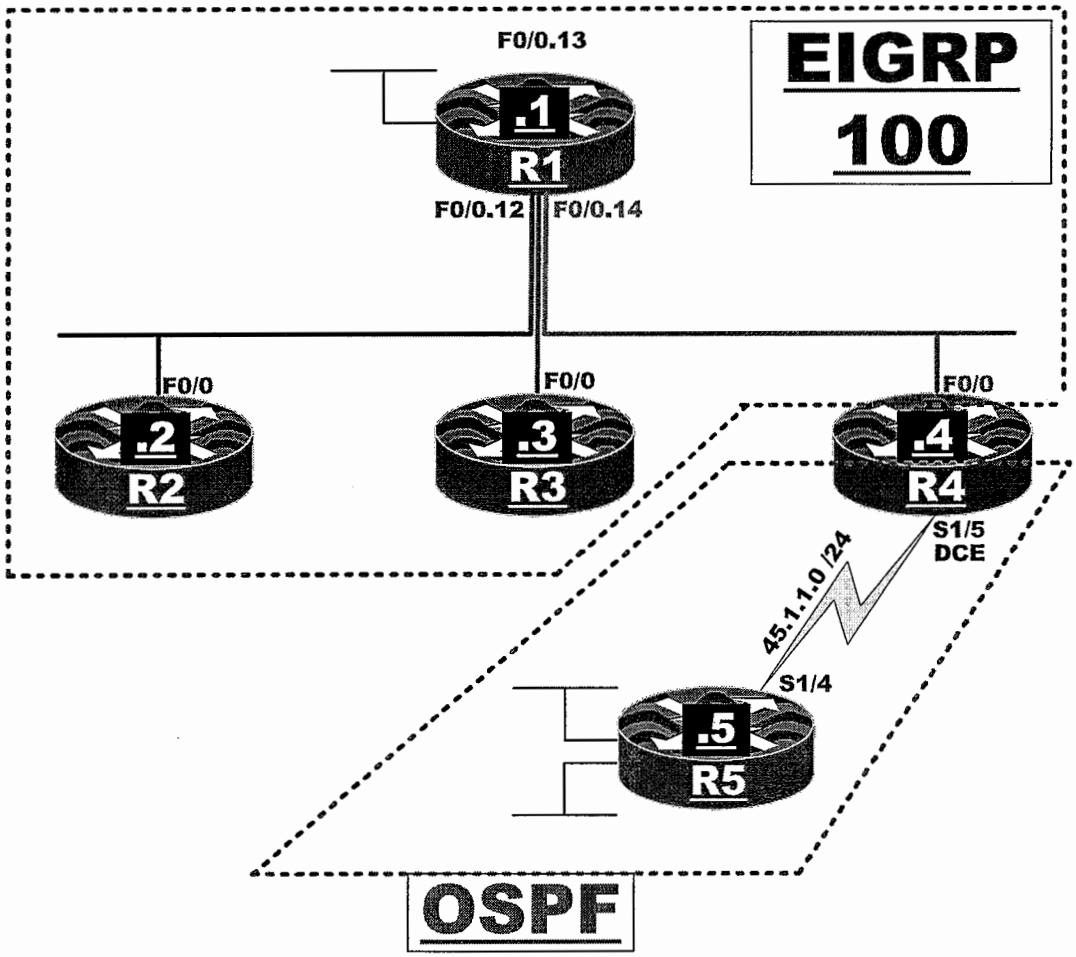
Task 11

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 3 – Advanced EIGRP Stub Configuration



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "EIGRP"→ "Lab-3".

Task 1

Configure OSPF area 0 on the following routers/interfaces; ensure that the loopback interfaces are advertised with their correct mask:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Router	Interface
R4	S1/5
R5	S1/4 Loopback 2 Loopback 3

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#Network 45.1.1.4 0.0.0.0 area 0
```

On R5:

```
R5(config)#int lo2
R5(config-if)#ip ospf netw point-to-point

R5(config-if)#int lo3
R5(config-if)#ip ospf netw point-to-point

R5(config)#Router ospf 1
R5(config-router)#router-id 0.0.0.5
R5(config-router)#Network 45.1.1.5 0.0.0.0 area 0
R5(config-router)#Network 2.2.2.2 0.0.0.0 area 0
R5(config-router)#Network 3.3.3.3 0.0.0.0 area 0
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | B Gate
Gateway of last resort is not set

O     2.0.0.0/8 [110/782] via 45.1.1.5, 00:00:13, Serial1/5
O     3.0.0.0/8 [110/782] via 45.1.1.5, 00:00:03, Serial1/5
```

Task 2

Configure EIGRP 100 on the following routers/interfaces.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Router	Interface
R1	F0/0.12 F0/0.13 F0/0.14 Loopback 0
R2	F0/0
R3	F0/0
R4	F0/0

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 12.1.1.1 0.0.0.0
R1(config-router)#Network 13.1.1.1 0.0.0.0
R1(config-router)#Network 14.1.1.1 0.0.0.0
R1(config-router)#Network 1.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp 100
R2(config-router)#Network 12.1.1.2 0.0.0.0
```

On R3:

```
R3(config)#Router eigrp 100
R3(config-router)#Network 13.1.1.3 0.0.0.0
```

On R4:

```
R4(config)#Router eigrp 100
R4(config-router)#Network 14.1.1.4 0.0.0.0
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     1.0.0.0/8 [90/156160] via 14.1.1.1, 00:00:35, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 14.1.1.1, 00:00:35, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 14.1.1.1, 00:00:35, FastEthernet0/0
```

On R3:

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

D     1.0.0.0/8 [90/156160] via 13.1.1.1, 00:01:27, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 13.1.1.1, 00:01:27, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 13.1.1.1, 00:01:27, FastEthernet0/0
```

On R2:

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

D     1.0.0.0/8 [90/156160] via 12.1.1.1, 00:02:18, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 12.1.1.1, 00:02:18, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 12.1.1.1, 00:02:18, FastEthernet0/0
```

Task 3

Configure mutual redistribution between OSPF and EIGRP on R4; use a metric of your choice.

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#Redistribute eigrp 100 subnets

R4(config)#Router eigrp 100
R4(config-router)#Redistribute ospf 1 metric 1 1 1 1 1
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

O E2  1.0.0.0/8 [110/20] via 45.1.1.4, 00:00:51, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O E2     12.1.1.0 [110/20] via 45.1.1.4, 00:00:51, Serial1/4
      13.0.0.0/24 is subnetted, 1 subnets
O E2     13.1.1.0 [110/20] via 45.1.1.4, 00:00:51, Serial1/4
      14.0.0.0/24 is subnetted, 1 subnets
O E2     14.1.1.0 [110/20] via 45.1.1.4, 00:00:51, Serial1/4
```

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

D EX  2.0.0.0/8 [170/2560002816] via 14.1.1.4, 00:01:12, FastEthernet0/0.14
D EX  3.0.0.0/8 [170/2560002816] via 14.1.1.4, 00:01:12, FastEthernet0/0.14
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/2560002816] via 14.1.1.4, 00:01:12, FastEthernet0/0.14
```

On R2:

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

D     1.0.0.0/8 [90/156160] via 12.1.1.1, 00:05:18, FastEthernet0/0
D EX  2.0.0.0/8 [170/2560005376] via 12.1.1.1, 00:01:57, FastEthernet0/0
D EX  3.0.0.0/8 [170/2560005376] via 12.1.1.1, 00:01:57, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 12.1.1.1, 00:05:18, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 12.1.1.1, 00:05:18, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/2560005376] via 12.1.1.1, 00:01:57, FastEthernet0/0
```

On R3:

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

D     1.0.0.0/8 [90/156160] via 13.1.1.1, 00:05:28, FastEthernet0/0
D EX  2.0.0.0/8 [170/2560005376] via 13.1.1.1, 00:02:36, FastEthernet0/0
D EX  3.0.0.0/8 [170/2560005376] via 13.1.1.1, 00:02:36, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 13.1.1.1, 00:05:28, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 13.1.1.1, 00:05:28, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/2560005376] via 13.1.1.1, 00:02:36, FastEthernet0/0
```

Note: R2 and R3 have networks 2.2.2.0/24 and 3.3.3.0/24.

Task 4

Configure "EIGRP stub" on R1 such that it **only** advertises its directly connected interfaces that are advertised with a "network" command to its EIGRP neighbors.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Eigrp stub connected
```

To verify the configuration:

On R2:

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

D     1.0.0.0/8 [90/156160] via 12.1.1.1, 00:00:12, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 12.1.1.1, 00:00:12, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 12.1.1.1, 00:00:12, FastEthernet0/0
```

On R3:

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

D     1.0.0.0/8 [90/156160] via 13.1.1.1, 00:00:39, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 13.1.1.1, 00:00:39, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D        14.1.1.0 [90/30720] via 13.1.1.1, 00:00:39, FastEthernet0/0
```

Note: R2 and R3 don't have the redistributed routes; this is because of R1's "stub connected" configuration.

Task 5

Configure R1 such that routers R2 and R3 have networks 2.0.0.0 /8 and 3.0.0.0 /8 in their routing table. **Do not** remove the "eigrp stub connected" configuration from R1 or summarize to accomplish this task.

To accomplish this task a "leak-map" is referenced in the "eigrp stub connected" command. The leak-map references a route-map called "TST", the route-map references an access-list. Any IP address/es that are permitted in the access-list is leaked along the connected networks.

On R1:

```
R1(config)#Access-list 1 permit 2.0.0.0 0.0.0.255
R1(config)#Access-list 1 permit 3.0.0.0 0.0.0.255

R1(config)#Route-map TST permit 10
R1(config-route-map)#Match ip addr 1

R1(config)#Router eigrp 100
R1(config-router)#Eigrp stub connected leak-map TST
```

To verify the configuration:

On R2:

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

D     1.0.0.0/8 [90/156160] via 12.1.1.1, 00:00:14, FastEthernet0/0
D EX  2.0.0.0/8 [170/2560005376] via 12.1.1.1, 00:00:14, FastEthernet0/0
D EX  3.0.0.0/8 [170/2560005376] via 12.1.1.1, 00:00:14, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 12.1.1.1, 00:00:14, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 12.1.1.1, 00:00:14, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

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

D     1.0.0.0/8 [90/156160] via 13.1.1.1, 00:01:19, FastEthernet0/0
D EX  2.0.0.0/8 [170/2560005376] via 13.1.1.1, 00:01:18, FastEthernet0/0
D EX  3.0.0.0/8 [170/2560005376] via 13.1.1.1, 00:01:18, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 13.1.1.1, 00:01:19, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 13.1.1.1, 00:01:19, FastEthernet0/0
```

Note: Both R2 and R3 have both networks (2.0.0.0 /8 and 3.0.0.0 /8) in their routing table.

Task 6

Re-configure R1 such that R2 gets network 2.0.0.0 /8 and R3 gets network 3.0.0.0 /8. **Do not** remove the "eigrp stub connected" configuration from R1 to accomplish this task.

The access-list and the route-map TST should be removed before proceeding further:

On R1:

```
R1(config)#No access-list 1
R1(config)#No route-map TST
```

We should identify the two networks using two different access-lists, in this case access-list 2 identifies network 2.0.0.0 /8 and access-list 3 identifies network 3.0.0.0 /8:

```
R1(config)#access-list 2 permit 2.0.0.0 0.255.255.255
R1(config)#access-list 3 permit 3.0.0.0 0.255.255.255
```

The next step is to configure a new route-map as follows:

```
R1(config)#Route-map TST permit 10
R1(config-route-map)#Match ip addr 2
R1(config-route-map)#Match inter F0/0.12
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-route-map)#Route-map TST permit 20
R1(config-route-map)#Match ip addr 3
R1(config-route-map)#Match inter F0/0.13
```

To verify the configuration:

On R2:

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

D     1.0.0.0/8 [90/156160] via 12.1.1.1, 00:06:16, FastEthernet0/0
D EX  2.0.0.0/8 [170/2560005376] via 12.1.1.1, 00:01:23, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/30720] via 12.1.1.1, 00:06:16, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 12.1.1.1, 00:06:16, FastEthernet0/0
```

On R3:

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

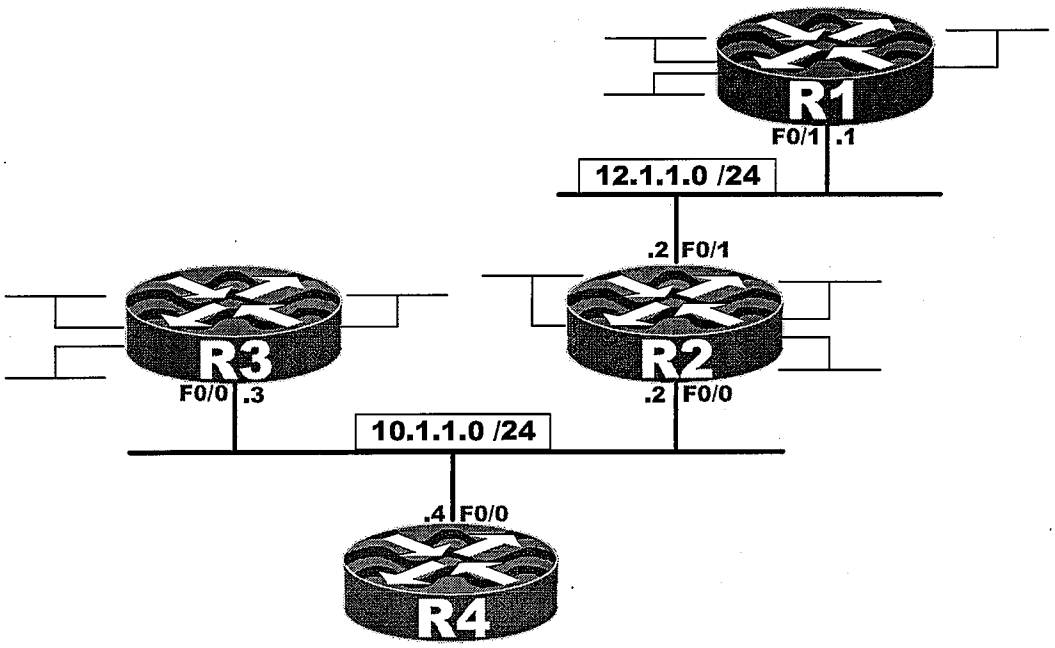
D     1.0.0.0/8 [90/156160] via 13.1.1.1, 00:06:38, FastEthernet0/0
D EX  3.0.0.0/8 [170/2560005376] via 13.1.1.1, 00:01:00, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 13.1.1.1, 00:06:38, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/30720] via 13.1.1.1, 00:06:38, FastEthernet0/0
```

Task 7

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

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 4 – EIGRP Filtering



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "EIGRP"→ "Lab-4".

Task 1

Configure EIGRP 100 on all routers and advertise their directly connected links.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 1.1.1.1 0.0.0.0
R1(config-router)#Network 11.1.1.1 0.0.0.0
R1(config-router)#Network 12.1.1.1 0.0.0.0
R1(config-router)#Network 111.1.1.1 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#Router eigrp 100
R2(config-router)#Network 12.1.1.2 0.0.0.0
R2(config-router)#Network 10.1.1.2 0.0.0.0
R2(config-router)#Network 200.1.1.1 0.0.0.0
R2(config-router)#Network 200.2.2.2 0.0.0.0
R2(config-router)#Network 2.2.2.2 0.0.0.0
```

On R3:

```
R3(config)#Router eigrp 100
R3(config-router)#Network 10.1.1.3 0.0.0.0
R3(config-router)#Network 200.1.1.1 0.0.0.0
R3(config-router)#Network 200.2.2.2 0.0.0.0
R3(config-router)#Network 3.3.3.3 0.0.0.0
```

On R4:

```
R4(config)#Router eigrp 100
R4(config-router)#Network 10.1.234.4 0.0.0.0
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      1.0.0.0/8 [90/158720] via 10.1.1.2, 00:00:33, FastEthernet0/0
D      2.0.0.0/8 [90/156160] via 10.1.1.2, 00:00:33, FastEthernet0/0
D      3.0.0.0/8 [90/156160] via 10.1.1.3, 00:00:33, FastEthernet0/0
D      11.0.0.0/8 [90/158720] via 10.1.1.2, 00:00:33, FastEthernet0/0
       12.0.0.0/24 is subnetted, 1 subnets
D         12.1.1.0 [90/30720] via 10.1.1.2, 00:00:33, FastEthernet0/0
D      111.0.0.0/8 [90/158720] via 10.1.1.2, 00:00:33, FastEthernet0/0
D      200.1.1.0/24 [90/156160] via 10.1.1.3, 00:00:33, FastEthernet0/0
                    [90/156160] via 10.1.1.2, 00:00:33, FastEthernet0/0
D      200.2.2.0/24 [90/156160] via 10.1.1.3, 00:00:33, FastEthernet0/0
                    [90/156160] via 10.1.1.2, 00:00:33, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 2

Configure R4 such that it filters existing (1.0.0.0/8, 11.0.0.0/8 and 111.0.0.0/8) and future network behind R1. **Do not** use "distribute-list", Access-list, Prefix-list, or "route-map" to accomplish this task

By default, EIGRP will discard routes that have a hop count of 101 or more, this behavior can be utilized to accomplish this task, as follows:

On R4:

```
R4#Sh ip eigrp topology 11.0.0.0 255.0.0.0 | inc Hop

        Hop count is 2

R4#Sh ip eigrp topology 1.0.0.0 255.0.0.0 | inc Hop

        Hop count is 2

R4#Sh ip eigrp topology 111.0.0.0 255.0.0.0 | inc Hop

        Hop count is 2
```

Note: The routes behind R1 have a hop count of 2, whereas, the other routes advertised in this topology have a hop count of 1:

```
R4#Sh ip eigrp topology 2.0.0.0 255.0.0.0 | inc Hop

   Hop count is 1

R4#Sh ip eigrp topology 3.0.0.0 255.0.0.0 | inc Hop

   Hop count is 1
```

Therefore, we should reject routes that have a hop count greater than 1.

On R4:

```
R4(config)#router eigrp 100
R4(config-router)#metric maximum-hops 1
```

Note: You should see the following console messages, this is because the policy for EIGRP is changed from 100 (default hop count) to only 1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.2 (FastEthernet0/0) is
down: Max hopcount changed
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.3 (FastEthernet0/0) is
down: Max hopcount changed
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.2 (FastEthernet0/0) is
up: new adjacency
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.3 (FastEthernet0/0) is
up: new adjacency
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      2.0.0.0/8 [90/156160] via 10.1.1.2, 00:03:20, FastEthernet0/0
D      3.0.0.0/8 [90/156160] via 10.1.1.3, 00:03:20, FastEthernet0/0
       12.0.0.0/24 is subnetted, 1 subnets
D         12.1.1.0 [90/30720] via 10.1.1.2, 00:03:20, FastEthernet0/0
D      200.1.1.0/24 [90/156160] via 10.1.1.3, 00:03:20, FastEthernet0/0
                    [90/156160] via 10.1.1.2, 00:03:20, FastEthernet0/0
D      200.2.2.0/24 [90/156160] via 10.1.1.3, 00:03:20, FastEthernet0/0
                    [90/156160] via 10.1.1.2, 00:03:20, FastEthernet0/0
```

Task 3

Configure R4 such that it **only** takes R2 to reach network 200.1.1.0 /24. R4 should **never** go through R3 to reach network 200.1.1.0/24. You should use an access-list to accomplish this task.

On R4:

Note: In this case an extended access-list can be used to filter a prefix from a given route-source; in the following extended access-list configuration, the source address in the ACL references the advertising neighbor, whereas, the destination address in the ACL references the actual Network:

```
R4(config)#Access-list 100 deny ip host 10.1.1.3 host 200.1.1.0
R4(config)#Access-list 100 permit ip any any
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config)#Router eigrp 100
R4(config-router)#Distribute-list 100 in f0/0
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      2.0.0.0/8 [90/156160] via 10.1.1.2, 00:10:42, FastEthernet0/0
D      3.0.0.0/8 [90/156160] via 10.1.1.3, 00:01:25, FastEthernet0/0
       12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 10.1.1.2, 00:10:42, FastEthernet0/0
D     200.1.1.0/24 [90/156160] via 10.1.1.2, 00:00:52, FastEthernet0/0
D     200.2.2.0/24 [90/156160] via 10.1.1.3, 00:01:25, FastEthernet0/0
                   [90/156160] via 10.1.1.2, 00:01:25, FastEthernet0/0
```

Task 4

Configure R4 such that it takes R3 to reach network 200.2.2.0 /24. R4 should ONLY use R2 as the next-hop to reach network 200.2.2.0/24 when R3 is down. You should use an access-list to accomplish this task.

In this case the distance can be manipulated to accomplish this task:

Step 1:
Configure an access-list to identify the network:

```
R4(config)#Access-list 1 permit 200.2.2.0 0.0.0.255
```

Step 2:
Utilizing the distance command the AD for network 200.2.2.0/24 through R3 is set to a lower value than the default AD of 90:

```
R4(config)#Router eigrp 100
R4(config-router)#Distance 89 10.1.1.3 0.0.0.0 1
```

The above command sets the AD for network 200.2.2.0/24 through R3 to 89.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      2.0.0.0/8 [90/156160] via 10.1.1.2, 00:01:06, FastEthernet0/0
D      3.0.0.0/8 [90/156160] via 10.1.1.3, 00:01:06, FastEthernet0/0
       12.0.0.0/24 is subnetted, 1 subnets
D         12.1.1.0 [90/30720] via 10.1.1.2, 00:01:06, FastEthernet0/0
D      200.1.1.0/24 [90/156160] via 10.1.1.2, 00:01:06, FastEthernet0/0
D      200.2.2.0/24 [89/156160] via 10.1.1.3, 00:01:06, FastEthernet0/0
```

Task 5

Filter network 2.0.0.0/8 on R4; **do not** use the "distribute-list" command to accomplish this task.

Once again, the distance command can be used to accomplish this task. The difference between the solution used in this task and the solution used in the previous task is that the AD is set to a value that is unreachable (255).

On R4:

```
R4(config)#Access-list 2 permit 2.0.0.0 0.255.255.255

R4(config)#Router eigrp 100
R4(config-router)#Distance 255 10.1.1.2 0.0.0.0 2
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      3.0.0.0/8 [90/156160] via 10.1.1.3, 00:00:35, FastEthernet0/0
       12.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D        12.1.1.0 [90/30720] via 10.1.1.2, 00:00:35, FastEthernet0/0
D     200.1.1.0/24 [90/156160] via 10.1.1.2, 00:00:35, FastEthernet0/0
D     200.2.2.0/24 [89/156160] via 10.1.1.3, 00:00:35, FastEthernet0/0
```

Task 6

Configure R4 to filter network 3.0.0.0/8 to filter network 3.0.0.0, DO NOT use any of the previously used solution to resolve this task..

On R4:

```
R4(config)#Access-list 3 deny 3.0.0.0 0.255.255.255
R4(config)#Access-list 3 permit any

R4(config)#Router eigrp 100
R4(config-router)#Distribute-list 3 in F0/0
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/30720] via 10.1.1.2, 00:05:34, FastEthernet0/0
D     200.1.1.0/24 [90/156160] via 10.1.1.3, 00:00:30, FastEthernet0/0
                   [90/156160] via 10.1.1.2, 00:00:30, FastEthernet0/0
D     200.2.2.0/24 [89/156160] via 10.1.1.3, 00:04:48, FastEthernet0/0
```

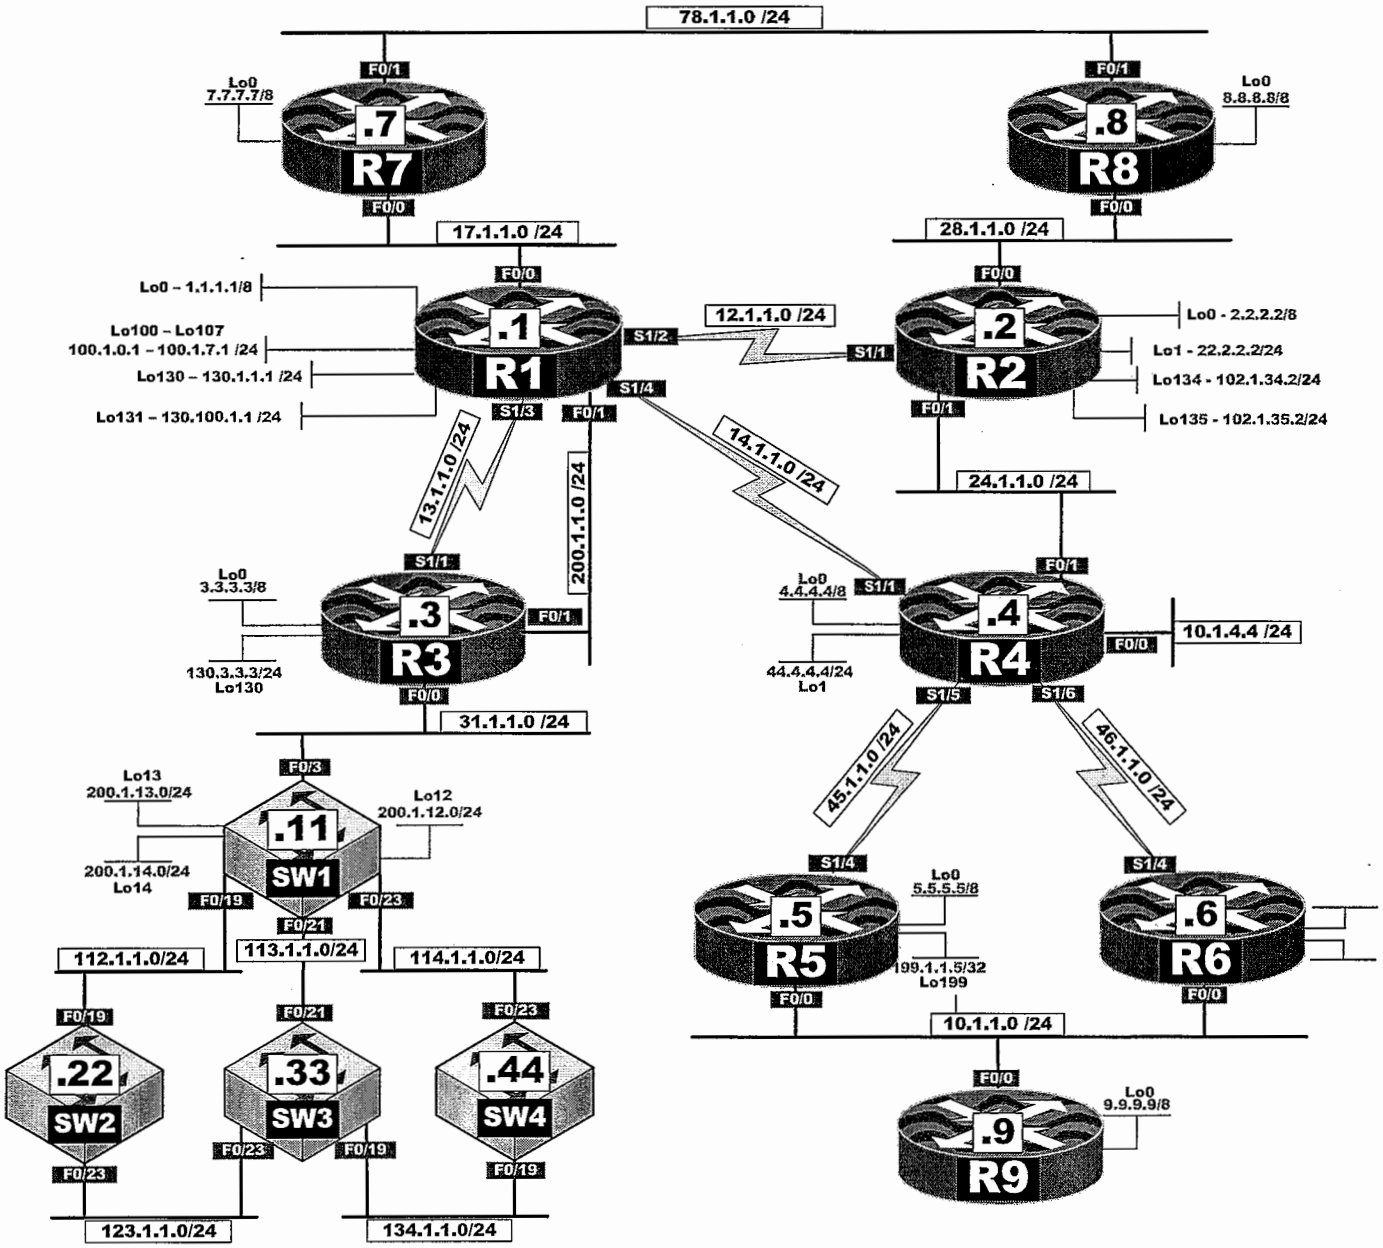
Task 7

Erase the startup config and reload the routers before proceeding to the next task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 5
EIGRP Mock Lab



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "EIGRP"→ "Lab-5".

Task 1

Configure the F0/0 interface of routers R5, R6 and R9 in AS 100; these routers should be configured to advertise their Lo0 and Lo199 interfaces in this AS using the following policy:

R9 and R6 should be configured to reach each other's loopback interface/s by going through R5. DO NOT use PBR or configure another AS to accomplish this task.

On R5:

RIPv2, Eigrp and BGP are classless routing protocols and by default, they have a Classful nature/behavior, meaning that they automatically summarize the networks that they advertise to their major network boundary. This behavior can be disabled by using the "No Auto-summary" router configuration command:

```
R5(config)#Router eigrp 100
R5(config-router)#Network 5.5.5.5 0.0.0.0
R5(config-router)#Network 199.1.1.5 0.0.0.0
R5(config-router)#Network 10.1.1.5 0.0.0.0
```

On R6:

```
R6(config)#Router eigrp 100
R6(config-router)#Network 6.6.6.6 0.0.0.0
R6(config-router)#Network 199.1.1.6 0.0.0.0
R6(config-router)#Network 10.1.1.6 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.5 (FastEthernet0/0) is
up: new adjacency
```

On R9:

```
R9(config)#Router eigrp 100
R9(config-router)#Network 9.9.9.9 0.0.0.0
R9(config-router)#Network 10.1.1.9 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

You should see the following two console messages stating that the local router has established two neighbor adjacency one with R5 and the second one with R6:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.5 (FastEthernet0/0) is
up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.6 (FastEthernet0/0) is
up: new adjacency
```

To verify the configuration:

On R5:

```
R5#Show ip eigrp neighbors

IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.1.9                Fa0/0            12 00:04:28    5   200  0  6
0   10.1.1.6                Fa0/0            11 00:08:44    2   200  0  5
```

To see which interfaces have EIGRP enabled:

```
R5#Sh ip eigrp interfaces

EIGRP-IPv4 Interfaces for AS(100)
                        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
Lo199              0        0/0        0       0/1           0           0
Fa0/0              2        0/0        3       0/1           50          0

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

D      6.0.0.0/8 [90/156160] via 10.1.1.6, 00:15:55, FastEthernet0/0
D      9.0.0.0/8 [90/156160] via 10.1.1.9, 00:11:38, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.6 [90/156160] via 10.1.1.6, 00:01:19, FastEthernet0/0
```

On R6:

```
R6#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(100)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.1.9                Fa0/0             11 00:07:15    4   200  0  6
0   10.1.1.5                Fa0/0             14 00:11:32 1276  5000  0  5

R6#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      5.0.0.0/8 [90/156160] via 10.1.1.5, 00:15:30, FastEthernet0/0
D      9.0.0.0/8 [90/156160] via 10.1.1.9, 00:11:13, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.5 [90/156160] via 10.1.1.5, 00:00:37, FastEthernet0/0
```

On R9:

```
R9#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(100)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   10.1.1.6                Fa0/0             13 00:08:22  816  4896  0  6
0   10.1.1.5                Fa0/0             10 00:08:22 1020  5000  0  6

R9#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      5.0.0.0/8 [90/156160] via 10.1.1.5, 00:10:44, FastEthernet0/0
D      6.0.0.0/8 [90/156160] via 10.1.1.6, 00:10:44, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.5 [90/156160] via 10.1.1.5, 00:00:08, FastEthernet0/0
D         199.1.1.6 [90/156160] via 10.1.1.6, 00:00:25, FastEthernet0/0
```

In order for R9 and R6 to go through R5 to reach each other's loopback interfaces R9 should have R5 as the next hop router to reach networks 6.0.0.0 /8 and 199.1.1.6/32, and R6 should have R5 as the next hop to reach networks 9.0.0.0 /8; To accomplish this task R6 and R9 can NOT have a neighbor adjacency, or else they will exchange routes and advertise themselves as the next hop for the routes they originate, therefore, R6 and R9 should both be configured with a neighbor command pointing to R5; R5 should also be configured to have a neighbor command for R6 and another for R9:

On R5:

```
R5(config)#Router eigrp 100
R5(config-router)#Neighbor 10.1.1.6 F0/0
R5(config-router)#Neighbor 10.1.1.9 F0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

NOTE: In order for two Eigrp routers to establish neighbor adjacency, they MUST both use Unicast or Multicast; you can NOT have one end using multicast and the other end using unicast, or if the local router has two Eigrp neighbors, it can NOT establish a Unicast neighbor adjacency to one and establish Multicast neighbor adjacency with the second one, unless the link connecting them is point-to-point.

In the following configuration R6 is configured to establish neighbor adjacency with R5 using Unicast, because it is configured to use Unicast, it will drop adjacency to R9, This is because all three routers are connected to the same Ethernet segment.

When configuring Neighbor command you MUST specify the interface through which the Neighbor adjacency will be established.

On R6 and R9:

```
Rx(config)#Router eigrp 100
Rx(config-router)#Neighbor 10.1.1.5 F0/0
```

To verify the configuration:

On R6:

```
R6#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     5.0.0.0/8 [90/156160] via 10.1.1.5, 00:00:40, FastEthernet0/0
      199.1.1.0/32 is subnetted, 2 subnets
D        199.1.1.5 [90/156160] via 10.1.1.5, 00:00:40, FastEthernet0/0
```

On R9:

```
R9#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     5.0.0.0/8 [90/156160] via 10.1.1.5, 00:01:06, FastEthernet0/0
      199.1.1.0/32 is subnetted, 1 subnets
D        199.1.1.5 [90/156160] via 10.1.1.5, 00:01:06, FastEthernet0/0
```

On R5:

```
R5#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     6.0.0.0/8 [90/156160] via 10.1.1.6, 00:01:51, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D      9.0.0.0/8 [90/156160] via 10.1.1.9, 00:01:37, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.6 [90/156160] via 10.1.1.6, 00:01:51, FastEthernet0/0
```

R5 sees both routes in its routing table whereas, R9 and R6 can ONLY see R5's routes. The reason for this behavior is "Split-Horizon"; this rule states not to advertise a route out of the interface through which it was learned.

The following command disables the split-horizon for Eigrp AS 100:

On R5:

```
R5(config)#Int F0/0
R5(config-if)#No IP split-horizon eigrp 100
```

You should see the following console messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.9 (FastEthernet0/0) is
resync: split horizon changed

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.6 (FastEthernet0/0) is
resync: split horizon changed
```

To test the configuration:

On R6:

```
R6#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      5.0.0.0/8 [90/156160] via 10.1.1.5, 00:07:27, FastEthernet0/0
D      9.0.0.0/8 [90/158720] via 10.1.1.5, 00:01:31, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.5 [90/156160] via 10.1.1.5, 00:07:27, FastEthernet0/0
```

On R9:

```
R9#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      5.0.0.0/8 [90/156160] via 10.1.1.5, 00:07:44, FastEthernet0/0
D      6.0.0.0/8 [90/158720] via 10.1.1.5, 00:02:02, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.5 [90/156160] via 10.1.1.5, 00:07:44, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D        199.1.1.6 [90/158720] via 10.1.1.5, 00:02:02, FastEthernet0/0
```

Task 2

Configure the serial interfaces of R4 (S1/5, and S1/6), R5's S1/4, and R6's S1/4 in AS 200.

On R4:

```
R4(config)#Router eigrp 200
R4(config-router)#Network 45.1.1.4 0.0.0.0
R4(config-router)#Network 46.1.1.4 0.0.0.0
```

On R5:

```
R5(config)#Router eigrp 200
R5(config-router)#Network 45.1.1.5 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 200: Neighbor 45.1.1.4 (Serial1/4) is up:
new adjacency
```

On R6:

```
R6(config)#Router eigrp 200
R6(config-router)#Network 46.1.1.6 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 200: Neighbor 46.1.1.4 (Serial1/4) is up:
new adjacency
```

To verify the configuration:

On R4:

```
R4#Show ip eigrp neighbors

R4#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(200)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
H   Address                 Interface       Hold Uptime   SRTT    RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   46.1.1.6                Se1/6            13 00:02:49   26   1140  0  3
0   45.1.1.5                Se1/5            12 00:03:31   20   1140  0  3
```

Task 3

Configure a mutual redistribution on R5 between the two ASes (100 and 200) and ensure that R4 and R9 receive all routes.

When redistributing between two Eigrp routing protocols, it is NOT necessary to configure the "Bandwidth, Delay, Reliability, Load and MTU", unless otherwise stated.

On R5:

```
R5(config)#Router Eigrp 100
R5(config-router)#redistribute eigrp 200

R5(config)#Router Eigrp 200
R5(config-router)#redistribute eigrp 100
```

To verify the configuration:

On R9:

```
R9#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      5.0.0.0/8 [90/156160] via 10.1.1.5, 00:23:22, FastEthernet0/0
D      6.0.0.0/8 [90/158720] via 10.1.1.5, 00:17:40, FastEthernet0/0
       45.0.0.0/24 is subnetted, 1 subnets
D EX      45.1.1.0 [170/20514560] via 10.1.1.5, 00:00:35, FastEthernet0/0
       46.0.0.0/24 is subnetted, 1 subnets
D EX      46.1.1.0 [170/21026560] via 10.1.1.5, 00:00:35, FastEthernet0/0
       199.1.1.0/32 is subnetted, 2 subnets
D         199.1.1.5 [90/156160] via 10.1.1.5, 00:23:22, FastEthernet0/0
D         199.1.1.6 [90/158720] via 10.1.1.5, 00:17:40, FastEthernet0/0
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D EX  5.0.0.0/8 [170/20640000] via 45.1.1.5, 00:01:42, Serial1/5
D EX  6.0.0.0/8 [170/20642560] via 45.1.1.5, 00:01:42, Serial1/5
D EX  9.0.0.0/8 [170/20642560] via 45.1.1.5, 00:01:42, Serial1/5
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     10.1.1.0/24 [170/20514560] via 45.1.1.5, 00:01:42, Serial1/5
      199.1.1.0/32 is subnetted, 2 subnets
D EX     199.1.1.5 [170/20640000] via 45.1.1.5, 00:01:42, Serial1/5
D EX     199.1.1.6 [170/20642560] via 45.1.1.5, 00:01:42, Serial1/5
```

Task 4

Configure the following serial interfaces in AS 300:

R1 - S1/2, S1/3, and S1/4
R2 – S1/1
R3 – S1/1
R4 – S1/1

On R1:

```
R1(config)#Router eigrp 300
R1(config-router)#Network 12.1.1.1 0.0.0.0
R1(config-router)#Network 13.1.1.1 0.0.0.0
R1(config-router)#Network 14.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router Eigrp 300
R2(config-router)#Network 12.1.1.2 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 300: Neighbor 12.1.1.1 (Serial1/1) is up:
new adjacency
```

On R3:

```
R3(config)#Router Eigrp 300
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config-router)#Network 13.1.1.3 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 300: Neighbor 13.1.1.1 (Serial1/1) is up:
new adjacency
```

On R4:

```
R4(config)#Router eigrp 300
R4(config-router)#Network 14.1.1.4 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 300: Neighbor 14.1.1.1 (Serial1/1) is up:
new adjacency
```

To verify the configuration

On R1:

```
R1#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(300)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
2   14.1.1.4                Se1/4            11 00:00:51   28   1140  0  3
1   13.1.1.3                Se1/3            14 00:02:04   26   1140  0  3
0   12.1.1.2                Se1/2            13 00:03:16   26   1140  0  3
```

On R3:

```
R3#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(300)

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   13.1.1.1                Se1/1            14 00:02:55   27   1140  0  4
```

On R4:

```
R4#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(200)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                                                (sec)          (ms)       Cnt Num
1   46.1.1.6                  Se1/6               14 00:14:25    32   1140  0   3
0   45.1.1.5                  Se1/5               14 00:15:06    28   1140  0   4

EIGRP-IPv4 Neighbors for AS(300)
H   Address                   Interface       Hold Uptime    SRTT   RTO  Q  Seq
                                              (sec)          (ms)       Cnt Num
0   14.1.1.1                  Se1/1            12 00:02:44 1600   5000  0  6
```

Task 5

Configure R1 in AS 300 to summarize its Lo100 – 107. The summary route should be advertised to R2, R3 and R4 based on the following policy:

1. R2 should receive the summary ONLY.
2. R3 should receive the summary plus Network 100.1.3.0 /24.
3. R4 should receive the summary plus all the specific routes.
4. Configure minimum number of "IP Summary-address" commands possible to accomplish this task.

To summarize in any routing protocol, the specific routes must be known to that routing protocol. The first step is to verify that the Loopback interfaces exist:

On R1:

```
R1#Sh run | Inc interface Loopback10|ip address 100.

interface Loopback100
 ip address 100.1.0.1 255.255.255.0
interface Loopback101
 ip address 100.1.1.1 255.255.255.0
interface Loopback102
 ip address 100.1.2.1 255.255.255.0
interface Loopback103
 ip address 100.1.3.1 255.255.255.0
interface Loopback104
 ip address 100.1.4.1 255.255.255.0
interface Loopback105
 ip address 100.1.5.1 255.255.255.0
interface Loopback106
 ip address 100.1.6.1 255.255.255.0
interface Loopback107
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
 ip address 100.1.7.1 255.255.255.0
```

In the second step, we have to run EIGRP on these loopback interfaces:

```
R1(config)#Router eigrp 300
R1(config-router)#Netw 100.1.0.1 0.0.0.0
R1(config-router)#Netw 100.1.1.1 0.0.0.0
R1(config-router)#Netw 100.1.2.1 0.0.0.0
R1(config-router)#Netw 100.1.3.1 0.0.0.0
R1(config-router)#Netw 100.1.4.1 0.0.0.0
R1(config-router)#Netw 100.1.5.1 0.0.0.0
R1(config-router)#Netw 100.1.6.1 0.0.0.0
R1(config-router)#Netw 100.1.7.1 0.0.0.0
```

Based on the requirements of this task, a single "summary-address" can be configured for R1's connection to R2; this should satisfy the first requirement.

On R1:

```
R1(config)#Int S1/2
R1(config-if)#IP summary-address eigrp 300 100.1.0.0 255.255.248.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp 300 | B Gate
Gateway of last resort is not set

      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/21024000] via 12.1.1.1, 02:18:59, Serial1/1
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/21024000] via 12.1.1.1, 02:18:58, Serial1/1
      100.0.0.0/21 is subnetted, 1 subnets
D        100.1.0.0 [90/20640000] via 12.1.1.1, 00:00:49, Serial1/1
```

To accomplish the second and the third requirements of this task, a "Leak-map" MUST be configured. Let's configure the "Leak-map" and verify:

On R1:

An access-list is configured to permit the prefix that should be advertised with the summary route to R3:

```
R1(config)#Access-list 1 permit 100.1.3.0 0.0.0.255
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

A route-map is configured to reference the access-list configured in the previous step:

```
R1(config)#Route-map TST3 permit 10
R1(config-route-map)#match ip addr 1
```

The summary route is configured under the S1/3 interface with the leak-map referencing the route-map:

```
R1(config)#Int S1/3
R1(config-if)#IP summary-add eigrp 300 100.1.0.0 255.255.248.0 leak-map
TST3
```

To verify the configuration:

On R3:

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

      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/21024000] via 13.1.1.1, 00:03:10, Serial1/1
      14.0.0.0/24 is subnetted, 1 subnets
D        14.1.1.0 [90/21024000] via 13.1.1.1, 00:03:06, Serial1/1
      100.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D        100.1.0.0/21 [90/20640000] via 13.1.1.1, 00:00:20, Serial1/1
D        100.1.3.0/24 [90/20640000] via 13.1.1.1, 00:03:10, Serial1/1
```

To configure the third policy, the behavior of Access-list/Prefix-list, Route-map and the Leak-map is utilized.
If a leak-map references a route-map that does Not exist, ONLY the summary is advertised, but if a leak-map references a route-map that does exist and the route-map is referencing an access-list, the network/s that are permited in the acess-list will be advertised with the summary.
But if the leak-map is referencing a route-map that does exist, but the route-map is NOT referencing an access-list, or it's referencing an access-list that does NOT exist, the summary plus all the specific routes are advertised. Let's configure and verify this behavior:

On R1:

```
R1(config)#Route-map TST4 permit 10

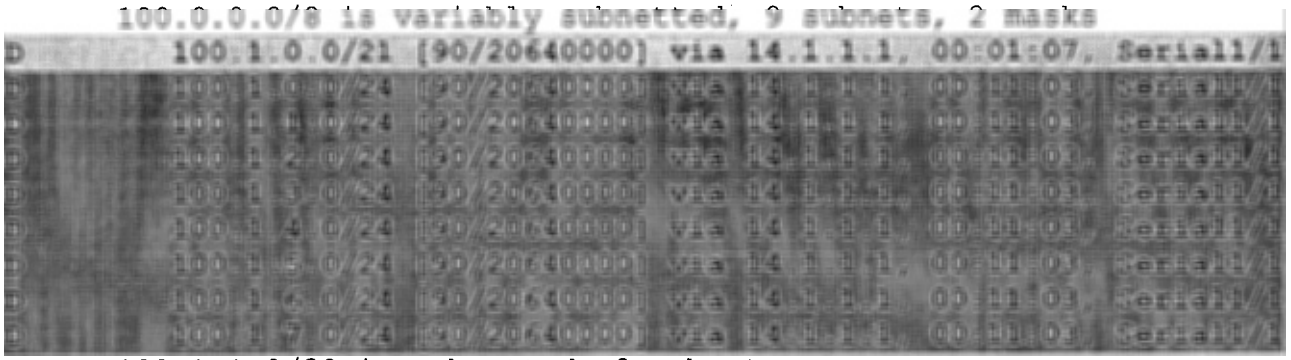
R1(config)#Int S1/4
R1(config-if)#IP summary-addr eig 300 100.1.0.0 255.255.248.0 leak-map TST4
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D EX  5.0.0.0/8 [170/20640000] via 45.1.1.5, 03:36:45, Serial1/5
D EX  6.0.0.0/8 [170/20642560] via 45.1.1.5, 03:36:45, Serial1/5
D EX  9.0.0.0/8 [170/20642560] via 45.1.1.5, 03:36:45, Serial1/5
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     10.1.1.0/24 [170/20514560] via 45.1.1.5, 03:36:45, Serial1/5
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/21024000] via 14.1.1.1, 00:11:03, Serial1/1
      13.0.0.0/24 is subnetted, 1 subnets
D        13.1.1.0 [90/21024000] via 14.1.1.1, 00:11:03, Serial1/1
      100.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D        100.1.0.0/21 [90/20640000] via 14.1.1.1, 00:01:07, Serial1/1

      199.1.1.0/32 is subnetted, 2 subnets
D EX     199.1.1.5 [170/20640000] via 45.1.1.5, 03:36:45, Serial1/5
D EX     199.1.1.6 [170/20642560] via 45.1.1.5, 03:36:45, Serial1/5
```

Task 6

R1 – R4 should be configured to advertise their Lo0 interface in AS 300.

On R1:

```
R1(config)#Router eigrp 300
R1(config-router)#Netw 1.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp 300
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Netw 2.2.2.2 0.0.0.0
```

On R3:

```
R3(config)#Router eigrp 300
R3(config-router)#Netw 3.3.3.3 0.0.0.0
```

On R4:

```
R4(config)#Router eigrp 300
R4(config-router)#Netw 4.4.4.4 0.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      2.0.0.0/8 [90/20640000] via 12.1.1.2, 00:00:42, Serial1/2
D      3.0.0.0/8 [90/20640000] via 13.1.1.3, 00:00:31, Serial1/3
D      4.0.0.0/8 [90/20640000] via 14.1.1.4, 00:00:19, Serial1/4
       100.0.0.0/8 is variably subnetted, 17 subnets, 3 masks
D         100.1.0.0/21 is a summary, 00:06:59, Null0
```

On R2:

```
R2#Sh ip route eigrp | B Gate
Gateway of last resort is not set

D      1.0.0.0/8 [90/20640000] via 12.1.1.1, 00:01:27, Serial1/1
D      3.0.0.0/8 [90/21152000] via 12.1.1.1, 00:00:56, Serial1/1
D      4.0.0.0/8 [90/21152000] via 12.1.1.1, 00:00:44, Serial1/1
       13.0.0.0/24 is subnetted, 1 subnets
D         13.1.1.0 [90/21024000] via 12.1.1.1, 00:20:45, Serial1/1
       14.0.0.0/24 is subnetted, 1 subnets
D         14.1.1.0 [90/21024000] via 12.1.1.1, 00:20:41, Serial1/1
       100.0.0.0/21 is subnetted, 1 subnets
D         100.1.0.0 [90/20640000] via 12.1.1.1, 00:20:50, Serial1/1
```

On R3:

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


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D      1.0.0.0/8 [90/20640000] via 13.1.1.1, 00:02:04, Serial1/1

D      2.0.0.0/8 [90/21152000] via 13.1.1.1, 00:01:45, Serial1/1
D      4.0.0.0/8 [90/21152000] via 13.1.1.1, 00:01:22, Serial1/1
       12.0.0.0/24 is subnetted, 1 subnets
D         12.1.1.0 [90/21024000] via 13.1.1.1, 00:21:23, Serial1/1
       14.0.0.0/24 is subnetted, 1 subnets
D         14.1.1.0 [90/21024000] via 13.1.1.1, 00:21:19, Serial1/1
       100.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D         100.1.0.0/21 [90/20640000] via 13.1.1.1, 00:18:33, Serial1/1
D         100.1.3.0/24 [90/20640000] via 13.1.1.1, 00:21:23, Serial1/1
```

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      1.0.0.0/8 [90/20640000] via 14.1.1.1, 00:02:35, Serial1/1
D      2.0.0.0/8 [90/21152000] via 14.1.1.1, 00:02:16, Serial1/1
D      3.0.0.0/8 [90/21152000] via 14.1.1.1, 00:02:04, Serial1/1
D EX   5.0.0.0/8 [170/20640000] via 45.1.1.5, 03:47:31, Serial1/5
D EX   6.0.0.0/8 [170/20642560] via 45.1.1.5, 03:47:31, Serial1/5
D EX   9.0.0.0/8 [170/20642560] via 45.1.1.5, 03:47:31, Serial1/5
       10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX      10.1.1.0/24 [170/20514560] via 45.1.1.5, 03:47:31, Serial1/5
       12.0.0.0/24 is subnetted, 1 subnets
D         12.1.1.0 [90/21024000] via 14.1.1.1, 00:21:49, Serial1/1
       13.0.0.0/24 is subnetted, 1 subnets
D         13.1.1.0 [90/21024000] via 14.1.1.1, 00:21:49, Serial1/1
       100.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D         100.1.0.0/21 [90/20640000] via 14.1.1.1, 00:11:53, Serial1/1
D         100.1.0.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.1.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.2.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.3.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.4.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.5.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.6.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
D         100.1.7.0/24 [90/20640000] via 14.1.1.1, 00:21:49, Serial1/1
       199.1.1.0/32 is subnetted, 2 subnets
D EX      199.1.1.5 [170/20640000] via 45.1.1.5, 03:47:31, Serial1/5
D EX      199.1.1.6 [170/20642560] via 45.1.1.5, 03:47:31, Serial1/5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 7

Configure a mutual redistribution between the two ASes (200 and 300) on R4.

On R4:

To verify the two ASes running on this router:

```
R4#Show ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 200"
Routing Protocol is "eigrp 300"
```

To redistribute between the two ASes:

```
R4(config)#router eigrp 200
R4(config-router)#redistribute eigrp 300

R4(config)#router eigrp 300
R4(config-router)#redistribute eigrp 200
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     2.0.0.0/8 [90/20640000] via 12.1.1.2, 00:04:49, Serial1/2
D     3.0.0.0/8 [90/20640000] via 13.1.1.3, 00:04:38, Serial1/3
D     4.0.0.0/8 [90/20640000] via 14.1.1.4, 00:04:26, Serial1/4
D EX  5.0.0.0/8 [170/21152000] via 14.1.1.4, 00:00:18, Serial1/4
D EX  6.0.0.0/8 [170/21154560] via 14.1.1.4, 00:00:18, Serial1/4
D EX  9.0.0.0/8 [170/21154560] via 14.1.1.4, 00:00:18, Serial1/4
      10.0.0.0/24 is subnetted, 1 subnets
D EX     10.1.1.0 [170/21026560] via 14.1.1.4, 00:00:18, Serial1/4
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/21024000] via 14.1.1.4, 00:00:18, Serial1/4
      46.0.0.0/24 is subnetted, 1 subnets
D EX     46.1.1.0 [170/21024000] via 14.1.1.4, 00:00:18, Serial1/4
      100.0.0.0/8 is variably subnetted, 17 subnets, 3 masks
D        100.1.0.0/21 is a summary, 00:11:06, Null0
      199.1.1.0/32 is subnetted, 2 subnets
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D EX      199.1.1.5 [170/21152000] via 14.1.1.4, 00:00:18, Serial1/4

D EX      199.1.1.6 [170/21154560] via 14.1.1.4, 00:00:18, Serial1/4
```

Task 8

Configure R2 to summarize its Loopback 134 and 135 and advertise a single summary in AS 300.

On R2:

```
R2#Sh ip int br | Inc Loopback13[4|5]

Loopback134                102.1.34.2      YES manual up                    up
Loopback135                102.1.35.2      YES manual up                    up
```

To configure the summary:

On R2:

```
R2(config)#Router eigrp 300
R2(config-router)#Netw 102.1.35.2 0.0.0.0
R2(config-router)#Netw 102.1.34.2 0.0.0.0
```

To summarize:

We can see that the first two octets of these two prefixes match (102.1.), the third octet is where they differ, if the third octet of these two prefixes are converted into binary we should see the following:

	128	64	32	16	8	4	2	1
34	0	0	1	0	0	0	1	0
35	0	0	1	0	0	0	1	1
	M	M	M	M	M	M	M	DM

M = Match
DM = Don't Match

If the last bit of these two prefixes are compared, you will see that the last bit can be either a "0" or a "1" which covers all possibilities; therefore, the Don't care bit is the last bit; as a result, the summary route will be 102.1.34.0 0.0.1.0 or 255.255.254.0:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#Int S1/1

R2(config-subif)#IP summary-address eigrp 300 102.1.34.0 255.255.254.0
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | Inc 102

D EX      10.1.1.0 [170/21026560] via 14.1.1.4, 00:07:04, Serial1/4
D EX      45.1.1.0 [170/21024000] via 14.1.1.4, 00:07:04, Serial1/4
D EX      46.1.1.0 [170/21024000] via 14.1.1.4, 00:07:04, Serial1/4
      102.0.0.0/23 is subnetted, 1 subnets
D        102.1.34.0 [90/20640000] via 12.1.1.2, 00:01:33, Serial1/2
```

To test the configuration:

On R1:

```
R1#Ping 102.1.34.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 102.1.34.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms

R1#Ping 102.1.35.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 102.1.35.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms
```

Task 9

Configure Authentication on R1, R2 and R4 using the following policy:

- The connection between R1 and R2 should be authenticated using "R1R2" as the password.
- The connection between R1 and R4 should be authenticated using "R1R4" as the password.
- You should MD5 authentication to accomplish this task.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

In the first step, R1 and R2 are configured for authentication; note the key chain names are locally significant:

On R1:

```
R1(config)#Key chain R1R2
R1(config-keychain)#key 1
R1(config-keychain-key)#key-string R1R2

R1(config)#Int S1/2
R1(config-if)#IP authentication key-chain eigrp 300 R1R2
R1(config-if)#IP authentication mode eigrp 300 MD5
```

To verify neighbor adjacency:

On R1:

```
R1#Show ip eigrp neighbors
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
2	14.1.1.4	Se1/4	10	00:34:05	26	1140	0	38
1	13.1.1.3	Se1/3	11	00:34:09	25	1140	0	49

On R2:

```
R2(config)#Key chain R2R1
R2(config-keychain)#Key 1
R2(config-keychain-key)#Key-string R1R2

R2(config)#Int S1/1
R2(config-if)#IP authentication key-chain eigrp 300 R2R1
R2(config-if)#IP authentication mode eigrp 300 MD5
```

You should get the following console message:

```
IP-EIGRP(0) 300: Neighbor 12.1.1.1 (Serial1/1) is up: new adjacency
```

To verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Show ip eigrp neighbors
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   12.1.1.1                Se1/1            11 00:01:17  119  1140  0  47
```

In the last item of this task, authentication is configured between R1 and R4:

On R1:

```
R1(config)#Key chain R1R4
R1(config-keychain)#key 1
R1(config-keychain-key)#Key-string R1R4

R1(config)#int S1/4
R1(config-if)#IP authentication key-chain eigrp 300 R1R4
R1(config-if)#IP authentication mode eigrp 300 MD5
```

On R4:

```
R4(config)#Key chain R1R4
R4(config-keychain)#Key 1
R4(config-keychain-key)#Key-String R1R4

R4(config)#Int S1/1
R4(config-if)#IP Authentication key-chain eigrp 300 R1R4
R4(config-if)#IP authentication mode eigrp 300 MD5
```

To verify the configuration:

On R4:

```
R4#Show ip eigrp 300 neighbors

EIGRP-IPv4 Neighbors for AS(300)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   14.1.1.1                Se1/1            12 00:00:20  106  1140  0  54
```

Task 10

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Configure the F0/1 interfaces of R1 and R3 in Eigrp AS 500, use the following policy to advertise EIGRP on their loopback interfaces:

1. R1 should be configured to advertise its loopback 130 using a route-map such that a "Show ip route eigrp 500" on R3 resembles the following, DO NOT use an access-list or a prefix-list to accomplish this task:

   D EX 130.1.1.0 [170/156160] via 200.1.1.1, 00:00:19, FastEthernet0/1

2. R1 should also advertise its Loopback 131 in this routing domain such that the network is injected as an external route, you should use an access-list and a route-map to accomplish this task.

3. R3 should be configured to advertise its loopback 130 such that a "Show ip route eigrp 500" on R1 resembles the following:

   D EX 130.3.0.0/16 [170/156160] via 200.1.1.3, 00:00:30, FastEthernet0/1

   R3 should use "Redistribute static" to accomplish this task. DO NOT configure a Static route to accomplish this task:

The neighbor adjacency of these two routers should be configured first:

On R1:

```
R1(config)#Router eigrp 500
R1(config-router)#Network 200.1.1.1 0.0.0.0
```

On R3:

```
R3(config)#Router eigrp 500
R3(config-router)#Network 200.1.1.3 0.0.0.0
```

To verify the configuration:

On R3:

```
R3#Show ip eigrp 500 neighbors

EIGRP-IPv4 Neighbors for AS(500)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   200.1.1.1               Fa0/1            12 00:00:31    1  4500  0  1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Show ip eigrp 500 neighbors

EIGRP-IPv4 Neighbors for AS(500)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   200.1.1.3               Fa0/1            11 00:01:13 1996  5000  0  1
```

Now that R1 and R3 have established an adjacency, we can proceed with the rest of the policies of this task:

On R1:

```
R1(config)#Route-map Task10 permit 10
R1(config-route-map)#match int lo130

R1(config-route-map)#Router eigrp 500
R1(config-router)#redistribute connected route-map Task10
```

To verify the configuration:

On R3:

```
R3#Show ip route eigrp 500 | B Gate
Gateway of last resort is not set

      130.1.0.0/24 is subnetted, 1 subnets
D EX     130.1.1.0 [170/156160] via 200.1.1.1, 00:00:19, FastEthernet0/1
```

To accomplish the second item:

An access-list is configured, BUT before configuring any access-list a show access-list is highly recommended so the access-list number of the existing access-list is NOT used in the new one:

On R1:

```
R1#Show access-list

Standard IP access list 1

    10 permit 100.1.3.0, wildcard bits 0.0.0.255 (2 matches)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config)#Access-list 2 permit 130.100.1.0 0.0.0.255
```

Since we can NOT have multiple redistribute connected commands, the new access-list is referenced in the same route-map using permit 20:

```
R1(config)#route-map Task10 permit 20
R1(config-route-map)#match ip addr 2
```

NOTE: R1 is already configured with redistribute connected referencing the route-map, therefore, no need to do it again.

To verify the configuration:

On R3:

```
R3#Show ip route eigrp 500 | B Gate
Gateway of last resort is not set

      130.1.0.0/24 is subnetted, 1 subnets
D EX    130.1.1.0 [170/156160] via 200.1.1.1, 00:06:19, FastEthernet0/1
      130.100.0.0/24 is subnetted, 1 subnets
D EX    130.100.1.0 [170/156160] via 200.1.1.1, 00:01:21, FastEthernet0/1
```

To accomplish the last item:

This task can be tricky, but since a static route can NOT be used, the behavior of a command is used to accomplish this task. If the "IP default-network" command references a subnetted network, IOS will automatically inject a static route, let's test this:

On R3:

```
R3(config)#IP default-network 130.3.3.0
```

To verify the configuration:

On R3:

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

      130.3.0.0/16 is variably subnetted, 3 subnets, 3 masks

S        130.3.0.0/16 [1/0] via 130.3.3.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Now that we see the default route automatically configured, we can use the redistribute static command to accomplish this task, in the prior version of IOS you could actually do a "Show run | Inc ip route"

command and see the static route, but in "c2800nm-adventerprisek9-mz.151-4.M7.bin" you can NOT see the actual static route:

On R3:

```
R3(config)#Router eigrp 500
R3(config-router)#Redistribute static
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp 500 | B Gate
Gateway of last resort is not set

D EX  130.3.0.0/16 [170/156160] via 200.1.1.3, 00:00:30, FastEthernet0/1
```

Task 11

Configure R2 and R4's F0/1, and the F0/0 interface of R4 in AS 400.

On R2:

```
R2(config)#Router eigrp 400
R2(config-router)#Netw 24.1.1.2 0.0.0.0
```

On R4:

```
R4(config)#Router eigrp 400
R4(config-router)#Netw 24.1.1.4 0.0.0.0
R4(config-router)#Netw 10.1.4.4 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp 400 | B Gate
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Gateway of last resort is not set

      10.0.0.0/24 is subnetted, 2 subnets

D        10.1.4.0 [90/30720] via 24.1.1.4, 00:01:08, FastEthernet0/1
```

Task 12

Redistribute AS 400 into the ASes configured on R2 and R4.

To see the ASes that are configured on R2 and R4:

On R4:

```
R4#Show ip proto | Inc Routing Protocol

Routing Protocol is "eigrp 200"
Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 400"

R4(config)#router eigrp 200
R4(config-router)#Redistribute eigrp 400

R4(config)#router eigrp 300
R4(config-router)#Redistribute eigrp 400
```

On R2:

```
R2#Show ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 400"

R2(config)#Router eigrp 300
R2(config-router)#Redistribute eigrp 400
```

To verify the configuration:

On R9:

```
R9#Show ip route eigrp | Inc EX
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D EX  1.0.0.0/8 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX  2.0.0.0/8 [170/21666560] via 10.1.1.5, 02:16:31, FastEthernet0/0

D EX  3.0.0.0/8 [170/21666560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX  4.0.0.0/8 [170/20642560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     10.1.4.0/24 [170/20517120] via 10.1.1.5, 00:00:59, FastEthernet0/0
D EX     12.1.1.0 [170/21538560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     13.1.1.0 [170/21538560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     14.1.1.0 [170/21026560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     24.1.1.0 [170/20517120] via 10.1.1.5, 00:00:59, FastEthernet0/0
D EX     45.1.1.0 [170/20514560] via 10.1.1.5, 02:16:32, FastEthernet0/0
D EX     46.1.1.0 [170/21026560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.0.0/21 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.0.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.1.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.2.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.3.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.4.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.5.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.6.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     100.1.7.0/24 [170/21154560] via 10.1.1.5, 02:16:31, FastEthernet0/0
D EX     102.1.34.0 [170/21666560] via 10.1.1.5, 02:16:31, FastEthernet0/0
```

Task 13

Configure the F0/1 interfaces of R7 and R8 in AS 600 and advertise their Lo0 interfaces in this AS such that they appear as external routes.

On R7:

```
R7(config)#Router eigrp 600
R7(config-router)#Netw 78.1.1.7 0.0.0.0

R7(config)#Route-map TST permit 10
R7(config-route-map)#match int lo0

R7(config-route-map)#Router eigrp 600
R7(config-router)#Redistribute connected route-map TST
```

On R8:

```
R8(config)#router eigrp 600
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R8(config-router)#Netw 78.1.1.8 0.0.0.0

R8(config)#Route-map TST permit 10

R8(config-route-map)#Match int lo0

R8(config-route-map)#Router eigrp 600
R8(config-router)#Redistribute connected route-map TST
```

To verify the configuration:

On R7:

```
R7#Show ip route eigrp 600 | B Gate
Gateway of last resort is not set

D EX  8.0.0.0/8 [170/156160] via 78.1.1.8, 00:01:04, FastEthernet0/1
```

On R8:

```
R8#Show ip route eigrp 600 | B Gate
Gateway of last resort is not set

D EX  7.0.0.0/8 [170/156160] via 78.1.1.7, 00:02:05, FastEthernet0/1
```

Task 14

Configure the F0/0 interface of R8, R2, R7 and R1 in AS 600.

On R1:

```
R1(config)#Router eigrp 600
R1(config-router)#Netw 17.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp 600
R2(config-router)#Network 28.1.1.2 0.0.0.0
```

On R7:

```
R7(config)#Router eigrp 600
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7(config-router)#Netw 17.1.1.7 0.0.0.0
```

On R8:

```
R8(config)#router eigrp 600
R8(config-router)#Netw 28.1.1.8 0.0.0.0
```

To verify the configuration;

On R1:

```
R1#Show ip route eigrp 600 | B Gate
Gateway of last resort is not set

D EX  7.0.0.0/8 [170/156160] via 17.1.1.7, 00:00:44, FastEthernet0/0
D EX  8.0.0.0/8 [170/158720] via 17.1.1.7, 00:00:44, FastEthernet0/0
      28.0.0.0/24 is subnetted, 1 subnets
D        28.1.1.0 [90/33280] via 17.1.1.7, 00:00:44, FastEthernet0/0
      78.0.0.0/24 is subnetted, 1 subnets
D        78.1.1.0 [90/30720] via 17.1.1.7, 00:00:44, FastEthernet0/0
```

On R2:

```
R2#Show ip route eigrp 600 | B Gate
Gateway of last resort is not set

D EX  7.0.0.0/8 [170/158720] via 28.1.1.8, 00:02:34, FastEthernet0/0
D EX  8.0.0.0/8 [170/156160] via 28.1.1.8, 00:02:34, FastEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D        17.1.1.0 [90/33280] via 28.1.1.8, 00:01:16, FastEthernet0/0
      78.0.0.0/24 is subnetted, 1 subnets

D        78.1.1.0 [90/30720] via 28.1.1.8, 00:02:34, FastEthernet0/0
```

Task 15

Redistribute mutually between AS 600 and the Eigrp ASes running on R1, and R2.

On R1:

```
R1#Sh ip protocol | Inc Routing Protocol
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 500"
Routing Protocol is "eigrp 600"

R1(config)#Router eigrp 300
R1(config-router)#redistribute eigrp 600

R1(config)#Router eigrp 500
R1(config-router)#redistribute eigrp 600

R1(config)#Router eigrp 600
R1(config-router)#redistribute eigrp 500
R1(config-router)#redistribute eigrp 300
```

On R2:

```
R2#Show ip protocol | Inc Routing Protocol

Routing Protocol is "eigrp 300"
Routing Protocol is "eigrp 400"
Routing Protocol is "eigrp 600"

R2(config)#Router eigrp 300
R2(config-router)#redistribute eigrp 600

R2(config)#Router eigrp 400
R2(config-router)#redistribute eigrp 600

R2(config)#router eigrp 600
R2(config-router)#redistribute eigrp 300
R2(config-router)#redistribute eigrp 400
```

To verify the configuration:

On R9:

```
R9#Show ip route eigrp | Inc EX

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
D EX  1.0.0.0/8 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX  2.0.0.0/8 [170/21666560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX  3.0.0.0/8 [170/21666560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX  4.0.0.0/8 [170/20642560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX  7.0.0.0/8 [170/20650240] via 10.1.1.5, 00:01:06, FastEthernet0/0
D EX  8.0.0.0/8 [170/20647680] via 10.1.1.5, 00:01:06, FastEthernet0/0
D EX     10.1.4.0/24 [170/20517120] via 10.1.1.5, 00:31:27, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D EX     12.1.1.0 [170/21538560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     13.1.1.0 [170/21538560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     14.1.1.0 [170/21026560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     17.1.1.0 [170/20524800] via 10.1.1.5, 00:01:06, FastEthernet0/0

D EX     24.1.1.0 [170/20517120] via 10.1.1.5, 00:31:27, FastEthernet0/0
D EX     28.1.1.0 [170/20519680] via 10.1.1.5, 00:01:06, FastEthernet0/0
D EX     45.1.1.0 [170/20514560] via 10.1.1.5, 02:47:00, FastEthernet0/0
D EX     46.1.1.0 [170/21026560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     78.1.1.0 [170/20522240] via 10.1.1.5, 00:01:06, FastEthernet0/0
D EX     100.1.0.0/21 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.0.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.1.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.2.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.3.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.4.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.5.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.6.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     100.1.7.0/24 [170/21154560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     102.1.34.0/23 [170/21666560] via 10.1.1.5, 02:46:59, FastEthernet0/0
D EX     102.1.34.0/24 [170/21162240] via 10.1.1.5, 00:00:52, FastEthernet0/0
D EX     102.1.35.0/24 [170/21162240] via 10.1.1.5, 00:00:52, FastEthernet0/0
D EX  130.3.0.0/16 [170/20655360] via 10.1.1.5, 00:01:06, FastEthernet0/0
D EX  200.1.1.0/24 [170/20527360] via 10.1.1.5, 00:01:06, FastEthernet0/0
```

Task 16

Configure the routers in AS 500 such that they log neighbor warning messages and repeat the warning messages every 10 minutes. You should disable logging of neighbor changes for this AS.

On R1 and R3:

```
Rx(config)#Router eigrp 500
Rx(config-router)#Eigrp log-neighbor-warnings 600
Rx(config-router)#No Eigrp log-neighbor-changes
```

To verify the configuration:

On R1:

```
R1#Sh run | s router eigrp 500
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
router eigrp 500
 redistribute connected route-map TASK10
 redistribute eigrp 600
 network 10.1.100.1 0.0.0.0

 no eigrp log-neighbor-changes
 eigrp log-neighbor-warnings 600
```

Task 17

Configure R7 and R8 in AS 600 such that they NEVER exchange existing and/or future external Eigrp routes, DO NOT configure an Access-list, Prefix-list or a route-map to accomplish this task.

On R7:

```
R7#Show ip route 8.0.0.0

Routing entry for 8.0.0.0/8
  Known via "eigrp 600", distance 170, metric 156160, type external
  Redistributing via eigrp 600
  Last update from 78.1.1.8 on FastEthernet0/1, 02:43:13 ago
  Routing Descriptor Blocks:
  * 78.1.1.8, from 78.1.1.8, 02:43:13 ago, via FastEthernet0/1
      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
```

Note, R7 can reach network 8.0.0.0/8 via R8 (78.1.1.8).

On R7:

```
R7(config)#Router eigrp 600
R7(config-router)#Eigrp router-id 11.11.11.11
```

To verify the configuration:

On R7:

```
R7#Show ip eigrp topology 7.0.0.0/8 | Inc 11.11.11.11
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
EIGRP-IPv4 Topology Entry for AS(600)/ID(11.11.11.11) for 7.0.0.0/8
        Originating router is 11.11.11.11
```

In the output of the above show command we can see that the topology table states that the originating router is 11.11.11.11, therefore, this router will discard any prefix that was originated by another router with a router-id of 11.11.11.11.

Let's configure R8 with the same router-id and verify the result:

On R8:

```
R8(config)#Router eigrp 600
R8(config-router)#Eigrp router-id 11.11.11.11
```

To verify the configuration:

On R7:

```
R7#Show ip route 8.0.0.0

Routing entry for 8.0.0.0/8
  Known via "eigrp 600", distance 170, metric 20645120, type external
  Redistributing via eigrp 600
  Last update from 17.1.1.1 on FastEthernet0/0, 00:01:16 ago
  Routing Descriptor Blocks:
  * 17.1.1.1, from 17.1.1.1, 00:01:16 ago, via FastEthernet0/0
      Route metric is 20645120, traffic share count is 1
      Total delay is 25200 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 3
```

Task 18

Configure an HDLC connection between R3 and R4. R3 should use its S1/4 with an IP address 34.1.1.3 /24, and R4 should use its S1/3 interface with an IP address of 34.1.1.4 /24 for this connection.
These two routers should run EIGRP AS 300 on this link.

On R3:

© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

```
R3(config)#Int S1/4
R3(config-if)#Clock rate 64000
R3(config-if)#IP addr 34.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Router eigrp 300
R3(config-router)#Netw 34.1.1.3 0.0.0.0
```

On R4:

```
R4(config)#Int S1/3
R4(config-if)#Ip addr 34.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config-fr-dlci)#Router eigrp 300
R4(config-router)#Netw 34.1.1.4 0.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route 34.1.1.0

Routing entry for 34.1.1.0/24
  Known via "eigrp 300", distance 90, metric 21024000, type internal
  Redistributing via eigrp 300, eigrp 600
  Advertised by eigrp 600
  Last update from 14.1.1.4 on Serial1/4, 00:01:07 ago
  Routing Descriptor Blocks:
    14.1.1.4, from 14.1.1.4, 00:01:07 ago, via Serial1/4
      Route metric is 21024000, traffic share count is 1
      Total delay is 40000 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
  * 13.1.1.3, from 13.1.1.3, 00:01:07 ago, via Serial1/3

      Route metric is 21024000, traffic share count is 1
      Total delay is 40000 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Task 19

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R1 should be configured to load balance in a ratio of 1:3 for network 34.1.1.0 /24.

When we see load balancing and ratios, we immediately think of offset-list. Remember that offset-list only affects the advertised distance and if we add to the advertised distance the prefix may NOT even meet the Feasibility Condition. Therefore, the delay value should be manipulated to achieve this ratio, the following steps outlines the formula and steps to accomplish this task:

Let's check the routing table of R1 for this prefix:

On R1:

```
R1#Show ip route 34.1.1.0 | Inc from|Route metric

  Last update from 14.1.1.4 on Serial1/4, 00:01:32 ago
    14.1.1.4, from 14.1.1.4, 00:01:32 ago, via Serial1/4
      Route metric is 21024000, traffic share count is 1
  * 13.1.1.3, from 13.1.1.3, 00:01:32 ago, via Serial1/3
      Route metric is 21024000, traffic share count is 1
```

We can see that R1 is performing an equal cost load balancing using a ratio of 1:1. If we need to change the ratio to 1:3, we should take the composite metric and multiply it by 3. It means that if the composite metric of R4 is three times as much as the composite metric of R3 then, R1 can do an unequal-cost load balancing using a ratio of 1:3:

21,024,000 * 3 = 63,072,000

We know that by default EIGRP uses the following formula:

Composite value = 256 (BW + DLY)

Therefore, 63,072,000 should equal to 256 multiply by bandwidth plus the dealy value.

63,072,000 = 256(BW + DLY)

How is bandwidth calculated? How is the delay value calculated?
When it comes to bandwidth and delay values, EIGRP looks at the vector metric, because the vector metric gives us the slowest bandwidth along the path to a given destination, and for the delay, it calculates the sum of all delays along the path for a given destination. So let's look at the composite metric values.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Show ip eigrp topology 34.1.1.0 255.255.255.0

EIGRP-IPv4 Topology Entry for AS(300)/ID(130.100.1.1) for 34.1.1.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is

21024000
  Descriptor Blocks:
  13.1.1.3 (Serial1/3), from 13.1.1.3, Send flag is 0x0
      Composite metric is (21024000/20512000), route is Internal
      Vector metric:
        Minimum bandwidth is 128 Kbit
        Total delay is 40000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 130.3.3.3
  14.1.1.4 (Serial1/4), from 14.1.1.4, Send flag is 0x0
      Composite metric is (21024000/20512000), route is Internal
      Vector metric:
        Minimum bandwidth is 128 Kbit
        Total delay is 40000 microseconds
        Reliability is 255/255
        Load is 3/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 44.4.4.4
(The rest of the output is omitted for brevity)
```

We can see the minimum bandwidth along the path is 128 Kbit, EIGRP uses 10,000,000 / minimum bandwidth, let's figure this out:

10,000,000 / 128 = 78,125

Let's calculate the delay value. EIGRP uses the sum of all the delays along the path divided by ten, but the vector metric makes this process VERY easy; we can see that the topology table gives us the "total delay", this is the sum of all delays along the path divided by ten. Let's calculate the delay value:

40,000 / 10 = 4000

So we have two choices, we can manipulate the bandwidth or the delay value to get the desired composite metric of 63,072,000. But manipulating the bandwidth is not the best choice because this can affect QoS, or other routing protocols like OSPF, so the best choice is to manipulate delay, since not many protocols use this value, therefore, the goal here is to find out the delay value, let's put this in a formula:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

In order to get a composite metric of 63,072,000, we need to:

63,072,000 = 256(BW + DLY +X)

Let's go through this calculation:

63,072,000 = 256(78,125 + 4000 + X)

Where "X" is the delay value that we need to add to the delay value of the S1/4 interface of R1 facing R4:

63,072,000 = 256(82,125 + X)

Now, if 256 is moved to the left side of the equation, we should see the following:

63,072,000 / 256 = 82,125 + X

This gives us the following result:

246,375 = 82,125 + X

Now, all we need to do is move the 82,125 to the left side of the equation. Once this is done we should get the following:

246,375 - 82,125 = X

Now it should be easy to figure out the value of "X":

164,250 = X

Now that we know the value of "X" all we need to do is add this to the delay value of the S1/4 interface, let's find out the delay value of S1/4:

```
R1#Show interface S1/4 | Inc DLY

  MTU 1500 bytes, BW 128 Kbit/sec, DLY 20000 usec,
```

164,250 + 20,000 = 184,250

Let's add this value as the delay value of the S1/4 interface:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config)#int S1/4
R1(config-if)#Delay 184250

R1#Clear ip route *
R1#Clear ip eigrp neighbors
```

Let's see the result:

```
R1#Show ip eigrp topology 34.1.1.0 255.255.255.0

EIGRP-IPv4 Topology Entry for AS(300)/ID(130.100.1.1) for 34.1.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is
21024000
  Descriptor Blocks:
  13.1.1.3 (Serial1/3), from 13.1.1.3, Send flag is 0x0
      Composite metric is (21024000/20512000), route is Internal
      Vector metric:
        Minimum bandwidth is 128 Kbit
        Total delay is 40000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 130.3.3.3
  14.1.1.4 (Serial1/4), from 14.1.1.4, Send flag is 0x0
      Composite metric is (67680000/20512000), route is Internal
      Vector metric:
        Minimum bandwidth is 128 Kbit
        Total delay is 1862500 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
        Originating router is 44.4.4.4
 (The rest of the output is omitted for brevity)
```

But remember that EIGRP will NOT have the second best route in its routing table unless the "Variance" command in router configuration mode is configured. To calculate the variance, divid the lower value by the higher value:

67680000 / 21024000 = 3.219 → This value is rounded up to a variance of 4.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#router eigrp 300
R1(config-router)#Variance 4
```

To enforce the changes, let's clear the neighbor adjacency and the routing table:

On R1:

```
R1#Clear ip eigrp neighbors

R1#Clear ip route *
```

To verify the configuration:

On R1:

```
R1#Show ip route 34.1.1.0 | Inc from|metric

  Known via "eigrp 300", distance 90, metric 21024000, type internal
  Last update from 14.1.1.4 on Serial1/4, 00:01:47 ago
    14.1.1.4, from 14.1.1.4, 00:01:47 ago, via Serial1/4
      Route metric is 67680000, traffic share count is 5

  * 13.1.1.3, from 13.1.1.3, 00:01:47 ago, via Serial1/3
      Route metric is 21024000, traffic share count is 16
```

If we do the following calculation:

16 / 5 = 3.2

Remember that we need to round up.

Task 20

Configure the routers in AS 500 such that a dead neighbor is detected within 3 seconds.

On R1 and R3:

```
Rx(config)#int F0/1
Rx(config-if)#ip hello-interval eigrp 500 1
Rx(config-if)#ip hold-time eigrp 500 3
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 21

Routers in AS 400 should be configured to use "Bandwidth" and NOT "Bandwidth + DLY" when calculating their composite metric.

On R2:

```
R2(config)#router eigrp 400
R2(config-router)#Metric weights 0 1 0 0 0 0
```

You should see the following message:

```
IP-EIGRP(0) 400: Neighbor 24.1.1.4 (FastEthernet0/1) is down: metric
changed

IP-EIGRP(0) 400: Neighbor 24.1.1.4 (FastEthernet0/1) is down: K-value
mismatch
```

On R4:

```
R4(config)#router eigrp 400
R4(config-router)#Metric weights 0 1 0 0 0 0
```

You should see the following console message:

```
IP-EIGRP(0) 400: Neighbor 24.1.1.2 (FastEthernet0/1) is up: new adjacency
```

To verify the configuration:

On R4:

```
R4#Show ip protocol | Inc Metric weight|Routing Protocol

Routing Protocol is "eigrp 200"
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0

Routing Protocol is "eigrp 300"
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0

Routing Protocol is "eigrp 400"
    Metric weight K1=1, K2=0, K3=0, K4=0, K5=0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Note: the K value for AS 400 is changed.

Task 22

Routers in AS 500 should be configured to use Delay when calculating their composite metric.

The ToS value is always set to zero, and the next five values are K1 – K5 respectively. Each "K" value has a range of 0 – 255.

On R1 and R3:




```
Rx(config)#Router eigrp 500
Rx(config-router)#Metric weight 0 0 0 1 0 0
```

To verify the configuration:

On R1:

```
R1#Show ip protocol | Inc Metric weight|Routing Protocol

Routing Protocol is "eigrp 300"
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0

Routing Protocol is "eigrp 500"
    Metric weight K1=0, K2=0, K3=1, K4=0, K5=0

Routing Protocol is "eigrp 600"
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
```

Task 23

R9 should be configured to filter existing and future Eigrp routes that have a composite metric ranging 21,154,560 to 21,666,560 inclusively.

To accomplish this task you must deny the routes that have a composite metric of 21,154,560 to 21,666,560 using a route-map; to deny the routes in the specified range, you must use the following formula to calculate the metric and the deviation value:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The two numbers are added and then, divided by 2:

21,154,560 + 21,666,560 = 42,821,120
42,821,120 / 2 = 21,410,560 → This gives us the metric value

To calculate the deviation value:
Subtract the start of the range number (The lower value or the *"from"* value, in this case 21,154,560) from the *"to"* value, in this case 21,666,560 and then, divide the result by two:

21,154,560 - 21,666,560 = 512,000
512,000 / 2 = 256,000

Therefore, the "match metric" command will have the following values:
Match metric 21,410,560 +- 256,000

The above statement instructs the router to filter prefix that has a metric value of (21,410,560 – 256,000 = 21,154,560) to any prefix that has a metric value of (21,410,560 + 256,000 = 21,666,560) inclusively which means that the prefix/es with a metric value of 21,154,560 and/or 21,666,560 are also included.

On R9:

```
R9#Show ip route eigrp | B Gate
Gateway of last resort is not set

D EX  1.0.0.0/8 [170/21154560] via 10.1.1.5, 02:53:01, FastEthernet0/0
D EX  2.0.0.0/8 [170/21666560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX  3.0.0.0/8 [170/21154560] via 10.1.1.5, 04:43:13, FastEthernet0/0
D EX  4.0.0.0/8 [170/20642560] via 10.1.1.5, 04:43:13, FastEthernet0/0
D     5.0.0.0/8 [90/156160] via 10.1.1.5, 12:00:40, FastEthernet0/0
D     6.0.0.0/8 [90/158720] via 10.1.1.5, 12:00:40, FastEthernet0/0
D EX  7.0.0.0/8 [170/21157120] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX  8.0.0.0/8 [170/20647680] via 10.1.1.5, 02:36:25, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     10.1.4.0/24 [170/20517120] via 10.1.1.5, 04:43:13, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets

D EX     12.1.1.0 [170/21538560] via 10.1.1.5, 02:53:01, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D EX     13.1.1.0 [170/21538560] via 10.1.1.5, 04:43:13, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D EX     14.1.1.0 [170/21026560] via 10.1.1.5, 04:26:24, FastEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D EX     17.1.1.0 [170/21029120] via 10.1.1.5, 02:53:00, FastEthernet0/0
      24.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D EX      24.1.1.0 [170/20517120] via 10.1.1.5, 04:43:13, FastEthernet0/0
      28.0.0.0/24 is subnetted, 1 subnets
D EX      28.1.1.0 [170/20519680] via 10.1.1.5, 02:36:25, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets

D EX      34.1.1.0 [170/21026560] via 10.1.1.5, 04:43:13, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
D EX      45.1.1.0 [170/20514560] via 10.1.1.5, 04:43:16, FastEthernet0/0
      46.0.0.0/24 is subnetted, 1 subnets
D EX      46.1.1.0 [170/21026560] via 10.1.1.5, 04:43:13, FastEthernet0/0
      78.0.0.0/24 is subnetted, 1 subnets
D EX      78.1.1.0 [170/20522240] via 10.1.1.5, 02:36:25, FastEthernet0/0
      100.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D EX      100.1.0.0/21 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.0.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.1.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.2.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.3.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.4.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.5.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.6.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      100.1.7.0/24 [170/21154560] via 10.1.1.5, 02:53:00, FastEthernet0/0
      102.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX      102.1.34.0/23 [170/21666560] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      102.1.34.0/24 [170/21162240] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX      102.1.35.0/24 [170/21162240] via 10.1.1.5, 02:53:00, FastEthernet0/0
D EX  130.3.0.0/16 [170/20655360] via 10.1.1.5, 00:04:51, FastEthernet0/0
      199.1.1.0/32 is subnetted, 2 subnets
D        199.1.1.5 [90/156160] via 10.1.1.5, 12:00:40, FastEthernet0/0
D        199.1.1.6 [90/158720] via 10.1.1.5, 12:00:40, FastEthernet0/0
D EX  200.1.1.0/24 [170/20527360] via 10.1.1.5, 02:36:25, FastEthernet0/0
```

The routes that are highlighted are the ones that R9 should deny:

On R9:

```
R9(config)#Route-map TST deny 10
R9(config-route-map)#match metric 21410560 +- 256000
R9(config)#Route-map TST permit 90

R9(config)#Router eigrp 100
R9(config-router)#Distribute-list route-map TST in F0/0
```

To verify the configuration:

On R9:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R9#Show ip route eigrp | B Gate
Gateway of last resort is not set

D EX  4.0.0.0/8 [170/20642560] via 10.1.1.5, 05:07:51, FastEthernet0/0

D     5.0.0.0/8 [90/156160] via 10.1.1.5, 12:25:18, FastEthernet0/0
D     6.0.0.0/8 [90/158720] via 10.1.1.5, 12:25:18, FastEthernet0/0
D EX  8.0.0.0/8 [170/20647680] via 10.1.1.5, 03:01:03, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX   10.1.4.0/24 [170/20517120] via 10.1.1.5, 05:07:51, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
D EX     14.1.1.0 [170/21026560] via 10.1.1.5, 04:51:02, FastEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D EX     17.1.1.0 [170/21029120] via 10.1.1.5, 03:17:38, FastEthernet0/0
      24.0.0.0/24 is subnetted, 1 subnets
D EX     24.1.1.0 [170/20517120] via 10.1.1.5, 05:07:51, FastEthernet0/0
      28.0.0.0/24 is subnetted, 1 subnets
D EX     28.1.1.0 [170/20519680] via 10.1.1.5, 03:01:03, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
D EX     34.1.1.0 [170/21026560] via 10.1.1.5, 05:07:51, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/20514560] via 10.1.1.5, 05:07:54, FastEthernet0/0
      46.0.0.0/24 is subnetted, 1 subnets
D EX     46.1.1.0 [170/21026560] via 10.1.1.5, 05:07:51, FastEthernet0/0
      78.0.0.0/24 is subnetted, 1 subnets
D EX     78.1.1.0 [170/20522240] via 10.1.1.5, 03:01:03, FastEthernet0/0
D EX  130.3.0.0/16 [170/20655360] via 10.1.1.5, 00:29:29, FastEthernet0/0
      199.1.1.0/32 is subnetted, 2 subnets
D        199.1.1.5 [90/156160] via 10.1.1.5, 12:25:18, FastEthernet0/0
D        199.1.1.6 [90/158720] via 10.1.1.5, 12:25:18, FastEthernet0/0
D EX  200.1.1.0/24 [170/20527360] via 10.1.1.5, 03:01:03, FastEthernet0/0
```

Task 24

Configure R8 such that it NEVER has external Eigrp routes in its routing table. DO NOT use any access-list, route-map or prefix-list to accomplish this task.

To see the existing routes in R8's routing table:

On R8:

```
R8#Show ip route eigrp | B Gate
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Gateway of last resort is not set

D EX  1.0.0.0/8 [170/158720] via 78.1.1.7, 03:34:55, FastEthernet0/1
D EX  2.0.0.0/8 [170/156160] via 28.1.1.2, 05:53:19, FastEthernet0/0

D EX  3.0.0.0/8 [170/20645120] via 78.1.1.7, 03:34:55, FastEthernet0/1
D EX  4.0.0.0/8 [170/67301120] via 78.1.1.7, 03:34:55, FastEthernet0/1
D EX  5.0.0.0/8 [170/22178560] via 28.1.1.2, 03:01:27, FastEthernet0/0
D EX  6.0.0.0/8 [170/22181120] via 28.1.1.2, 03:01:27, FastEthernet0/0
D EX  7.0.0.0/8 [170/20645120] via 28.1.1.2, 03:34:55, FastEthernet0/0
D EX  9.0.0.0/8 [170/22181120] via 28.1.1.2, 03:01:27, FastEthernet0/0
      10.0.0.0/24 is subnetted, 2 subnets
D EX     10.1.1.0 [170/22053120] via 28.1.1.2, 03:01:27, FastEthernet0/0
D EX     10.1.4.0 [170/33280] via 28.1.1.2, 03:18:20, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D EX     12.1.1.0 [170/20514560] via 28.1.1.2, 03:34:56, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D EX     13.1.1.0 [170/20517120] via 78.1.1.7, 03:34:55, FastEthernet0/1
      14.0.0.0/24 is subnetted, 1 subnets
D EX     14.1.1.0 [170/67173120] via 78.1.1.7, 03:34:55, FastEthernet0/1
      17.0.0.0/24 is subnetted, 1 subnets
D EX     17.1.1.0 [170/20517120] via 28.1.1.2, 03:34:55, FastEthernet0/0
      24.0.0.0/24 is subnetted, 1 subnets
D EX     24.1.1.0 [170/30720] via 28.1.1.2, 05:53:19, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
D EX     34.1.1.0 [170/21538560] via 28.1.1.2, 03:34:55, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/67685120] via 78.1.1.7, 03:34:56, FastEthernet0/1
      46.0.0.0/24 is subnetted, 1 subnets
D EX     46.1.1.0 [170/67685120] via 78.1.1.7, 03:34:56, FastEthernet0/1
      100.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D EX     100.1.0.0/21 [170/158720] via 78.1.1.7, 03:34:55, FastEthernet0/1
D EX     100.1.0.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.1.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.2.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.3.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.4.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.5.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.6.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
D EX     100.1.7.0/24 [170/158720] via 78.1.1.7, 03:34:56, FastEthernet0/1
      102.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D EX     102.1.34.0/23 [170/156160] via 28.1.1.2, 05:53:19, FastEthernet0/0
D EX     102.1.34.0/24 [170/156160] via 28.1.1.2, 05:53:19, FastEthernet0/0
D EX     102.1.35.0/24 [170/156160] via 28.1.1.2, 05:53:19, FastEthernet0/0
D EX  130.3.0.0/16 [170/161280] via 78.1.1.7, 00:46:46, FastEthernet0/1
      199.1.1.0/32 is subnetted, 2 subnets
D EX     199.1.1.5 [170/22178560] via 28.1.1.2, 03:01:27, FastEthernet0/0
D EX     199.1.1.6 [170/22181120] via 28.1.1.2, 03:01:27, FastEthernet0/0
D EX  200.1.1.0/24 [170/33280] via 78.1.1.7, 03:34:56, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To configure the task:

On R8:

```
R8(config)#Router eigrp 600
R8(config-router)#Distance Eigrp 90 255
```

The above command assigns an administrative distance of 90 to all internal EIGRP routes and the second value assigns an administrative distance of 255 to the existing and future external routes, which means that the external routes will NOT be injected into the routing table.

To verify the configuration:

On R8:

```
R8#Show ip route eigrp | B Gate
Gateway of last resort is not set
R8#
```

Task 25

Configure R2 such that Eigrp NEVER uses more than 25% of its S1/1 link's bandwidth.

On R2:

```
R2(config)#Int S1/1
R2(config-subif)#IP bandwidth-percent eigrp 300 25
```

Task 26

Configure Eigrp AS 700 on SW1, SW2, SW3 and SW4 and advertise the following interfaces in this AS:

On SW1 – F0/19, F0/21, and F0/23
On SW2 – F0/19, and F0/23
On SW3 – F0/19, F0/21, and F0/23
On SW4 – F0/19 and F0/23

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

By default, IP Routing is disabled on switches and the following command is neded for the switches to

perform IP routing. In IOS release 15, the autosummarization is disabled by default. But on the following switches is NOT disabled by default and you must disable the auto-summarization if the task requires it.

```
SW1(config)#ip routing

SW1(config)#Router eigrp 700
SW1(config-router)#Netw 112.1.1.11 0.0.0.0
SW1(config-router)#Netw 113.1.1.11 0.0.0.0
SW1(config-router)#Netw 114.1.1.11 0.0.0.0
```

On SW2:

```
SW2(config)#IP routing

SW2(config)#Router eigrp 700
SW2(config-router)#Netw 123.1.1.22 0.0.0.0
SW2(config-router)#Netw 112.1.1.22 0.0.0.0
```

On SW3:

```
SW3(config)#IP routing

SW3(config)#Router eigrp 700
SW3(config-router)#Netw 113.1.1.33 0.0.0.0
SW3(config-router)#Netw 123.1.1.33 0.0.0.0
SW3(config-router)#Netw 134.1.1.33 0.0.0.0
```

On SW4:

```
SW4(config)#IP routing

SW4(config)#Router eigrp 700
SW4(config-router)#Netw 114.1.1.44 0.0.0.0
SW4(config-router)#Netw 134.1.1.44 0.0.0.0
```

To verify the configuration:

On SW2:

```
SW2#Show ip route eigrp
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
     114.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D      114.1.1.0/24 [90/30720] via 112.1.1.11, 00:00:04, FastEthernet0/19
D      114.0.0.0/8 [90/61440] via 123.1.1.33, 00:00:01, FastEthernet0/23
     113.0.0.0/24 is subnetted, 1 subnets

D      113.1.1.0 [90/30720] via 123.1.1.33, 00:00:04, FastEthernet0/23
                 [90/30720] via 112.1.1.11, 00:00:04, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D      134.1.1.0 [90/30720] via 123.1.1.33, 00:00:04, FastEthernet0/23
```

On SW3:

```
SW3#Show ip route eigrp

     114.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       114.1.1.0/24 [90/30720] via 134.1.1.44, 00:00:03, FastEthernet0/19
                     [90/30720] via 113.1.1.11, 00:00:03, FastEthernet0/21
D       114.0.0.0/8 [90/51200] via 134.1.1.44, 00:00:03, FastEthernet0/19
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 123.1.1.22, 00:00:03, FastEthernet0/23
                  [90/30720] via 113.1.1.11, 00:00:03, FastEthernet0/21
```

On SW4:

```
SW4#Show ip route eigrp

     114.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       114.0.0.0/8 [90/56320] via 114.1.1.11, 00:00:03, FastEthernet0/23
     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 134.1.1.33, 00:00:01, FastEthernet0/19
                  [90/30720] via 114.1.1.11, 00:00:01, FastEthernet0/23
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 114.1.1.11, 00:00:01, FastEthernet0/23
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.33, 00:00:01, FastEthernet0/19
```

On SW1:

```
SW1#Show ip route eigrp

     114.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D       114.0.0.0/8 [90/53760] via 113.1.1.33, 00:00:03, FastEthernet0/21
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 113.1.1.33, 00:00:01, FastEthernet0/21
                  [90/30720] via 112.1.1.22, 00:00:01, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D       134.1.1.0 [90/30720] via 114.1.1.44, 00:00:03, FastEthernet0/23
                  [90/30720] via 113.1.1.33, 00:00:03, FastEthernet0/21
```

Task 27

Configure SW1 to advertise its Lo12, Lo13 and Lo14 as candidate default routes.

This task calls for the "IP default-network" global config command; this command can ONLY be used when "ip routing" is enabled. When this command is configured, the router considers routes to the configured network for installation as the gateway of last resort; the receiving router will flag the route/s as a candidate default route ONLY if it has reachability to the specified network/s.

Step One:
In order to provide reachability to the rest of the switches, the three networks MUST be advertised in Eigrp:

On SW1:

```
SW1(config)#Router eigrp 700
SW1(config-router)#Netw 200.1.12.11 0.0.0.0
SW1(config-router)#Netw 200.1.13.11 0.0.0.0
SW1(config-router)#Netw 200.1.14.11 0.0.0.0
```

Step Two:
Configure "IP Default-Network" for the three networks:

On SW1:

```
SW1(config)#IP default-network 200.1.12.0
SW1(config)#IP default-network 200.1.13.0
SW1(config)#IP default-network 200.1.14.0
```

To verify the configuration:

On SW1:

```
SW1#Show ip route | Inc 200

C*    200.1.12.0/24 is directly connected, Loopback12
C*    200.1.13.0/24 is directly connected, Loopback13
C*    200.1.14.0/24 is directly connected, Loopback14
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW2:

```
SW2#Sh ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 112.1.1.11, 01:15:48, FastEthernet0/19
     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 123.1.1.33, 01:15:48, FastEthernet0/23
                  [90/30720] via 112.1.1.11, 01:15:48, FastEthernet0/19
D*   200.1.12.0/24 [90/156160] via 112.1.1.11, 00:00:55, FastEthernet0/19
D*   200.1.13.0/24 [90/156160] via 112.1.1.11, 00:00:51, FastEthernet0/19
D*   200.1.14.0/24 [90/156160] via 112.1.1.11, 00:00:47, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.33, 01:15:48, FastEthernet0/23
```

On SW3:

```
SW3#Sh ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 134.1.1.44, 01:16:32, FastEthernet0/19
                  [90/30720] via 113.1.1.11, 01:16:32, FastEthernet0/21
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 123.1.1.22, 01:16:32, FastEthernet0/23
                  [90/30720] via 113.1.1.11, 01:16:32, FastEthernet0/21
D*   200.1.12.0/24 [90/156160] via 113.1.1.11, 00:01:41, FastEthernet0/21
D*   200.1.13.0/24 [90/156160] via 113.1.1.11, 00:01:37, FastEthernet0/21
D*   200.1.14.0/24 [90/156160] via 113.1.1.11, 00:01:34, FastEthernet0/21
```

On SW4:

```
SW4#Show ip route eigrp

     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 134.1.1.33, 01:17:26, FastEthernet0/19
                  [90/30720] via 114.1.1.11, 01:17:26, FastEthernet0/23
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 114.1.1.11, 01:17:26, FastEthernet0/23
D*   200.1.12.0/24 [90/156160] via 114.1.1.11, 00:02:35, FastEthernet0/23
D*   200.1.13.0/24 [90/156160] via 114.1.1.11, 00:02:31, FastEthernet0/23
D*   200.1.14.0/24 [90/156160] via 114.1.1.11, 00:02:27, FastEthernet0/23
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.33, 01:17:27, FastEthernet0/19
```


© 2014 Narbik Kocharians, All rights reserved

https://t.me/learningnets

Task 28

Configure an HDLC connection between R2 and R3. R2 should use its S1/3 interface and assign an IP address of 23.1.1.2 /24 and R3 should use its S1/2 interface and assign an IP address of 23.1.1.3 /24. These routers should advertise this IP address in Eigrp AS 300.

On R2:

```
R2(config)#Int S1/3
R2(config-if)#Clock rate 64000
R2(config-if)#IP addr 23.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#router eigrp 300
R2(config-router)#Netw 23.1.1.2 0.0.0.0
```

On R3:

```
R3(config)#Int S1/2
R3(config-if)#Ip addr 23.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config-subif)#Router eigrp 300
R3(config-router)#Netw 23.1.1.3 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip eigrp neighbors

EIGRP-IPv4 Neighbors for AS(300)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   23.1.1.3                Se1/3            14 00:00:51   82   1140  0  12276
0   12.1.1.1                Se1/1            10 05:15:04   59   1140  0  20241

EIGRP-IPv4 Neighbors for AS(400)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
0   24.1.1.4                Fa0/1           11 04:58:27    1   200  0  4768

EIGRP-IPv4 Neighbors for AS(600)
H   Address                 Interface       Hold Uptime    SRTT   RTO  Q  Seq
                                            (sec)          (ms)       Cnt Num

0   28.1.1.8                Fa0/0           10 01:38:40    1   200  0  16331
```

Task 29

Configure SW2, SW3 and SW4 based on the following policy:

1. SW2 should have 200.1.12.0 /24, 200.1.13.0 /24 and 200.1.14.0 /24 in its routing table, BUT it should ONLY use 200.1.12.0 /24 as the candidate default route. SW2 should be configured to accomplish this task.

2. SW3 should have 200.1.12.0 /24, 200.1.13.0 /24 and 200.1.14.0 /24 in its routing table, BUT it should ONLY use 200.1.13.0 /24 as the candidate default route.. SW3 should be configured to accomplish this task.

3. SW4 should have 200.1.12.0 /24, 200.1.13.0 /24 and 200.1.14.0 /24 in its routing table, BUT it should ONLY use 200.1.14.0 /24 as the candidate default route. SW4 should be configured to accomplish this task.

To configure the first policy:

On SW2:

Note SW2 has three candidate default routes in its routing table.

```
SW2#Show ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 112.1.1.11, 01:20:21, FastEthernet0/19
     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 123.1.1.33, 01:20:21, FastEthernet0/23
                  [90/30720] via 112.1.1.11, 01:20:21, FastEthernet0/19
D*   200.1.12.0/24 [90/156160] via 112.1.1.11, 00:05:28, FastEthernet0/19
D*   200.1.13.0/24 [90/156160] via 112.1.1.11, 00:05:24, FastEthernet0/19
D*   200.1.14.0/24 [90/156160] via 112.1.1.11, 00:05:20, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.33, 01:20:21, FastEthernet0/23
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

In order to accomplish this task, you must configure an access-list to identify the network that must be used as a candidate default route:

```
SW2(config)#Access-list 12 permit 200.1.12.0 0.0.0.255
```

Once the access-list is configured, you must configure the "Default-information allowed in" command to reference the access-list. Basically this command will ONLY allow the specified Network to be used as a candidate default route.

```
SW2(config)#Router eigrp 700
SW2(config-router)#default-information allowed in 12

SW2#Clear ip eigrp neighbors

SW2#Show ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 112.1.1.11, 00:00:08, FastEthernet0/19
     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 123.1.1.33, 00:00:08, FastEthernet0/23
                  [90/30720] via 112.1.1.11, 00:00:08, FastEthernet0/19
D*   200.1.12.0/24 [90/156160] via 112.1.1.11, 00:00:08, FastEthernet0/19
D    200.1.13.0/24 [90/156160] via 112.1.1.11, 00:00:08, FastEthernet0/19
D    200.1.14.0/24 [90/156160] via 112.1.1.11, 00:00:08, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.33, 00:00:08, FastEthernet0/23
```

We can see that SW2 is using 200.1.12.0/24 as the candidate default.

To configure the second policy:

On SW3:

```
SW3#Show ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 134.1.1.44, 00:01:25, FastEthernet0/19
                  [90/30720] via 113.1.1.11, 00:01:25, FastEthernet0/21
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 123.1.1.22, 00:01:29, FastEthernet0/23
                  [90/30720] via 113.1.1.11, 00:01:29, FastEthernet0/21
D*   200.1.12.0/24 [90/156160] via 113.1.1.11, 00:01:25, FastEthernet0/21
D*   200.1.13.0/24 [90/156160] via 113.1.1.11, 00:01:25, FastEthernet0/21
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D*    200.1.14.0/24 [90/156160] via 113.1.1.11, 00:01:26, FastEthernet0/21
```

```
SW3(config)#Access-list 13 permit 200.1.13.0 0.0.0.255
```

Since the "allowed" keyword is the default, it can be omitted:

```
SW3(config)#router eigrp 700
SW3(config-router)#Default-information in 13

SW3#Clear ip eigrp neighbors

SW3#Show ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 134.1.1.44, 00:00:03, FastEthernet0/19
                  [90/30720] via 113.1.1.11, 00:00:03, FastEthernet0/21
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 123.1.1.22, 00:00:03, FastEthernet0/23
                  [90/30720] via 113.1.1.11, 00:00:03, FastEthernet0/21
D    200.1.12.0/24 [90/156160] via 113.1.1.11, 00:00:03, FastEthernet0/21
D*   200.1.13.0/24 [90/156160] via 113.1.1.11, 00:00:03, FastEthernet0/21
D    200.1.14.0/24 [90/156160] via 113.1.1.11, 00:00:04, FastEthernet0/21
```

To configure the third policy:

On SW4:

```
SW4#Sh ip route eigrp

     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 134.1.1.33, 00:01:37, FastEthernet0/19
                  [90/30720] via 114.1.1.11, 00:01:37, FastEthernet0/23
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 114.1.1.11, 00:01:37, FastEthernet0/23
D*   200.1.12.0/24 [90/156160] via 114.1.1.11, 00:01:37, FastEthernet0/23
D*   200.1.13.0/24 [90/156160] via 114.1.1.11, 00:01:37, FastEthernet0/23
D*   200.1.14.0/24 [90/156160] via 114.1.1.11, 00:01:37, FastEthernet0/23
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.33, 00:01:38, FastEthernet0/19

SW4(config)#Access-list 14 permit 200.1.14.0 0.0.0.255

SW4(config)#router eigrp 700
SW4(config-router)#default-information in 14

SW4#Clear ip eigrp neighbors
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On SW4:

```
SW4#Show ip route eigrp

     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 134.1.1.33, 00:00:09, FastEthernet0/19
                  [90/30720] via 114.1.1.11, 00:00:09, FastEthernet0/23
     112.0.0.0/24 is subnetted, 1 subnets
D       112.1.1.0 [90/30720] via 114.1.1.11, 00:00:09, FastEthernet0/23
D    200.1.12.0/24 [90/156160] via 114.1.1.11, 00:00:09, FastEthernet0/23
D    200.1.13.0/24 [90/156160] via 114.1.1.11, 00:00:09, FastEthernet0/23
D*   200.1.14.0/24 [90/156160] via 114.1.1.11, 00:00:09, FastEthernet0/23
     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 134.1.1.33, 00:00:10, FastEthernet0/19
```

Task 30

Configure the link between R3 and SW1 in AS 800 and redistribute the routes from AS 700 into this AS.

On SW1:

```
SW1(config)#Router eigrp 800
SW1(config-router)#Netw 31.1.1.11 0.0.0.0
SW1(config-router)#Redistribute eigrp 700
```

On R3:

```
R3(config)#Router eigrp 800
R3(config-router)#Netw 31.1.1.3 0.0.0.0
```

To verify the configuration:

On R3:

```
R3#Show ip rou eigrp 800 | B Gate
Gateway of last resort is not set
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      112.0.0.0/24 is subnetted, 1 subnets
D EX     112.1.1.0 [170/30720] via 31.1.1.11, 00:00:09, FastEthernet0/0
      113.0.0.0/24 is subnetted, 1 subnets
D EX     113.1.1.0 [170/30720] via 31.1.1.11, 00:00:09, FastEthernet0/0

      114.0.0.0/24 is subnetted, 1 subnets
D EX     114.1.1.0 [170/30720] via 31.1.1.11, 00:00:09, FastEthernet0/0
      123.0.0.0/24 is subnetted, 1 subnets
D EX     123.1.1.0 [170/33280] via 31.1.1.11, 00:00:09, FastEthernet0/0
      134.1.0.0/24 is subnetted, 1 subnets
D EX     134.1.1.0 [170/33280] via 31.1.1.11, 00:00:09, FastEthernet0/0
D*EX  200.1.12.0/24 [170/156160] via 31.1.1.11, 00:00:09, FastEthernet0/0
D*EX  200.1.13.0/24 [170/156160] via 31.1.1.11, 00:00:09, FastEthernet0/0
D*EX  200.1.14.0/24 [170/156160] via 31.1.1.11, 00:00:09, FastEthernet0/0
```

Task 31

Configure injection of default routes on R3 based on the following policy:

1. For R3's neighbor/s in AS 300:
   R3 should inject a default route for every neighbor in this AS. You should use an interface configuration mode to accomplish this task.

2. For R3's neighbor/s in AS 500:
   R3 should be configured to advertise a default route plus networks 200.1.12.0/24, 200.1.13.0 /24 and 200.1.14.0 /24 from AS 800.

3. For R3's neighbor/s in AS 800:
   R3 should inject a default route into this AS.

4. R1 should always prefer the default route through AS 300 and NOT any other AS.

To configure the first item:

On R3:

R3 has three neighbors in AS 300 and they are reached via S1/2, S1/4 and S1/1 interfaces, therefore, the "IP summary-address" interface configuration mode command can be configured on these three interfaces. To see the interfaces in AS 300:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip eigrp 300 interfaces

EIGRP-IPv4 Interfaces for AS(300)
                        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
Se1/1              1        0/0       57       5/190        418          0
Se1/4              1        0/0       64       5/190        454          0
Se1/2              1        0/0      165       5/190        958          0
```

On R3:

```
R3(config)#int S1/1
R3(config-if)#ip summary-address eigrp 300 0.0.0.0 0.0.0.0

R3(config)#int S1/2
R3(config-if)#ip summary-address eigrp 300 0.0.0.0 0.0.0.0

R3(config)#Int S1/4
R3(config-if)#ip summary-address eigrp 300 0.0.0.0 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp 300 | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/20640000] via 23.1.1.3, 00:00:53, Serial1/3
```

On R4:

```
R4#Show ip route eigrp 300 | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/20640000] via 34.1.1.3, 00:01:07, Serial1/3
```

On R1:

```
R1#Show ip route eigrp 300 | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/20640000] via 13.1.1.3, 00:01:46, Serial1/3
```

To configure the second item:

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

By default, if an "IP summary-address" is configured, then the more specific routes are suppressed, unless a "leak-map" is configured.

Since the routes are ONLY advertised in AS 700, they should be redistributed into AS 500, as follows:

On R3:

```
R3(config)#Access-list 2 permit 200.1.12.0 0.0.0.255
R3(config)#Access-list 2 permit 200.1.13.0 0.0.0.255
R3(config)#Access-list 2 permit 200.1.14.0 0.0.0.255

R3(config)#route-map intoAS500
R3(config-route-map)#match ip addr 2

R3(config)#router eigrp 500
R3(config-router)#redistribute Eigrp 800 route-map intoAS500
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp 500 | B Gate
Gateway of last resort is 13.1.1.3 to network 0.0.0.0

D EX    130.3.0.0/16 [170/130560] via 200.1.1.3, 00:00:25, FastEthernet0/1
D*EX  200.1.12.0/24 [170/133120] via 200.1.1.3, 00:00:25, FastEthernet0/1
D*EX  200.1.13.0/24 [170/133120] via 200.1.1.3, 00:00:25, FastEthernet0/1
D*EX  200.1.14.0/24 [170/133120] via 200.1.1.3, 00:00:25, FastEthernet0/1
```

The highlighted routes are the ones coming from AS 800. Network 130.3.0.0 /16 was redistributed into this AS in Task 10. Therefore, another access-list statement is configured to include this network, as follows:

On R3:

```
R3(config)#access-list 2 permit 130.3.0.0 0.0.255.255

R3(config)#Route-map TST permit 10
R3(config-route-map)#match ip addr 2

R3(config-route-map)#int F0/1
R3(config-if)#ip summary-address eigrp 500 0.0.0.0 0.0.0.0 leak-map TST
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip route eigrp 500 | B Gate
Gateway of last resort is 200.1.1.3 to network 0.0.0.0

D*    0.0.0.0/0 [90/5120] via 200.1.1.3, 00:01:04, FastEthernet0/1
D EX  130.3.0.0/16 [170/130560] via 200.1.1.3, 01:32:38, FastEthernet0/1
D*EX  200.1.12.0/24 [170/133120] via 200.1.1.3, 01:32:33, FastEthernet0/1
D*EX  200.1.13.0/24 [170/133120] via 200.1.1.3, 01:32:33, FastEthernet0/1
D*EX  200.1.14.0/24 [170/133120] via 200.1.1.3, 01:32:33, FastEthernet0/1
```

To configure the third item:

On R3:

```
R3(config)#int f0/0
R3(config-if)#ip summary-address eigrp 800 0.0.0.0 0.0.0.0
```

To verify the configuration:

On SW1:

```
SW1#Show ip route eigrp 800

D*    0.0.0.0/0 [90/30720] via 10.1.113.3, 00:04:43, FastEthernet0/3
```

To configure the fourth item:

Let's check the routing table of R1:

On R1:

```
R1#Show ip route eigrp | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/5120] via 200.1.1.3, 00:05:28, FastEthernet0/1
```

R1 is getting the default route through its F0/1 interface which is in AS 500. This task's requirement states that R1 should always prefer the default route through AS 300 and NOT any other AS.

Is R1 receiving the default route through its S1/3 interface? If R1 was receiving a default route through its S1/3 interface, the default route would be in the Topology table, let's verify that information:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1#Show ip eigrp 300 topology | S 0.0.0.0/0

P 0.0.0.0/0, 0 successors, FD is Inaccessible
        via 12.1.1.2 (21152000/20640000), Serial1/2
        via 14.1.1.4 (67808000/20640000), Serial1/4

        via 13.1.1.3 (20640000/128256), Serial1/3
```

The output of the above show command reveals that the default route was advertised to the local router by R3; the default route was injected with a cost of 20,640,000. If the routing table of R1 is checked again, you will see that R1's cost through AS 500 (F0/1) is 5,120, and that's the reason R1 chooses the default route through AS 500 and NOT AS 300.

Administrative distance can be manipulated to influence R1's decision, let's configure and verify:

```
R1(config)#router eigrp 500
R1(config-router)#distance 91 10.1.100.3 0.0.0.0

R1#Clear ip eigrp neighbors
```

To verify the configuration:

On R1:

```
R1#Sh ip route eigrp | Inc 0.0.0.0/0

D*    0.0.0.0/0 [90/20640000] via 13.1.1.3, 00:00:26, Serial1/3
```

Task 32

R4 in AS 400 should ONLY receive routes from R2 on 24.1.1.0 /24 segment, and it should NOT advertise any route/s to any existing and future router in this AS.

On R2:

```
R2(config)#Router eigrp 400
R2(config-router)#Neighbor 24.1.1.4 F0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R4:

```
R4(config)#Router eigrp 400
R4(config-router)#Neighbor 24.1.1.2 F0/1
R4(config-router)#eigrp stub receive-only
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp 400 | B Gate
Gateway of last resort is 23.1.1.3 to network 0.0.0.0
```

On R4:

```
R4#Show ip route eigrp 400 | B Gate
Gateway of last resort is 34.1.1.3 to network 0.0.0.0

D EX  8.0.0.0/8 [170/25600] via 24.1.1.2, 00:01:39, FastEthernet0/1
      28.0.0.0/24 is subnetted, 1 subnets
D EX     28.1.1.0 [170/25600] via 24.1.1.2, 00:01:39, FastEthernet0/1
      78.0.0.0/24 is subnetted, 1 subnets
D EX     78.1.1.0 [170/25600] via 24.1.1.2, 00:01:39, FastEthernet0/1
```

Task 33

SW2 should be configured to reach network 200.1.13.0 /24 through SW3. SW2 should ONLY use SW1 if SW3 is down.

On SW2:

```
SW2#Show ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 112.1.1.11, 02:14:07, FastEthernet0/19
     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 123.1.1.33, 02:14:07, FastEthernet0/23
                  [90/30720] via 112.1.1.11, 02:14:07, FastEthernet0/19
D    200.1.12.0/24 [90/156160] via 112.1.1.11, 02:14:07, FastEthernet0/19
D    200.1.13.0/24 [90/156160] via 112.1.1.11, 02:14:07, FastEthernet0/19
D    200.1.14.0/24 [90/156160] via 112.1.1.11, 02:14:07, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D       134.1.1.0 [90/30720] via 123.1.1.33, 02:14:08, FastEthernet0/23
```

The output of the above show command displays that SW2 goes directly to SW1 to reach network 200.1.13.0 /24.
By looking at the diagram we can see that there are two paths to reach this network, let's verify this information by looking at the topology table:

```
SW2#Show ip eigrp topology 200.1.13.0 255.255.255.0

IP-EIGRP (AS 700): Topology entry for 200.1.13.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 156160
  Routing Descriptor Blocks:
  112.1.1.11 (FastEthernet0/19), from 112.1.1.11, Send flag is 0x0
      Composite metric is (156160/128256), Route is Internal
      Vector metric:
        Minimum bandwidth is 100000 Kbit
        Total delay is 5100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1504
        Hop count is 1
  123.1.1.33 (FastEthernet0/23), from 123.1.1.33, 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 1504
        Hop count is 2
```

Sure enough there are two paths to reach network 200.1.13.0/24, one through SW1 directly and the second path is through SW3 and the reason SW2 is going directly to SW1 to reach network 200.1.13.0/24 is because it has a lower cost. Let's resolve this by manipulating the administrative distance:

On SW2:

```
SW2(config)#Access-list 13 permit 200.1.13.0 0.0.0.255

SW2(config)#Router eigrp 700
SW2(config-router)#Distance 91 112.1.1.11 0.0.0.0 13
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2#Show ip route eigrp

     114.0.0.0/24 is subnetted, 1 subnets
D       114.1.1.0 [90/30720] via 112.1.1.11, 00:00:10, FastEthernet0/19
     113.0.0.0/24 is subnetted, 1 subnets
D       113.1.1.0 [90/30720] via 123.1.1.33, 00:00:10, FastEthernet0/23

                  [90/30720] via 112.1.1.11, 00:00:10, FastEthernet0/19
D    200.1.12.0/24 [90/156160] via 112.1.1.11, 00:00:10, FastEthernet0/19
D    200.1.13.0/24 [90/158720] via 123.1.1.33, 00:00:10, FastEthernet0/23
D    200.1.14.0/24 [90/156160] via 112.1.1.11, 00:00:10, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 123.1.1.33, 00:00:11, FastEthernet0/23
```

To test the configuration:

On SW2:

```
SW2#Traceroute 200.1.13.11

Type escape sequence to abort.
Tracing the route to 200.1.13.11
  1 123.1.1.33 0 msec 0 msec 0 msec
  2 113.1.1.11 0 msec *  0 msec
```

Task 34

Configure R9 to filter network 17.1.1.0 /24. You should use the existing "distribute-list route-map" to accomplish this task.

Let's check the routing table of R9 for network 17.1.1.0 /24:

On R9:

```
R9#Show ip route 17.1.1.0

Routing entry for 17.1.1.0/24
  Known via "eigrp 100", distance 170, metric 21029120, type external
  Redistributing via eigrp 100
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
  Last update from 10.1.1.5 on FastEthernet0/0, 01:04:20 ago
  Routing Descriptor Blocks:
  * 10.1.1.5, from 10.1.1.5, 01:04:20 ago, via FastEthernet0/0
      Route metric is 21029120, traffic share count is 1
      Total delay is 40200 microseconds, minimum bandwidth is 128 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 3
```

Since the existing route-map MUST be used to accomplish this task, our ONLY choice here is to use the "Continue" clause.

A route-map evaluates match clauses until a successful match occurs. After a match occurs, the route-map STOPS evaluating match clauses; when a continue clause is configured, the route-map continues to evaluate and execute match clauses in the specified route-map entry after a successful match occurs.

```
R9(config)#Access-list 17 permit 17.1.1.0 0.0.0.255

R9(config)#Route-map TST deny 10
R9(config-route-map)#continue 20

R9(config)#Route-map TST deny 20
R9(config-route-map)#Match ip addr 17
```

To see the complete route-map:

On R9:

```
R9#Show run | s route-map

 distribute-list route-map TST in FastEthernet0/0

route-map TST deny 10
 match metric 21410560 +- 256000
 continue 20
route-map TST deny 20
 match ip address 17
route-map TST permit 90
```

We can see how the additional commands were added to the existing route-map.

To verify the configuration:

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9#Show ip route 17.1.1.0
% Network not in table
```

Task 35

Configure R1 to filter Network 2.0.0.0 /8. DO NOT use distribute-list to accomplish this task.

On R1:

```
R1#Show ip route 2.0.0.0

Routing entry for 2.0.0.0/8
  Known via "eigrp 300", distance 90, metric 20640000, type internal
  Redistributing via eigrp 300, eigrp 600
  Advertised by eigrp 600
  Last update from 12.1.1.2 on Serial1/2, 01:13:53 ago
  Routing Descriptor Blocks:
  * 12.1.1.2, from 12.1.1.2, 01:13:53 ago, via Serial1/2
      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
```

The administrative distance can be used to configure this task. A route with an administrative distance of 255 will NOT get injected into the routing table. In EIGRP, the administrative distance can ONLY be used in the following manner for internal prefixes and NOT external.

On R1:

```
R1(config)#Access-list 3 permit 2.0.0.0 0.255.255.255

R1(config)#Router eigrp 300
R1(config-router)#Distance 255 12.1.1.2 0.0.0.0 3
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1#Show ip route 2.0.0.0
% Network not in table
```

Task 36

R7 should be configured to filter any route/s that have a hop count of 4 or more. DO NOT use an access-list or a prefix-list to accomplish this task.

Let's check some of the prefixes that are the advertised on the other end of the diagram, like networks advertised by R9, R6 and SW1:

On R7:

```
R7#Show ip eigrp topology 200.1.12.0 | Inc Hop
        Hop count is 3

R7#Show ip eigrp topology 5.0.0.0 | Inc Hop
        Hop count is 3

R7#Show ip eigrp topology 9.0.0.0 | Inc Hop
        Hop count is 4

R7#Show ip eigrp topology 6.0.0.0 | Inc Hop
        Hop count is 4
```

To configure this task we can manipulate the "Maximum-hops" command:

```
R7(config)#router eigrp 600
R7(config-router)#metric maximum-hops 3
```

To verify the configuration:

On R7:

```
R7#Show ip eigrp topology 6.0.0.0 | Inc Hop
R7#
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7#Show ip eigrp topology 9.0.0.0 | Inc Hop
R7#

R7#Show ip eigrp topology 5.0.0.0 | Inc Hop
        Hop count is 3

R7#Show ip eigrp topology 200.1.12.0 | Inc Hop
        Hop count is 3
```

Task 37

Configure the appropriate routers/switches such that if SW1 loses its Successor and has no feasible successor it will NOT send a query to SW2, SW3 or SW4.

SW1 will NOT send queries to neighbors that are configured as stub. This can be verified by running "debug eigrp packets" or "Debug eigrp packets terse" and then, shutting down a loopback on SW1 and observing the output of the debug. The difference between the "Debug eigrp packet" and "Debug eigrp packet Terse" is that the later one will NOT show the hello packets in the output of the debug.

Let's see the default behavior:

On SW1:

```
SW1#Debug eigrp packets terse
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)

SW1(config)#Int Lo12
SW1(config-if)#Shut
```

We should see the follwiong console messages stating that the local router has enqueued and then sent

queries out of its F0/3, F0/19, F0/21, and F0/23 interfaces:

```
EIGRP: Enqueueing QUERY on FastEthernet0/21 iidbQ un/rely 0/1 serno 16-16
EIGRP: Enqueueing QUERY on FastEthernet0/23 iidbQ un/rely 0/1 serno 16-16
EIGRP: Enqueueing QUERY on FastEthernet0/19 iidbQ un/rely 0/1 serno 16-16
EIGRP: Enqueueing QUERY on FastEthernet0/3 iidbQ un/rely 0/1 serno 16-16
EIGRP: Sending QUERY on FastEthernet0/21
EIGRP: Sending QUERY on FastEthernet0/23
EIGRP: Sending QUERY on FastEthernet0/19
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
EIGRP: Sending QUERY on FastEthernet0/3
```

Let's "No Shut" the loopback interface and configure SW2, SW3 and SW4 as "Eigrp Stub" and test again:

```
SW1(config)#Int Lo12
SW1(config-if)#No Shut
```

On SW2, SW3, and SW4:

```
SWx(config)#router eigrp 700

SWx(config-router)#Eigrp stub connected
```

To verify and test the configuration:

On SW1:

```
SW1#Show ip route eigrp 700

     123.0.0.0/24 is subnetted, 1 subnets
D       123.1.1.0 [90/30720] via 113.1.1.33, 00:00:25, FastEthernet0/21
                  [90/30720] via 112.1.1.22, 00:00:25, FastEthernet0/19
     134.1.0.0/24 is subnetted, 1 subnets
D       134.1.1.0 [90/30720] via 114.1.1.44, 00:00:13, FastEthernet0/23
                  [90/30720] via 113.1.1.33, 00:00:13, FastEthernet0/21

SW1#Debug eigrp packets terse
EIGRP Packets debugging is on

    (UPDATE, REQUEST, QUERY, REPLY, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)

SW1(config)#Int Lo12
SW1(config-if)#Shut
```

You should see the following console messages stating that the SW1 has enqueued and sent updates out of its F0/19, F0/21 and F0/23 interfaces, but the local Switch (SW1) has enqueued and sent a query out of its F0/3 interface to R3, this happened because R3 was NOT configured as a stub router.

```
EIGRP: Enqueueing UPDATE on FastEthernet0/19 iidbQ un/rely 0/1 serno 31-31
EIGRP: Enqueueing UPDATE on FastEthernet0/21 iidbQ un/rely 0/1 serno 31-31
EIGRP: Enqueueing UPDATE on FastEthernet0/23 iidbQ un/rely 0/1 serno 31-31

EIGRP: Enqueueing QUERY on FastEthernet0/3 iidbQ un/rely 0/1 serno 20-20
```


© 2014 Narbik Kocharians, All rights reserved

https://t.me/learningnets

```
EIGRP: Sending UPDATE on FastEthernet0/19
EIGRP: Sending UPDATE on FastEthernet0/21
EIGRP: Sending UPDATE on FastEthernet0/23

EIGRP: Sending QUERY on FastEthernet0/3
```

Task 38

Routers in AS 400 should wait one hour after sending queries before declaring a non responding neighbor as stuck in active, whereas, routers in AS 500 should wait indefinitely.

On R2 and R4:

```
Rx(config)#Router eigrp 400
Rx(config-router)#Timers active-time 60
```

This command instructs the routers in AS 400 to wait for 60 minutes for query replies before considering the route as stuck in active.

On R1 and R3:

```
Rx(config)#Router eigrp 500
Rx(config-router)#Timers active-time disable
```

This command instructs the routers in AS 500 to wait indefinitely for query replies.

Task 39

Configure Loopback 999 on R9 using the following policy:
IP address: 200.200.200.3 /24
Tag: 999

On R9:

```
R9(config)#Int lo999
R9(config-if)#IP addr 200.200.200.9 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R9(config)#Route-map Task39 permit 10
R9(config-route-map)#Match inter lo999
R9(config-route-map)#Set tag 999
R9(config)#Route-map Task39 permit 90

R9(config-route-map)#Router eigrp 100
R9(config-router)#Redistribute connected route-map Task39 metric 1 1 1 1 1
```

To verify the configuration:

On R9:

```
R9#Show ip eigrp topology 200.200.200.0 255.255.255.0 | Inc Admin

        Administrator tag is 999 (0x000003E7)
```

On R1:

```
R1#Show ip eigrp topology 200.200.200.0 255.255.255.0 | Inc tag

        Administrator tag is 999 (0x000003E7)
        Administrator tag is 999 (0x000003E7)
```

Task 40

Configure Eigrp AS 900 on the F0/0 interface of R5, R6, and R9; configure the following Loopback interfaces and advertise them in this AS:

R5 – Lo209 – 209.1.1.5/32
R6 – Lo209 – 209.1.1.6/32
R9 – Lo209 – 209.1.1.9/32

On R5:

```
R5(config-router)#int lo209
R5(config-if)#ip addr 209.1.1.5 255.255.255.255

R5(config)#Router eigrp 900
R5(config-router)#Netw 10.1.1.5 0.0.0.0
R5(config-router)#Netw 209.1.1.5 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6(config-router)#int lo209
R6(config-if)#ip addr 209.1.1.6 255.255.255.255

R6(config)#Router eigrp 900
R6(config-router)#Netw 209.1.1.6 0.0.0.0

R6(config-router)#Netw 10.1.1.6 0.0.0.0
```

On R9:

```
R9(config-router)#int lo209
R9(config-if)#ip addr 209.1.1.9 255.255.255.255

R9(config)#Router eigrp 900
R9(config-router)#Netw 10.1.1.9 0.0.0.0
R9(config-router)#Netw 209.1.1.9 0.0.0.0
```

To verify the configuration:

On R9:

```
R9#Show ip route eigrp 900 | B Gate
Gateway of last resort is not set

      209.1.1.0/32 is subnetted, 3 subnets
D        209.1.1.5 [90/156160] via 10.1.1.5, 00:04:37, FastEthernet0/0
D        209.1.1.6 [90/156160] via 10.1.1.6, 00:05:06, FastEthernet0/0
```

Task 41

R7 should be configured NOT to have any external Eigrp Routes in its routing table.

On R7:

```
R7#Show ip route eigrp | B Gate
Gateway of last resort is 17.1.1.1 to network 0.0.0.0

D*EX  0.0.0.0/0 [170/20642560] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX  1.0.0.0/8 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX  4.0.0.0/8 [170/67298560] via 17.1.1.1, 02:39:20, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
D EX  5.0.0.0/8 [170/67810560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      10.0.0.0/24 is subnetted, 2 subnets
D EX     10.1.1.0 [170/67685120] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     10.1.4.0 [170/67173120] via 17.1.1.1, 02:39:20, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D EX     12.1.1.0 [170/20514560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
D EX     13.1.1.0 [170/20514560] via 17.1.1.1, 02:39:20, FastEthernet0/0

      14.0.0.0/24 is subnetted, 1 subnets
D EX     14.1.1.0 [170/67170560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
D EX     23.1.1.0 [170/21026560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      24.0.0.0/24 is subnetted, 1 subnets
D EX     24.1.1.0 [170/67173120] via 17.1.1.1, 02:39:20, FastEthernet0/0
      28.0.0.0/24 is subnetted, 1 subnets
D EX     28.1.1.0 [170/67175680] via 17.1.1.1, 02:39:20, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
D EX     34.1.1.0 [170/67682560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [170/67682560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      46.0.0.0/24 is subnetted, 1 subnets
D EX     46.1.1.0 [170/67682560] via 17.1.1.1, 02:39:20, FastEthernet0/0
      100.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
D EX     100.1.0.0/21 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.0.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.1.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.2.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.3.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.4.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.5.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.6.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX     100.1.7.0/24 [170/156160] via 17.1.1.1, 02:39:20, FastEthernet0/0
      102.0.0.0/23 is subnetted, 1 subnets
D EX     102.1.34.0 [170/20642560] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX  130.3.0.0/16 [170/158720] via 17.1.1.1, 02:39:20, FastEthernet0/0
      199.1.1.0/32 is subnetted, 1 subnets
D EX     199.1.1.5 [170/67810560] via 17.1.1.1, 02:39:20, FastEthernet0/0
D EX  200.1.1.0/24 [170/30720] via 17.1.1.1, 02:39:20, FastEthernet0/0
D*EX  200.1.12.0/24 [170/161280] via 17.1.1.1, 02:25:16, FastEthernet0/0
D*EX  200.1.13.0/24 [170/161280] via 17.1.1.1, 02:39:20, FastEthernet0/0
D*EX  200.1.14.0/24 [170/161280] via 17.1.1.1, 02:39:20, FastEthernet0/0
```

On R7:

Once again the administrative distance can be used to resolve this task. When it comes to configure administrative distance in EIGRP, there are two flavors:

Distance *AD-value Neighbor's IP address inverse-mask ACL*

The above configuration ONLY affects the internal routes received from a given neighbor, however, we can

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

manipulate the Neighbor's IP address and the inverse mask such that all neighbors form a given directly connected segment are affected.

Distance EIGRP *AD-value for internal routes AD for external routes*

With the above command we can affect we can change the administrative distance for all internal or external, since a route-map or an ACL can not be referenced, this command can change the administrative distance for ALL internal and/or ALL external routes.

```
R7(config)#Router eigrp 600
R7(config-router)#Distance Eigrp 255 170
```

To verify the configuration:

On R7:

```
R7#Show ip route eigrp | B Gate
Gateway of last resort is not set
R7#

R7#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(600)
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
1   17.1.1.1                Fa0/0             14 00:00:28    4   300  0  62
0   78.1.1.8                Fa0/1             13 00:00:28    4   300  0  37
```

Task 42

Configure R9 in AS 900 such that it filters prefix 209.1.1.5 /32 advertised by R5. DO NOT use a "route-map" or a "distance" command to accomplish this task.

On R9:

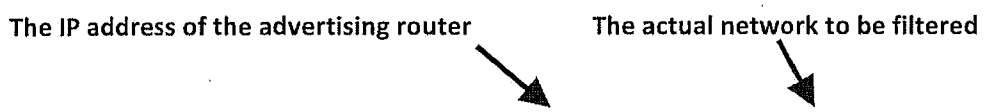
```
R9#Show ip route eigrp 900 | B Gate
Gateway of last resort is not set

      209.1.1.0/32 is subnetted, 3 subnets
D        209.1.1.5 [90/156160] via 10.1.1.5, 00:35:36, FastEthernet0/0
D        209.1.1.6 [90/156160] via 10.1.1.6, 00:36:05, FastEthernet0/0
```

Since R9 ONLY has a single interface through which it receives prefixes 209.1.1.6/32 and 209.1.1.5 /32 from R5 and R6, and the task states that a distance command or a route-map can NOT be used to

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

accomplish this task; an extended access-list can be used to match prefix 209.1.1.5 /32 advertised by R5, extended access-list can be used in this manner in Eigrp, RIPv2 and/or BGP:




```
R9(config)#Access-list 100 deny ip host 10.1.1.5 host 209.1.1.5
R9(config)#Access-list 100 permit ip any any

R9(config)#Router eigrp 900
R9(config-router)#Distribute-list 100 in F0/0
```

To verify the configuration:

On R9:

```
R9#Show ip route eigrp 900 | B Gate
Gateway of last resort is not set

      209.1.1.0/32 is subnetted, 2 subnets
D        209.1.1.6 [90/156160] via 10.1.1.6, 00:40:50, FastEthernet0/0
```

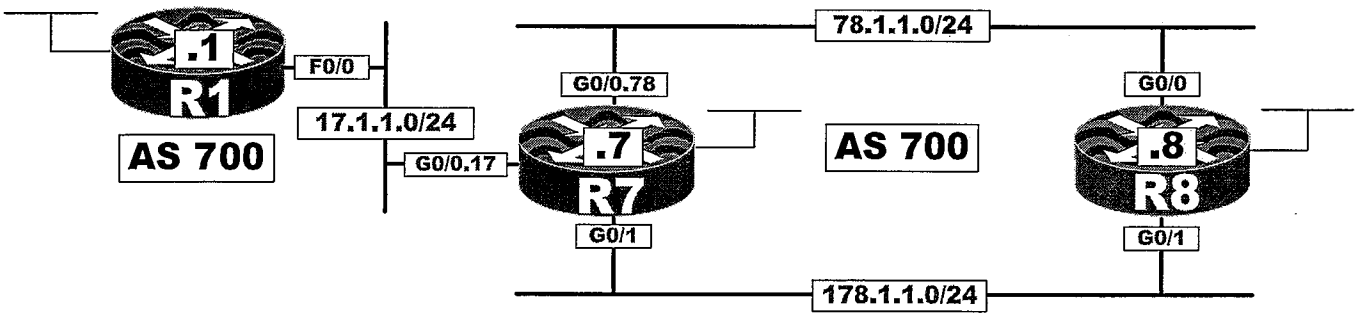
Task 43

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 6
Configuring EIGRP SHA-2 256 Authentication



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "EIGRP"→ "Lab-6".

Task 1

Configure EIGRP based on the above diagram. If this configuration is successful, these routers should be able to see and have reachability to each others routes. You should use a named EIGRP configuration when configuring R7 and R8, and Classic EIGRP configuration when configuring R1 to accomplish this task.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#Netw 1.1.1.1 0.0.0.0
R1(config-router)#Netw 17.1.1.1 0.0.0.0
```

On R7:

```
R7(config)#router eigrp TST
R7(config-router)#address-family ipv4 unicast autonomous-system 100
R7(config-router-af)#Netw 7.7.7.7 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R7(config-router-af)#Netw 178.1.1.7 0.0.0.0
R7(config-router-af)#Netw 78.1.1.7 0.0.0.0
R7(config-router-af)#Netw 17.1.1.7 0.0.0.0
```

On R8:

```
R8(config)#router eigrp TST
R8(config-router)#address-family ipv4 unicast autonomous-system 100
R8(config-router-af)#Netw 8.8.8.8 0.0.0.0
R8(config-router-af)#Netw 178.1.1.8 0.0.0.0
R8(config-router-af)#Netw 78.1.1.8 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 78.1.1.7 (GigabitEthernet0/0)
is up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 178.1.1.7 GigabitEthernet0/1)
is up: new adjacency
```

To verify the configuration:

On R8:

```
R8#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     1.0.0.0/8 [90/2713600] via 178.1.1.7, 00:01:16, GigabitEthernet0/1
                [90/2713600] via 78.1.1.7, 00:01:16, GigabitEthernet0/0
D     7.0.0.0/8 [90/103040] via 178.1.1.7, 00:01:16, GigabitEthernet0/1
                [90/103040] via 78.1.1.7, 00:01:16, GigabitEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D        17.1.1.0 [90/153600] via 178.1.1.7, 00:01:16, GigabitEthernet0/1
                  [90/153600] via 78.1.1.7, 00:01:16, GigabitEthernet0/0

R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     7.0.0.0/8 [90/28192] via 17.1.1.7, 00:09:11, FastEthernet0/0
D     8.0.0.0/8 [90/30752] via 17.1.1.7, 00:07:38, FastEthernet0/0
      78.0.0.0/24 is subnetted, 1 subnets
D        78.1.1.0 [90/30720] via 17.1.1.7, 00:09:11, FastEthernet0/0
      178.1.0.0/24 is subnetted, 1 subnets
D        178.1.1.0 [90/30720] via 17.1.1.7, 00:09:11, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 2

Configure R7 to authenticate all existing and future directly connected interface/s using the strongest authentication method available; use minimum number of commands and "CCIE" as the password to accomplish this task.
R7 should authenticate R1 using MD5 and "Cisco" as the password.
In the future R8 may have other unauthenticated interface/s connecting to other router.

To ensure authenticity of the received packets from an EIGRP neighbor, the EIGRP packets must be authenticated; this means that the two EIGRP neighbors must be configured with the same preshared authentication key.

In EIGRP, authentication is configured on a per interface basis which means that the packet exchange between the two neighbors through that given interface is authenticated.
When HMAC-SHA-256 authentication is configured, each EIGRP packet is authenticated using the preshared key. The preshared key is used to generate and verify a message digest that is added to the packet. The message digest is a one-way function of the packet and the secret key.

The sending router calculates the hash based on the two key parts and the data based on the following:

- **Key Part 1: The configured Shared secret**
- **Key Part 2: The interface's IP address through which the packet is sent**
- **Data: The EIGRP packet to be sent.**

The receiving router calculates the hash based on the following:

- **Key Part 1: The configured Shared secret**
- **Key Part 2: The IPv4/IPv6 source address in the packet header**
- **Data: The EIGRP packet received.**

If the hash value matches, the packet is accepted. Remember that this authentication is ONLY available through the named configuration and NOT the classic EIGRP configuration.

Let's configure the task:

On R7:

```
R7(config)#router eigrp TST
R7(config-router)#address-family ipv4 unicast autonomous-system 100
```

The following configuration applies the authentication to the "AF-interface default" which will be applied to all existing and future interfaces added to the local router. This configuration can only be achieved

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

through the name configuration; this option is not available in the classical configuration of EIGRP.

```
R7(config-router-af)#af-interface default
R7(config-router-af-interface)#authentication mode hmac-sha-256 CCIE
```

We can see that the MD5 authentication is applied to the G0/0.17 sub-interface, in this case the most specific configuration will be implemented, meaning that what we apply to the interface/sub-interface will override what the interface/sub-interface inherits through the "default" configuration.

```
R7(config-router-af-interface)#af-interface g0/0.17
R7(config-router-af-interface)#authentication mode md5
R7(config-router-af-interface)#authentication key-chain TST

R7(config)#Key chain TST
R7(config-keychain)#key 1
R7(config-keychain-key)#key-string Cisco
```

On R8:

On R8, the "af-interface default" configuration can not be used, because the task clearly states that in the future R8 may have other interface/s that will not need authentication.

```
R8(config)#router eigrp TST
R8(config-router)#address-family ipv4 unicast autonomous-system 100
R8(config-router-af)#af-interface g0/0
R8(config-router-af-interface)#authentication mode hmac-sha-256 CCIE
R8(config-router-af-interface)#af-interface g0/1
R8(config-router-af-interface)#authentication mode hmac-sha-256 CCIE
```

You should see the following console messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 78.1.1.7 (GigabitEthernet0/0)
is up: new adjacency

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 178.1.1.7(GigabitEthernet0/1)
is up: new adjacency
```

On R1:

```
R1(config)#key chain TST
R1(config-keychain)#key 1
R1(config-keychain-key)#key-string Cisco

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


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#ip authentication mode eigrp 100 md5
R1(config-if)#ip authentication key-chain eigrp 100 TST
```

You should see the following console messages:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 17.1.1.7 (FastEthernet0/0) is
up: new adjacency
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     7.0.0.0/8 [90/28192] via 17.1.1.7, 00:01:32, FastEthernet0/0
D     8.0.0.0/8 [90/30752] via 17.1.1.7, 00:01:32, FastEthernet0/0
      78.0.0.0/24 is subnetted, 1 subnets
D        78.1.1.0 [90/30720] via 17.1.1.7, 00:01:32, FastEthernet0/0
      178.1.0.0/24 is subnetted, 1 subnets
D        178.1.1.0 [90/30720] via 17.1.1.7, 00:01:32, FastEthernet0/0
```

On R8:

```
R8#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     1.0.0.0/8 [90/2713600] via 178.1.1.7, 00:02:01, GigabitEthernet0/1
                [90/2713600] via 78.1.1.7, 00:02:01, GigabitEthernet0/0
D     7.0.0.0/8 [90/103040] via 178.1.1.7, 00:04:11, GigabitEthernet0/1
                [90/103040] via 78.1.1.7, 00:04:11, GigabitEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D        17.1.1.0 [90/153600] via 178.1.1.7, 00:04:11, GigabitEthernet0/1
                  [90/153600] via 78.1.1.7, 00:04:11, GigabitEthernet0/0
```

Task 3

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

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Advanced
CCIE Routing & Switching
v5.0

www.MicronicsTraining.com

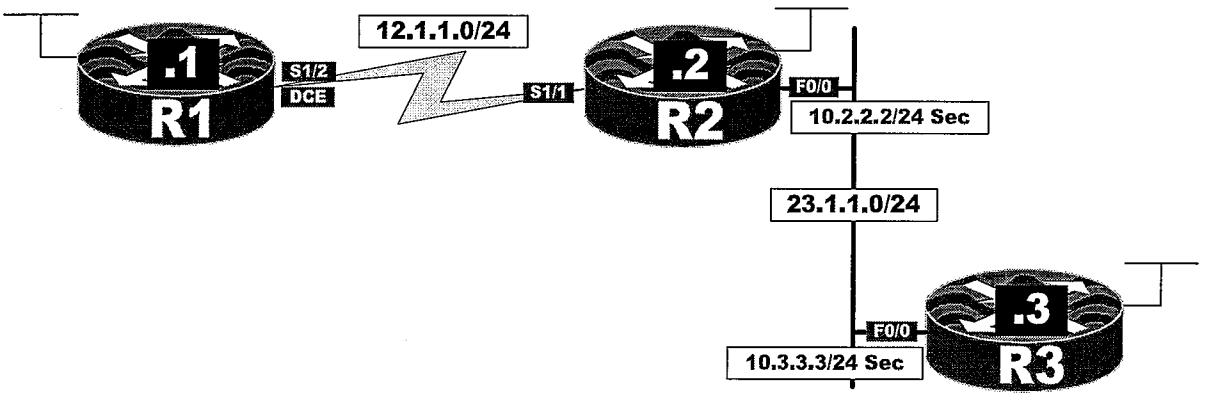
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

OSPF

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Running OSPF on the interfaces



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "OSPF"→ "Lab-1".

Task 1

Configure OSPF area 0 on the Lo0 interfaces and the links that connect the routers to each other, this includes the secondary ip addresses assigned to the F0/0 interfaces of R2 and R3. Do not use the "network" command to accomplish this task. The loopback 0 interfaces should be advertised with their correct mask. If this configuration is performed successfully, you should have reachability to all routes within the topology.

On R1:

```
R1(config)#int s1/2
R1(config-if)#ip ospf 1 area 0

R1(config)#int lo0
R1(config-if)#ip ospf 1 area 0
```

When configuring OSPF for an interface there are two choices:
Runnign OSPF directly under that interface – This is done by entering the "IP OSPF Area" command in the

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

interface configuration mode.

Configuring a Network command in the router configuration mode – This is achieved by configuring the "Network Area" command in the router configuration mode.

In this case OSPF is enabled under the Loopback 0 and S1/2 interfaces of R1.

To verify the configuration:

On R1:

```
R1#Show ip ospf interface brie

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/8          1     LOOP  0/0
Se1/2        1     0               12.1.1.1/24        781   P2P   0/0
```

When configuring OSPF directly under a given interface, an OSPF process is auto-started for every process ID specified in the command. Let's verify:

```
R1#Show run | S router ospf

router ospf 1
```

IF both (A "Network" command, and "IP OSPF area") is configured, the configuration of the interface supercedes the "Network" command configured in the router configuration mode. Let's verify:

```
R1(config)#Router ospf 1
R1(config-router)#Network 12.1.1.1 0.0.0.0 area 12
```

As we can see the following show command reveals that S1/2 interface is running in area 0.

```
R1#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/8          1     LOOP  0/0
Se1/2        1     0               12.1.1.1/24        781   P2P   0/0
```

Let's remove the "Network" command and continue:

```
R1(config)#router ospf 1
R1(config-router)#No Network 12.1.1.1 0.0.0.0 area 12
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's check the database for all the interfaces that R1 is running OSPF on:

```
R1#Show ip ospf database router self-originate

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Router Link States (Area 0)

  LS age: 283
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000004
  Checksum: 0x27CD
  Length: 48
  Number of Links: 2
```

The loopback interfaces are advertised as stub, and since the local router can not establish an OSPF adjacency through the loopback interface, it is advertised as the ONLY IP address on that interface and therefore the IP address is advertised as a host route (/32):

```
    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 1.1.1.1
     (Link Data) Network Mask: 255.255.255.255
      Number of MTID metrics: 0
       TOS 0 Metrics: 1
```

The serial interface of R1 is entered in the database as a stub network and the network IP address is advertised with its correct mask:

```
    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 12.1.1.0
     (Link Data) Network Mask: 255.255.255.0
      Number of MTID metrics: 0
       TOS 0 Metrics: 781
```

Let's see why the cost of this interface is 781. In OSPF cost is calculated based on the following formula:

```
Cost = Reference/Bandwidth
```

To find out the Bandwidth of this interface:

```
R1#Show interface S1/2 | Inc BW
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
  MTU 1500 bytes, BW 128 Kbit/sec, DLY 20000 usec,
```

Since the reference is not changed, in this case:

Cost = 100000/128 = 781.25

The fractions are dropped, therefore, the cost from OSPF's perspective is 781. To verify this:

```
R1#Show ip ospf interface s1/2 | Inc Cost

  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 781
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

But the task stated that the loopback interfaces should be advertised with their correct mask, to accomplish this task the network type of this interface is changed to "Point-To-Point" and since in Point-To-Point network types the interfaces can form an adjacency through that interface, the network is advertised with its correct mask. Let's change the network type and verify:

```
R1(config)#Int lo0
R1(config-if)#IP ospf network point-to-point
```

To verify the configuration:

```
R1#Show ip ospf database router

            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Router Link States (Area 0)

  LS age: 64
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 1.1.1.1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000006
  Checksum: 0xB20E
  Length: 60
  Number of Links: 3

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 1.0.0.0
     (Link Data) Network Mask: 255.0.0.0
      Number of MTID metrics: 0
       TOS 0 Metrics: 1
(The rest of the output is omitted for brevity)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#int lo0
R2(config-if)#ip ospf network point-to-point
R2(config-if)#ip ospf 1 area 0

R2(config)#int f0/0
R2(config-if)#ip ospf 1 area 0

R2(config)#int s1/1
R2(config-if)#ip ospf 1 area 0
```

On R3:

```
R3(config)#int lo0
R3(config-if)#ip ospf 1 area 0
R3(config-if)#ip ospf network point-to-point

R3(config)#int f0/0
R3(config-if)#ip ospf 1 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

O     2.0.0.0/8 [110/782] via 12.1.1.2, 00:02:30, Serial1/2
O     3.0.0.0/8 [110/783] via 12.1.1.2, 00:01:19, Serial1/2
      10.0.0.0/24 is subnetted, 2 subnets
O        10.2.2.0 [110/782] via 12.1.1.2, 00:02:30, Serial1/2
O        10.3.3.0 [110/783] via 12.1.1.2, 00:01:19, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/782] via 12.1.1.2, 00:01:29, Serial1/2
```

NOTE: The secondary IP addresses are also advertised.

```
R1#Show ip ospf route

            OSPF Router with ID (1.1.1.1) (Process ID 1)


                Base Topology (MTID 0)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
    Area BACKBONE(0)

    Intra-area Route List
*>  23.1.1.0/24, Intra, cost 782, area 0
      via 12.1.1.2, Serial1/2
*   1.0.0.0/8, Intra, cost 1, area 0, Connected
      via 1.1.1.1, Loopback0
*>  2.0.0.0/8, Intra, cost 782, area 0
      via 12.1.1.2, Serial1/2
*>  3.0.0.0/8, Intra, cost 783, area 0
      via 12.1.1.2, Serial1/2
*>  10.2.2.0/24, Intra, cost 782, area 0
      via 12.1.1.2, Serial1/2
*>  10.3.3.0/24, Intra, cost 783, area 0
      via 12.1.1.2, Serial1/2
*   12.1.1.0/24, Intra, cost 781, area 0, Connected
      via 12.1.1.1, Serial1/2
```

The above show command is a hidden one and it shows all the intra area routes, their cost, next-hop IP address and interface to reach that route and the area in which they are running in.

NOTE: the directly connected routes do not have the greater than (>) sign to the left of the network IP address.

Task 2

Configure R2 and R3 such that the secondary IP addresses are **not** advertised; **Do not** filtering, route-map, access-list, or Prefix-list, DO NOT remove them. You should use minimum number of commands to accomplish this task.

On R2 and R3:

The "secondaries none" keyword will not advertise the secondary IP address/es.

```
Rx(config)#int f0/0
Rx(config-if)#ip ospf 1 area 0 secondaries none
```

To verify the configuration:

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

O     2.0.0.0/8 [110/782] via 12.1.1.2, 00:06:31, Serial1/2
O     3.0.0.0/8 [110/783] via 12.1.1.2, 00:05:20, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/782] via 12.1.1.2, 00:05:30, Serial1/2

R1#Show ip ospf route | B Area
    Area BACKBONE(0)

    Intra-area Route List
*>  23.1.1.0/24, Intra, cost 782, area 0
      via 12.1.1.2, Serial1/2
*   1.0.0.0/8, Intra, cost 1, area 0, Connected
      via 1.1.1.1, Loopback0
*>  2.0.0.0/8, Intra, cost 782, area 0
      via 12.1.1.2, Serial1/2
*>  3.0.0.0/8, Intra, cost 783, area 0
      via 12.1.1.2, Serial1/2
*   12.1.1.0/24, Intra, cost 781, area 0, Connected
      via 12.1.1.1, Serial1/2
```

NOTE: The secondary prefixes are no longer advertised.

Task 3

R3 is getting flooded with LSA Type 6 packets; ensure that R3 does not generate a syslog message for these LSA types.

Cisco routers do not support LSA type 6 packets therefore each time an OSPF router receives an MOSPF (LSA type 6) packet it generates a syslog message. If the router/s receive many LSA type 6 packets they will generate a large number of syslog messages, this feature should be disabled to prevent this from occurring.

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#ignore lsa mospf
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

To ensure fast detection of a neighbor being down, configure R2 and R3 to send their hellos every 250ms with a hold time of one second for their Ethernet link.

The dead interval is advertised in OSPF hello packets. The values of this parameter must be the same for two routers in order for them to form a neighbor adjacency. By specifying the "minimal" and "hello-multiplier" keywords with a multiplier value, you are enabling OSPF fast hello packets, the "minimal" keyword sets the dead interval to 1 second and the "hello-multiplier" value sets the number of hello packets sent during the dead interval (1 second). Since one second is 1000ms, if 1000 is divided by 250 (the requested hello interval), the result should be the hello multiplier. Let's configure this feature and verify:

On R2 and R3:

```
Rx(config)#Int F0/0
Rx(config-if)#IP ospf dead-interval minimal hello-multiplier 4
```

To verify the configuration:

On R3:

```
R3#Show ip ospf interface F0/0 | Inc Hello

  Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 5
    Hello due in 18 msec
```

Task 5

Ensure that these routers lookup DNS names for use in most of the OSPF show commands, test this task to ensure proper operation. Since there are no DNS servers in this lab, you should use the routers for that purpose. Configure the router-ids of these routers to 1.1.1.1, 2.2.2.2 and 3.3.3.3 for R1, R2 and R3 respectively.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1

R1#Clear ip ospf process
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Reset ALL OSPF processes? [no]: Y
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 1.1.1.1

R2#Clear ip ospf process

Reset ALL OSPF processes? [no]: Y
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#router-id 1.1.1.1

R3#Clear ip ospf process

Reset ALL OSPF processes? [no]: Y
```

On R1:

```
R1#Show ip ospf database router | Inc Advert

  Advertising Router: 1.1.1.1
  Advertising Router: 2.2.2.2
  Advertising Router: 3.3.3.3
```

As we can see the router ID of the routers are used as the router identifier. Let's configure this task and verify:

On R1, R2 and R3:

```
Rx(config)#IP host R1 1.1.1.1
Rx(config)#IP host R2 2.2.2.2
Rx(config)#IP host R3 3.3.3.3

Rx(config)#IP ospf name-lookup
```

To verify the configuration:

On R1:

```
R1#Show ip ospf database router | inc Advert
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Advertising Router: R1
Advertising Router: R2
Advertising Router: R3
```

Task 6

Configure R2 such that if it does not receive an acknowledgment from R3 for a given LSA, it waits for 10 seconds before it resends that given LSA.

When an OSPF enabled router sends an LSA to its neighbor, it needs to receives an ACK from that given neighbor, if the retransmission timer expires and the router receives no ACKs, the router will resend that LSA. The default timer is set to 5 seconds, and the range is 1 – 65535.

Let's see the default value:

```
R2#Show ip ospf int f0/0 | Inc Timer

  Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 5
```

To change the value:

```
R2(config)#Int F0/0
R2(config-if)#IP ospf retransmit-interval 10
```

To verify the configuration:

```
R2#Show ip ospf int f0/0 | Inc Timer

 Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 10
```

Task 7

Configure R1 and R2 such that when there is a topology change in area 0 for LSAs type ones and twos, the entire SPF is NOT recomputed. This should **only** occur for the affected part/s of the tree.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

When any change occurs for an LSAs type-1 and/or Type-2, the entire SPF is recomputed. The incremental SPF allows the IOS to only recompute the affected part/s of the tree. This can result in faster OSPF convergence, and save resources.

On R1 and R2:

```
R2#Show ip ospf | Inc Increm

 Incremental-SPF disabled

R1(config)#router ospf 1
R1(config-router)#ispf
```

To verify the configuration:

On R1:

```
R1#Show ip ospf | Inc Increm

 Incremental-SPF enabled
```

Task 8

Configure R2 such that it limits the number of non-self generated LSAs that an OSPF routing process can keep in the OSPF LSDB to 10, a warning message should be generated when 80 percent of this threshold is reached, and an additional warning message should be generated for every prefix above the configured threshold.

Let's verify the existing number of non-self generated LSAs on R2:

On R2:

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

O     1.0.0.0/8 [110/782] via 12.1.1.1, 00:08:06, Serial1/1
O     3.0.0.0/8 [110/2] via 23.1.1.3, 00:08:06, FastEthernet0/0

R2#Show ip ospf database database-summary
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
                    OSPF Router with ID (2.2.2.2) (Process ID 1)

Area 0 database summary
  LSA Type       Count    Delete   Maxage
  Router         3        0        0
  Network        1        0        0
  Summary Net    0        0        0
  Summary ASBR   0        0        0
  Type-7 Ext     0        0        0
    Prefixes redistributed in Type-7  0
  Opaque Link    0        0        0
  Opaque Area    0        0        0
  Subtotal       4        0        0

Process 1 database summary
  LSA Type       Count    Delete   Maxage
  Router         3        0        0
  Network        1        0        0
  Summary Net    0        0        0
  Summary ASBR   0        0        0
  Type-7 Ext     0        0        0
  Opaque Link    0        0        0
  Opaque Area    0        0        0
  Type-5 Ext     0        0        0
      Prefixes redistributed in Type-5  0
  Opaque AS      0        0        0
  Non-self       2
  Total          4        0        0
```

As we can see that there are three router LSAs and one Network LSA, and the number of "Non-self" LSAs are toward the bottom of the output and it's currectly set to 2.

Let's configure this task and verify:

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#max-lsa 10 80 warning-only
```

To generate more LSAs; we will configure R1 to redistribute three of its loopback interfaces and then verify R2:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Route-map tst permit 10
R1(config-route-map)#match inter Lo1 Lo2 Lo3

R1(config)#Router ospf 1
R1(config-router)#Redistribute connected route-map tst sub
```

The above route-map can be modified to include more Loopback interfaces as we test. Let's verify the number of non-self generated LSAs On R2:

On R2:

```
R2#Sh ip route ospf | b Gate
Gateway of last resort is not set

O     1.0.0.0/8 [110/782] via 12.1.1.1, 00:49:16, Serial1/1
O     3.0.0.0/8 [110/2] via 23.1.1.3, 00:49:16, FastEthernet0/0
O E2  110.0.0.0/8 [110/20] via 12.1.1.1, 00:00:10, Serial1/1
O E2  120.0.0.0/8 [110/20] via 12.1.1.1, 00:00:09, Serial1/1
O E2  130.0.0.0/8 [110/20] via 12.1.1.1, 00:00:09, Serial1/1

R2#Show ip ospf database database-summary

            OSPF Router with ID (2.2.2.2) (Process ID 1)

Area 0 database summary
  LSA Type      Count    Delete   Maxage
  Router        3        0        0
  Network       1        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-7 Ext    0        0        0
    Prefixes redistributed in Type-7  0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
  Subtotal      4        0        0

Process 1 database summary
  LSA Type      Count    Delete   Maxage
  Router        3        0        0
  Network       1        0        0
  Summary Net   0        0        0
  Summary ASBR  0        0        0
  Type-7 Ext    0        0        0
  Opaque Link   0        0        0
  Opaque Area   0        0        0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Type-5 Ext     3          0          0
    Prefixes redistributed in Type-5   0
Opaque AS      0          0          0
Non-self       5
Total          7          0          0
```

The total number of non-self generated LSAs are 5, let's add another external prefix and see the result:

On R1:

```
R1(config)#Route-map tst per 10
R1(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4
```

To verify the configuration:

On R2:

```
R2#Show ip ospf database database-summary | Inc Non

  Non-self       6
```

Let's add another two, this should generate a warning message because the number of Non-self generated routes will reach 80 percent of the set threshold:

On R1:

```
R1(config)#Route-map tst per 10
R1(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4 Lo5 lo6
```

To verify the configuration:

On R2:

```
R2#Show ip ospf database database-summary | Inc Non-self

  Non-self       8
```

Sure enough when the 80 percent of the configured threshold is reached a console message is generated:

```
%OSPF-4-OSPF_MAX_LSA_THR: Threshold for maximum number of non self-
generated LSA has been reached "ospf 1" - 8 LSAs

R2#Sh ip ospf da database-summary | Inc Non-self
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Non-self      8
```

Let's add another two, this should reach the configured threshold:

On R1:

```
R1(config)#Route-map tst per 10
R1(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4 Lo5 Lo6 lo7 lo8
```

To verify the configuration:

On R2:

```
R2#Show ip ospf database database-summary | Inc Non-self

  Non-self      10
```

Let's add another one, this should exceed the threshold for number of Non-self generated routes:

On R1:

```
R1(config)#Route-map tst per 10
R1(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4 Lo5 Lo6 Lo7 lo8 lo9
```

On R2:

The following console message will be generated for every additional route above the configured threshold.

```
%OSPF-4-OSPF_MAX_LSA: Maximum number of non self-generated LSA has been
exceeded "ospf 1" - 11 LSAs

R2#Show ip ospf database database-summary | Inc Non-self

  Non-self      11
```

To verify the configuration:

On R2:

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


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
O     1.0.0.0/8 [110/782] via 12.1.1.1, 00:17:57, Serial1/1
O     3.0.0.0/8 [110/2] via 23.1.1.3, 00:40:59, FastEthernet0/0
O E2  110.0.0.0/8 [110/20] via 12.1.1.1, 00:17:57, Serial1/1
O E2  120.0.0.0/8 [110/20] via 12.1.1.1, 00:17:57, Serial1/1
O E2  130.0.0.0/8 [110/20] via 12.1.1.1, 00:17:57, Serial1/1
O E2  140.0.0.0/8 [110/20] via 12.1.1.1, 00:15:46, Serial1/1
O E2  150.0.0.0/8 [110/20] via 12.1.1.1, 00:13:56, Serial1/1
O E2  160.0.0.0/8 [110/20] via 12.1.1.1, 00:13:56, Serial1/1
O E2  170.0.0.0/8 [110/20] via 12.1.1.1, 00:06:16, Serial1/1
O E2  180.0.0.0/8 [110/20] via 12.1.1.1, 00:06:16, Serial1/1
O E2  190.0.0.0/8 [110/20] via 12.1.1.1, 00:02:12, Serial1/1

R2#Show ip ospf 1 | S non

 Maximum number of non self-generated LSA allowed 10 (warning-only)
    Current number of non self-generated LSA 11
    Threshold for warning message 80%
```

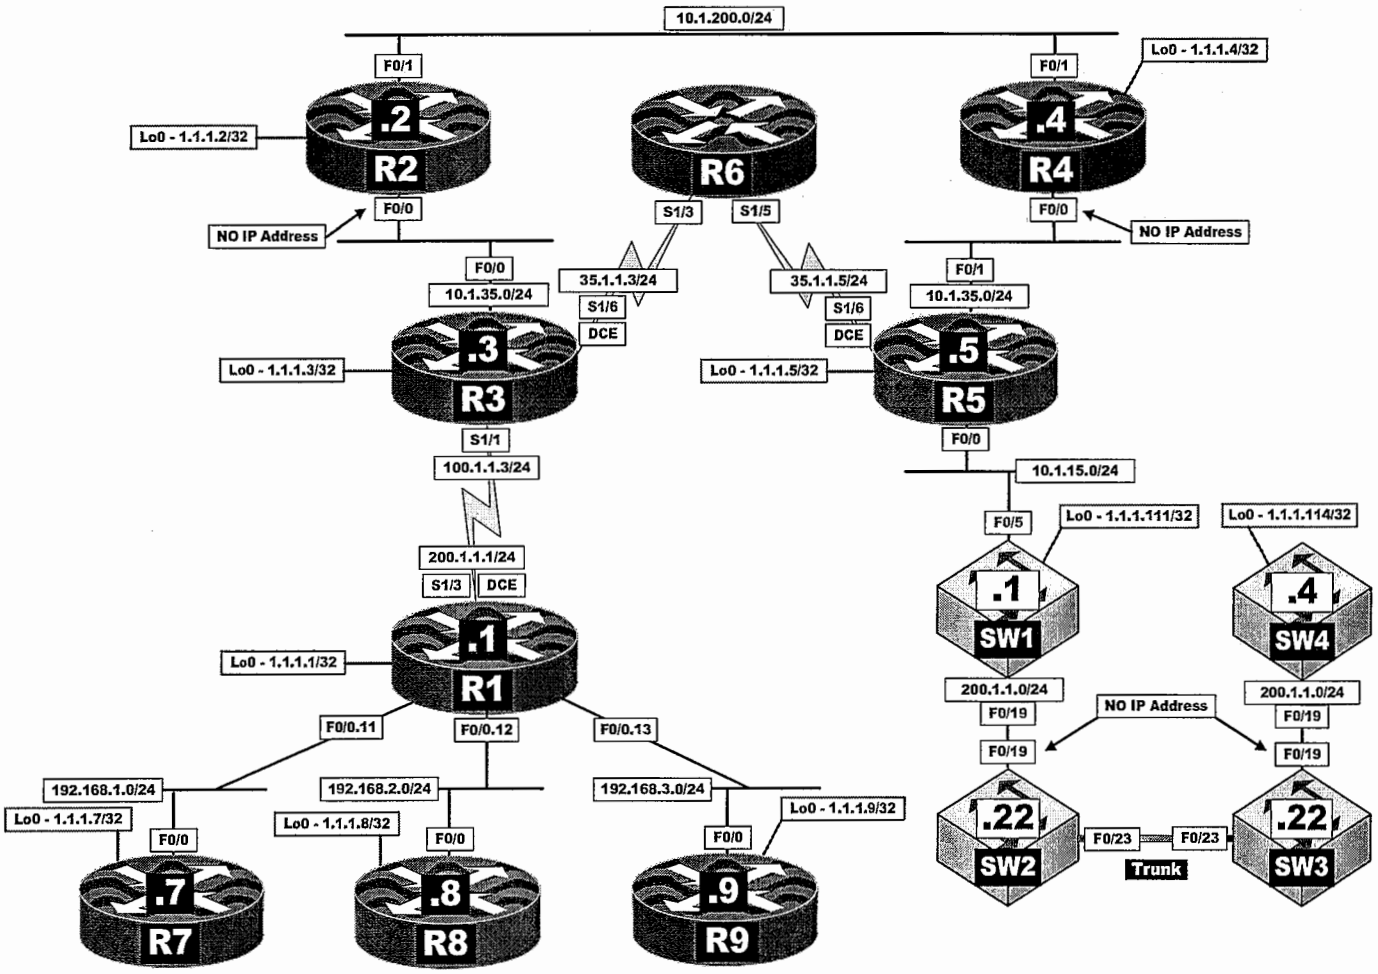
Task 9

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 2
Establishing OSPF Adjacency



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "OSPF"→ "Lab-2".

Lab Rules:

- DO NOT use the same solution more than once

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 1

Configure OSPF Area 0 on the link that connects R1 to R3 and their Lo0 interface, if this configuration is successful, these two routers should exchange hellos through their Serial interfaces and form an adjacency and exchange routes. DO NOT change R1 or R3's IP address to accomplish this task. These routers should use their Lo0 as their OSPF router-id.

In order for two OSPF speaking routers to form an adjacency they have to adhere to many rules, one of which states that they must be on the same subnet, this is like a sanity check. This rule can be bypassed if the "IP unnumbered" is used in a point-to-point configuration.

When a router receives routing update/s, it installs the source IP address of the update/s as the next hop in its routing table. Normally, the next hop is a directly-connected router. When the "IP unnumbered" is configured, the Serial interface "borrows" its IP address from the configured Loopback interface and routes learned through the IP unnumbered interface have the interface as the next hop instead of the source IP address of the routing update. Thus avoiding an invalid next hop address problem.

On R1:

```
R1#Show ip int brief S1/3

Interface              IP-Address      OK? Method Status                Protocol
Serial1/3              200.1.1.1       YES manual up                    up
```

On R3:

```
R3#Show ip int brief S1/1

Interface              IP-Address      OK? Method Status                Protocol
Serial1/1              100.1.1.3       YES manual up                    up
```

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 1.1.1.1
R1(config-router)#netw 200.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config)#int S1/3
R1(config-if)#IP unnumbered Lo1

R1(config-if)#int lo1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#ip addr 200.1.1.1 255.255.255.0
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 1.1.1.3
R3(config-router)#Netw 1.1.1.3 0.0.0.0 area 0
R3(config-router)#Netw 100.1.1.3 0.0.0.0 area 0

R3(config)#Int S1/1
R3(config-if)#IP unnumbered Lo1

R3(config-if)#int Lo1
R3(config-if)#ip addr 100.1.1.3 255.255.255.0
```

You should see the following message on R3:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R1:

```
R1#Show ip int brief S1/3

Interface        IP-Address      OK? Method Status                Protocol
Serial1/3        200.1.1.1       YES TFTP   up                    up
```

On R3

```
R3#Show ip int brief S1/1

Interface        IP-Address      OK? Method Status                Protocol
Serial1/1        100.1.1.3       YES TFTP   up                    up

R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.1           0   FULL/  -        00:00:38    200.1.1.1       Serial1/1

R3#Show ip route ospf | 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, 2 subnets
O        1.1.1.1 [110/782] via 200.1.1.1, 00:01:37, Serial1/1
      200.1.1.0/32 is subnetted, 1 subnets
O        200.1.1.1 [110/782] via 200.1.1.1, 00:01:37, Serial1/1
```

On R1:

To check the routing table:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.3 [110/782] via 100.1.1.3, 00:02:38, Serial1/3
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.3 [110/782] via 100.1.1.3, 00:02:38, Serial1/3
```

To test NLRI:

```
R1#Ping 1.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
```

Task 2

Configure OSPF Area 0 on the F0/1 and Lo0 interfaces of R2 and R4. These routers should use their Lo0 interface as their OSPF router-id.

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#Router-id 1.1.1.2
R2(config-router)#Netw 10.1.200.2 0.0.0.0 area 0
R2(config-router)#Netw 1.1.1.2 0.0.0.0 area 0
```

On R4:

```
R4(config)#Router ospf 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config-router)#router-id 1.1.1.4
R4(config-router)#Netw 10.1.200.4 0.0.0.0 area 0
R4(config-router)#Netw 1.1.1.4 0.0.0.0 area 0
```

To verify the configuration:

On R2:

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

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.4 [110/2] via 10.1.200.4, 00:00:03, FastEthernet0/1
```

On R4:

```
R4#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.2 [110/2] via 10.1.200.2, 00:00:48, FastEthernet0/1
```

Task 3

Configure OSPF area0 on R3's F0/0 and R5's F0/1 and their Lo0 interfaces. DO NOT configure extra IP addresses on another router, Change IP addresses, GRE tunnels or an "IPnIP" tunnel or IRB. You should NOT change the VLAN assignments or configure the switches to accomplish this task. R5 should use its Lo0 as its OSPF router-id. If this configuration is successful R5 should receive all the routes advertised by the other routers in this area.

AToM can be used to accomplish this task. AToM provides a layer two connectivity to the attached routers (R3 and R5). This could have been done using L2TPv3, but since we could NOT use or change the IP addressing on any of the devices AToM was our best choice. Both AToM and L2TPv3 provide a Point-to-Point tunnel called a "Pseudowire".

AToM uses targeted LDP sessions between the routers to form an AToM VC, for this to work the two routers have to go through a discovery process using UDP port 646 and once the discovery process is completed successfully, a TCP session is established using TCP port 646 on the server side and a TCP high port on the client side.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To configure AToM successfully, the VC identifiers MUST match on both R2 and R4's configuration and the configured IP address (The Lo0 in his case) must be reachable.

On R2 and R4:

```
Rx(config)#int F0/1
Rx(config-if)#MPLS IP

Rx(config)#MPLS LDP router-id Lo0
```

To verify the configuration:

On R2:

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 1.1.1.4:0 (1) is UP

R2#Show mpls ldp discovery detail

 Local LDP Identifier:
    1.1.1.2:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 1.1.1.2
            LDP Id: 1.1.1.4:0
              Src IP addr: 10.1.200.4; Transport IP addr: 1.1.1.4
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 1.1.1.4/32
```

On R4:

```
R4#Show mpls ldp discovery detail

 Local LDP Identifier:
    1.1.1.4:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 1.1.1.4
            LDP Id: 1.1.1.2:0
              Src IP addr: 10.1.200.2; Transport IP addr: 1.1.1.2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
            Hold time: 15 sec; Proposed local/peer: 15/15 sec
            Reachable via 1.1.1.2/32
```

On R2:

```
R2#Show ip route 1.1.1.4

Routing entry for 1.1.1.4/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.1.200.4 on FastEthernet0/1, 00:07:58 ago
  Routing Descriptor Blocks:
  * 10.1.200.4, from 1.1.1.4, 00:07:58 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1

R2#Ping 1.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R2(config)#Int F0/0
R2(config-if)#Xconnect 1.1.1.4 24 encapsulation MPLS
```

On R4:

```
R4#Show ip route 1.1.1.2

Routing entry for 1.1.1.2/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.1.200.2 on FastEthernet0/1, 00:07:44 ago
  Routing Descriptor Blocks:
  * 10.1.200.2, from 1.1.1.2, 00:07:44 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1

R4#Ping 1.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R4(config)#Int F0/0
R4(config-if)#Xconnect 1.1.1.2 24 encapsulation MPLS
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On R2:

```
R2#Show mpls l2transport vc

Local intf     Local circuit                 Dest address    VC ID      Status
-------------  ----------------------------  --------------  ---------  -------
Fa0/0          Ethernet                      1.1.1.4         24         UP
```

On R4:

```
R4#Show mpls l2transport vc

Local intf     Local circuit                 Dest address    VC ID      Status
-------------  ----------------------------  --------------  ---------  ----------
Fa0/0          Ethernet                      1.1.1.2         24         UP
```

Since the Layer 2 tunnel is configured successfully, OSPF can be configured on R3 and R5 so they can establish an OSPF adjacency:

On R5:

```
R5(config)#router ospf 1
R5(config-router)#router-id 1.1.1.5
R5(config-router)#Netw 1.1.1.5 0.0.0.0 area 0
R5(config-router)#Netw 10.1.35.5 0.0.0.0 area 0
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Netw 10.1.35.3 0.0.0.0 area 0
```

You should see the following console message on R3:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.5 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

To verify the configuration:

On R3:

```
R3#Show ip ospf neighbor
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Neighbor ID      Pri   State           Dead Time   Address         Interface
1.1.1.5            1   FULL/DR         00:00:32    10.1.35.5       FastEthernet0/0
1.1.1.1            0   FULL/  -        00:00:37    200.1.1.1       Serial1/1
```

```
R3#Show ip route ospf | B Gate

      1.0.0.0/32 is subnetted, 3 subnets
O        1.1.1.1 [110/782] via 200.1.1.1, 00:24:39, Serial1/1
O        1.1.1.5 [110/2] via 10.1.35.5, 00:01:49, FastEthernet0/0
      200.1.1.0/32 is subnetted, 1 subnets
O        200.1.1.1 [110/782] via 200.1.1.1, 00:24:39, Serial1/1
```

On R5:

```
R5#Show ip ospf neighbor

Neighbor ID      Pri   State           Dead Time   Address         Interface
1.1.1.3            1   FULL/BDR        00:00:33    10.1.35.3       FastEthernet0/1
```

```
R5#Show ip route ospf | B Gate

      1.0.0.0/32 is subnetted, 3 subnets
O        1.1.1.1 [110/783] via 10.1.35.3, 00:02:41, FastEthernet0/1
O        1.1.1.3 [110/2] via 10.1.35.3, 00:02:41, FastEthernet0/1
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.3 [110/2] via 10.1.35.3, 00:02:41, FastEthernet0/1
      200.1.1.0/32 is subnetted, 1 subnets
O        200.1.1.1 [110/783] via 10.1.35.3, 00:02:41, FastEthernet0/1
```

Task 4

Configure OSPF area 0 on R5's F0/0 and SW1's F0/5 and Lo0 interface. SW1 should use its Lo0 as its OSPF router-id. If this configuration is successful, SW1 should receive all the routes advertised in this routing domain. You should NOT make any additional changes on SW1 to accomplish this task.

On R5:

```
R5(config)#Router ospf 1
R5(config-router)#Netw 10.1.15.5 0.0.0.0 area 0
```

On SW1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1(config)#IP routing

SW1(config)#router ospf 1
SW1(config-router)#router-id 1.1.1.11
SW1(config-router)#Netw 10.1.15.1 0.0.0.0 area 0
SW1(config-router)#Netw 1.1.1.11 0.0.0.0 area 0
```

To verify the configuration:

On SW1:

You may see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.5 on FastEthernet0/5 from EXSTART to
DOWN, Neighbor Down: Too many retransmissions

%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.5 on FastEthernet0/5 from DOWN to
DOWN, Neighbor Down: Ignore timer expired

SW1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.5           1   EXSTART/DR      00:00:38    10.1.15.5       FastEthernet0/5
```

On R5:

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.11 on FastEthernet0/0 from EXSTART
to DOWN, Neighbor Down: Too many retransmissions

%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.11 on FastEthernet0/0 from DOWN to
DOWN, Neighbor Down: Ignore timer expired

R5#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.11          1   EXSTART/BDR     00:00:38    10.1.15.1       FastEthernet0/0
1.1.1.3           1   FULL/BDR        00:00:34    10.1.35.3       FastEthernet0/1
```

To see the problem, let's enable debugging on R5:

On R5:

```
R5#Debug ip ospf adj
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
OSPF-1 ADJ   Fa0/0: Send DBD to 1.1.1.11 seq 0x133E opt 0x52 flag 0x7 len
32
OSPF-1 ADJ   Fa0/0: Retransmitting DBD to 1.1.1.11 [1]
OSPF-1 ADJ   Fa0/0: Rcv DBD from 1.1.1.11 seq 0x24C6 opt 0x52 flag 0x7
len 32  mtu 1504 state EXSTART
OSPF-1 ADJ   Fa0/0: Nbr 1.1.1.11 has larger interface MTU
```

The output of the above debug command reveals that R5 receives the DBD packets but the MTU of the received packets are larger than the configured MTU size on R5. Remember in order for two OSPF speakers to establish an adjacency, the MTU size must match; in this case R5 can be configured to either ignore the received MTU size, or change the routing MTU on SW1 to match R5's F0/0's MTU size. Since the task states NOT to make any changes to SW1, R5 the device with a lower MTU size is configured to ignore the MTU.

```
R5(config)#int F0/0
R5(config-if)#IP OSPF mtu-ignore
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.11 on FastEthernet0/0 from LOADING
to FULL, Loading Done
```

To verify the configuration:

On R5:

```
R5#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.11          1   FULL/DR         00:00:35    10.1.15.1       FastEthernet0/0
1.1.1.3           1   FULL/BDR        00:00:37    10.1.35.3       FastEthernet0/1

R5#Show ip route ospf | Inc O
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 4 subnets
O        1.1.1.1 [110/783] via 10.1.35.3, 00:17:50, FastEthernet0/1
O        1.1.1.3 [110/2] via 10.1.35.3, 00:17:50, FastEthernet0/1
O        1.1.1.11 [110/2] via 10.1.15.1, 00:00:43, FastEthernet0/0
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.3 [110/2] via 10.1.35.3, 00:17:50, FastEthernet0/1
      200.1.1.0/32 is subnetted, 1 subnets
O        200.1.1.1 [110/783] via 10.1.35.3, 00:17:50, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW1:

```
SW1#Show ip route ospf | Inc O

O       1.1.1.1 [110/784] via 10.1.15.5, 00:01:28, FastEthernet0/5
O       1.1.1.3 [110/3] via 10.1.15.5, 00:01:28, FastEthernet0/5
O       1.1.1.5 [110/2] via 10.1.15.5, 00:01:28, FastEthernet0/5
O       100.1.1.3 [110/3] via 10.1.15.5, 00:01:28, FastEthernet0/5
O       10.1.35.0 [110/2] via 10.1.15.5, 00:01:28, FastEthernet0/5
```

Task 5

Configure OSPF area 0 on R1, R7, R8 and R9 based on the following policy:

- R1 should be configured such that it establishes a neighbor adjacency with R7 through its F0/0.11 sub-interface ONLY if R7's F0/0 (192.168.1.7) interface is up; if R1 looses its NLRI to R7's F0/0 interface, within 2 seconds it should establish an OSPF adjacency with R8 through its F0/0.12; R1 should ONLY establish an OSPF adjacency with R9 if it does NOT have NLRI to R7 and R8's F0/0 interfaces.
- The routers should exchange hellos every 250ms with a dead interval of one second.
- R7, R8 and R9 should use their Lo0 interface as their OSPF router-id.
- These routers MUST use Unicast and NOT Multicast for this task.

To configure this task we should configure OSPF and re-adjust the default timers and see the adjacencies between R1 and the other routers first before we configure the other policies.

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#Netw 192.168.1.1 0.0.0.0 area 0
R1(config-router)#Netw 192.168.2.1 0.0.0.0 area 0
R1(config-router)#Netw 192.168.3.1 0.0.0.0 area 0

R1(config)#Int F0/0.11
R1(config-subif)#IP OSPF dead-interval minimal hello-multiplier 4

R1(config)#Int F0/0.12
R1(config-subif)#IP OSPF dead-interval minimal hello-multiplier 4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Int F0/0.13
R1(config-subif)#IP OSPF dead-interval minimal hello-multiplier 4
```

On R7:

```
R7(config)#Router ospf 1
R7(config-router)#router-id 1.1.1.7
R7(config-router)#Netw 192.168.1.7 0.0.0.0 area 0
R7(config-router)#Netw 1.1.1.7 0.0.0.0 area 0

R7(config)#int F0/0
R7(config-if)#IP OSPF dead-interval minimal hello-multiplier 4
```

You should see the following console message stating that the adjacency with R1 is up:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

On R8:

```
R8(config)#Router ospf 1
R8(config-router)#router-id 1.1.1.8
R8(config-router)#Netw 1.1.1.22 0.0.0.0 area 0
R8(config-router)#Netw 192.168.2.8 0.0.0.0 area 0

R8(config)#int F0/0
R8(config-if)#IP OSPF dead-interval minimal hello-multiplier 4
```

You should see the following console message stating that the adjacency with R1 is up:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

On R9:

```
R9(config)#Router ospf 1
R9(config-router)#router-id 1.1.1.9
R9(config-router)#Netw 1.1.1.9 0.0.0.0 area 0
R9(config-router)#Netw 192.168.3.9 0.0.0.0 area 0

R9(config)#int F0/0
R9(config-if)#IP OSPF dead-interval minimal hello-multiplier 4
```

You should see the following console message stating that the adjacency with R1 is up:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on FastEthernet0/0 from LOADING to
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
FULL, Loading Done
```

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   FULL/BDR        756 msec    192.168.3.9     FastEthernet0/0.13
1.1.1.8           1   FULL/BDR        956 msec    192.168.2.8     FastEthernet0/0.12
1.1.1.7           1   FULL/BDR        880 msec    192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:33    100.1.1.3       Serial1/3

R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 7 subnets
O        1.1.1.3 [110/782] via 100.1.1.3, 00:57:44, Serial1/3
O        1.1.1.5 [110/783] via 100.1.1.3, 00:34:44, Serial1/3
O        1.1.1.7 [110/2] via 192.168.1.7, 00:05:14, FastEthernet0/0.11
O        1.1.1.8 [110/2] via 192.168.2.8, 00:03:24, FastEthernet0/0.12
O        1.1.1.9 [110/2] via 192.168.3.9, 00:00:46, FastEthernet0/0.13
O        1.1.1.11 [110/784] via 100.1.1.3, 00:17:38, Serial1/3
      10.0.0.0/24 is subnetted, 2 subnets
O        10.1.15.0 [110/783] via 100.1.1.3, 00:17:48, Serial1/3
O        10.1.35.0 [110/782] via 100.1.1.3, 00:34:54, Serial1/3
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.3 [110/782] via 100.1.1.3, 00:57:44, Serial1/3
```

To the routers to use unicast instead of multicast:

On R1:

```
R1(config)#Int F0/0.11
R1(config-subif)#IP OSPF Network Non-Broadcast

R1(config-subif)#Int F0/0.12
R1(config-subif)#IP OSPF Network Non-Broadcast

R1(config-subif)#Int F0/0.13
R1(config-subif)#IP OSPF Network Non-Broadcast

R1(config)#Router ospf 1
R1(config-router)#Neighbor 192.168.1.7
R1(config-router)#Neighbor 192.168.2.8
R1(config-router)#Neighbor 192.168.3.9
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   FULL/DR         964 msec    192.168.3.9     FastEthernet0/0.13
1.1.1.8           1   FULL/DR         988 msec    192.168.2.8     FastEthernet0/0.12
1.1.1.7           1   FULL/DR         940 msec    192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:32    100.1.1.3       Serial1/3
```

On R1:

```
R1(config)#Access-list 102 permit ospf host 192.168.2.1 host 192.168.2.8
R1(config)#Access-list 103 permit ospf host 192.168.3.1 host 192.168.3.9

R1(config)#IP SLA 11
R1(config-rtr)#icmp-echo 192.168.1.7 source-ip 192.168.1.1
R1(config-rtr-echo)#Freq 5

R1(config)#IP SLA schedule 11 start now life forever

R1(config)#Track 1 ip sla 11 reachability
R1(config-track)#delay down 2
R1(config-track)#delay up 2

R1(config)#IP SLA 12
R1(config-rtr)#icmp-echo 192.168.2.8 source-ip 192.168.2.1
R1(config-rtr-echo)#Freq 5

R1(config)#IP SLA schedule 12 start now life forever

R1(config)#Track 2 ip sla 12 reachability
R1(config-track)#delay down 2
R1(config-track)#delay up 2

R1(config)#Route-map TST permit 10
R1(config-route-map)#match ip addr 102 103

R1(config-route-map)#set ip next-hop verify-availability 192.168.1.7 1 track 1
R1(config-route-map)#set ip next-hop verify-availability 192.168.2.8 2 track 2
R1(config)#Route-map TST permit 90
```

NOTE: The set ip next-hop commands are order specific, and the order is configured using the value after

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

the IP address.

```
R1(config)#ip local policy route-map TST
```

To verify the configuration:

NOTE: As long as the connection between R1 and R7 is up, R1 will ONLY establish an OSPF adjacency with R7 and NO other router.

On R1:

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.9 on FastEthernet0/0.13 from FULL to
DOWN, Neighbor Down: Dead timer expired

%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.8 on FastEthernet0/0.12 from FULL to
DOWN, Neighbor Down: Dead timer expired

R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   INIT/DROTHER    868 msec    192.168.3.9     FastEthernet0/0.13
1.1.1.8           1   INIT/DROTHER    908 msec    192.168.2.8     FastEthernet0/0.12
1.1.1.7           1   FULL/DR         772 msec    192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:37    100.1.1.3       Serial1/3
```

To test the configuration:

NOTE: Once R1 loses its NLRI to R7, it establishes an adjacency with R8:

On R7:

```
R7(config)#Int F0/0
R7(config-if)#Shut
```

On R1:

You should see the following console messages stating that R1 has lost its OSPF adjacency to R7 and it has established an OSPF adjacency with R8:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.7 on FastEthernet0/0.11 from FULL to
DOWN, Neighbor Down: Dead timer expired

%TRACKING-5-STATE: 1 ip sla 11 state Up->Down
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.8 on FastEthernet0/0.12 from LOADING
to FULL, Loading Done

R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   INIT/DROTHER    788 msec    192.168.3.9     FastEthernet0/0.13
1.1.1.8           1   FULL/DR         852 msec    192.168.2.8     FastEthernet0/0.12
N/A               0   ATTEMPT/DROTHER     -       192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:38    100.1.1.3       Serial1/3
```

BUT if R1 loses its adjacency to R8 and R7, ONLY then it will establish an adjacency with R9, let's test this sceanrio:

On R8:

```
R8(config)#int f0/0
R8(config-if)#Shut
```

On R1:

We see similar console messages when R7 went down, the only difference here is that R1 is establishing an OSPF adjacency with R9:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.8 on FastEthernet0/0.12 from FULL to
DOWN, Neighbor Down: Dead timer expired

%TRACKING-5-STATE: 2 ip sla 12 reachability Up->Down

%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.9 on FastEthernet0/0.13 from LOADING
to FULL, Loading Done

R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   FULL/DR         936 msec    192.168.3.9     FastEthernet0/0.13
N/A               0   ATTEMPT/DROTHER     -       192.168.2.8     FastEthernet0/0.12
N/A               0   ATTEMPT/DROTHER     -       192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:33    100.1.1.3       Serial1/3
```

To test the failover:

NOTE: Once R8 is up and its reachable by R1, R1 drops its adjacency to R9 and it re-establishes its adjacency with R8, let's test this sceanrio:

On R8:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R8(config)#Int F0/0
R8(config-if)#No Shut
```

On R1:

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.8 on FastEthernet0/0.12 from LOADING
to FULL, Loading Done

%TRACKING-5-STATE: 2 ip sla 12 reachability Down->Up

R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   INIT/DROTHER    968 msec    192.168.3.9     FastEthernet0/0.13
1.1.1.8           1   FULL/BDR        928 msec    192.168.2.8     FastEthernet0/0.12
N/A               0   ATTEMPT/DROTHER     -       192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:30    100.1.1.3       Serial1/3
```

But if R7 comes up, R1 drops its adjacency to R8 and R9 and once again it re-establishes its OSPF adjacency with R7, let's test this sceanrio:

On R7:

```
R7(config)#Int F0/0
R7(config-if)#No Shut
```

On R1:

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.7 on FastEthernet0/0.11 from LOADING
to FULL, Loading Done

%TRACKING-5-STATE: 1 ip sla 11 state Down->Up

R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.9           1   INIT/DROTHER    880 msec    192.168.3.9     FastEthernet0/0.13
1.1.1.8           1   INIT/DROTHER    892 msec    192.168.2.8     FastEthernet0/0.12
1.1.1.7           1   FULL/BDR        824 msec    192.168.1.7     FastEthernet0/0.11
1.1.1.3           0   FULL/  -        00:00:34    100.1.1.3       Serial1/3
```

Task 6

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Configure OSPF area 0 on the F0/19 interfaces of SW1 and SW4, and the Loopback0 interface of SW4. You should NOT configure a layer 3 link between SW2 and SW3 to accomplish this task. If this configuration is done successfully, SW4 should receive all OSPF routes. You should use VLAN 100, DO NOT configure another VLAN, or configure additional IP address/es to accomplish this task.

This task calls for Dot1q Tunneling, 802.1q tunneling enables the service providers to use a single VLAN to support customers who have a single or multiple VLANs that need to connect across the provider's network while preserving their VLAN-IDs. The provider can use this feature to keep traffic from different customers segregated.

When configuring QinQ tunneling, a tunnel port must be defined, this port should be assigned to a VLAN, different customers must be assign to different tunnel ports and different tunnel ports should be configured in different provider VLANs, and this is how the traffic from different customers are segregated.

When a given tunnel port receives customer traffic, it adds a 2 Byte Ether-Type field of 0x8100 followed by a 2 Byte field containing the CoS and the VLAN and this traffic is then placed into the VLAN to which the tunnel port is assigned. The Egress tunnel port strips off the 4 Bytes that was added by the ingress tunnel port and transmits the traffic to the customer device.

When Dot1q tunneling is configured, a layer 2 protocol tunneling can also be configured, a layer 2 protocol tunneling allows layer 2 protocol data units (PDUs) to be tunneled through the network, the layer 2 protocols that can be tunneled are: CDP, STP and VTP and they need to be configured or else they will NOT get propagated across the tunnel.

To configure DOT1Q on SW2 and SW3:

On SW2:

```
SW2(config)#int f0/19
SW2(config-if)#Swi acc v 100
SW2(config-if)#swi mode dot1q-tunnel
```

On SW3:

```
SW3(config)#int f0/19
SW3(config-if)#swi acc v 100
SW3(config-if)#swi mode dot1q-tunnel
```

To verify the configuration:

On SW2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW2#Show dot1q-tunnel

dot1q-tunnel mode LAN Port(s)
------------------------------

Fa0/19
```

On SW3:

```
SW3#Show dot1q-tunnel

dot1q-tunnel mode LAN Port(s)
------------------------------
Fa0/19
```

To configure SW1 and SW4:

On SW1:

```
SW1(config)#Router ospf 1
SW1(config-router)#Netw 200.1.1.1 0.0.0.0 area 0
```

On SW4:

```
SW4(config)#ip routing

SW4(config)#Router ospf 1
SW4(config-router)#router-id 1.1.1.114
SW4(config-router)#Netw 1.1.1.114 0.0.0.0 area 0
SW4(config-router)#Netw 200.1.1.4 0.0.0.0 area 0
```

To verify the configuration:

On SW4:

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.11 on FastEthernet0/19 from LOADING
to FULL, Loading Done

SW4#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.11          1   FULL/DR         00:00:34    200.1.1.1       FastEthernet0/19

SW4#Show ip route ospf | Inc O
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O       1.1.1.1 [110/785] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       1.1.1.3 [110/4] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       1.1.1.5 [110/3] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       1.1.1.7 [110/786] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       1.1.1.11 [110/2] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       100.1.1.3 [110/4] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       200.1.1.1/32 [110/785] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       10.1.15.0 [110/2] via 200.1.1.1, 00:01:55, FastEthernet0/19
O       10.1.35.0 [110/3] via 200.1.1.1, 00:01:56, FastEthernet0/19
O    192.168.1.0/24 [110/785] via 200.1.1.1, 00:01:56, FastEthernet0/19
O    192.168.2.0/24 [110/785] via 200.1.1.1, 00:01:56, FastEthernet0/19
O    192.168.3.0/24 [110/785] via 200.1.1.1, 00:01:56, FastEthernet0/19
```

Task 7

Configure the second adjacency between R3 and R5 based on the following policy:
R3 and R5 should establish their second OSPF adjacency through their S1/6 interfaces.
You should use **an** IP address of your choice to accomplish this task.

On R6:

```
R6(config)#bridge irb

R6(config)#bridge 1 protocol ieee

R6(config)#int s1/3
R6(config-subif)#bridge-group 1

R6(config-subif)#int s1/5
R6(config-subif)#bridge-group 1

R6(config)#int bvi 1
R6(config-if)#mac-address 0000.6666.6666
R6(config-if)#ip addr 35.1.1.6 255.255.255.0

R6(config)#bridge 1 route ip
```

On R3:

```
R3(config)#bridge irb
R3(config)#bridge 1 protocol ieee
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config)#int S1/6
R3(config-subif)#No ip addr
R3(config-subif)#bridge-group 1

R3(config-subif)#int bvi 1
R3(config-if)#ip addr 35.1.1.3 255.255.255.0
R3(config-if)#mac-address 0000.3333.3333

R3(config)#bridge 1 route ip
```

On R5:

```
R5(config)#bridge irb
R5(config)#bridge 1 protocol ieee

R5(config)#int s1/6
R5(config-subif)#No ip addr
R5(config-subif)#bridge-group 1

R5(config)#int bvi 1
R5(config-if)#mac-address 0000.5555.5555
R5(config-if)#ip addr 35.1.1.5 255.255.255.0

R5(config)#bridge 1 route ip
```

To test the configuration:

On R6:

```
R6#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 80 percent (4/5), round-trip min/avg/max = 60/60/60 ms

R6#Ping 35.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 35.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 64/64/64 ms
```

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Ping 35.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 35.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 120/123/124 ms
```

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 100 percent (5/5), round-trip min/avg/max = 120/122/124 ms
```

To configure OSPF on R3 and R5:

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Netw 35.1.1.3 0.0.0.0 area 0
```

On R5:

```
R5(config)#Router ospf 1
R5(config-router)#Netw 35.1.1.5 0.0.0.0 area 0
```

Within a minute you should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.3 on BVI1 from LOADING to FULL, Loa
Done
```

To verify the configuration:

On R5:

```
R5#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.3           1   FULL/DR         00:00:39    35.1.1.3        BVI1
1.1.1.11          1   FULL/DR         00:00:39    10.1.15.1       FastEthernet0/0
1.1.1.3           1   FULL/BDR        00:00:37    10.1.35.3       FastEthernet0/1
```

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
1.1.1.5           1   FULL/BDR        00:00:37    35.1.1.5        BVI1
1.1.1.5           1   FULL/DR         00:00:35    10.1.35.5       FastEthernet0/0
1.1.1.1           0   FULL/  -        00:00:37    200.1.1.1       Serial1/1
```

To test the configuration:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 5 subnets
O        1.1.1.1 [110/783] via 10.1.35.3, 03:40:09, FastEthernet0/1
O        1.1.1.3 [110/2] via 10.1.35.3, 03:40:09, FastEthernet0/1
O        1.1.1.7 [110/784] via 10.1.35.3, 00:28:49, FastEthernet0/1
O        1.1.1.11 [110/2] via 10.1.15.1, 03:23:02, FastEthernet0/0
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.3 [110/2] via 10.1.35.3, 03:40:09, FastEthernet0/1
O     192.168.1.0/24 [110/783] via 10.1.35.3, 00:52:39, FastEthernet0/1
O     192.168.2.0/24 [110/783] via 10.1.35.3, 00:49:01, FastEthernet0/1
O     192.168.3.0/24 [110/783] via 10.1.35.3, 00:32:08, FastEthernet0/1
      200.1.1.0/32 is subnetted, 1 subnets
O        200.1.1.1 [110/783] via 10.1.35.3, 03:40:09, FastEthernet0/1
```

To test the new connection:

```
R5(config)#int f0/1
R5(config-if)#Shut

R5#Show ip route ospf | Inc O

      1.0.0.0/32 is subnetted, 5 subnets
O        1.1.1.1 [110/1563] via 35.1.1.3, 00:00:18, BVI1
O        1.1.1.3 [110/782] via 35.1.1.3, 00:00:18, BVI1
O        1.1.1.7 [110/1564] via 35.1.1.3, 00:00:18, BVI1
O        1.1.1.11 [110/2] via 10.1.15.1, 03:24:54, FastEthernet0/0
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.3 [110/782] via 35.1.1.3, 00:00:18, BVI1
O     192.168.1.0/24 [110/1563] via 35.1.1.3, 00:00:18, BVI1
O     192.168.2.0/24 [110/1563] via 35.1.1.3, 00:00:18, BVI1
O     192.168.3.0/24 [110/1563] via 35.1.1.3, 00:00:18, BVI1
      200.1.1.0/32 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O        200.1.1.1 [110/1563] via 35.1.1.3, 00:00:18, BVI1

R5#Ping 192.168.1.7

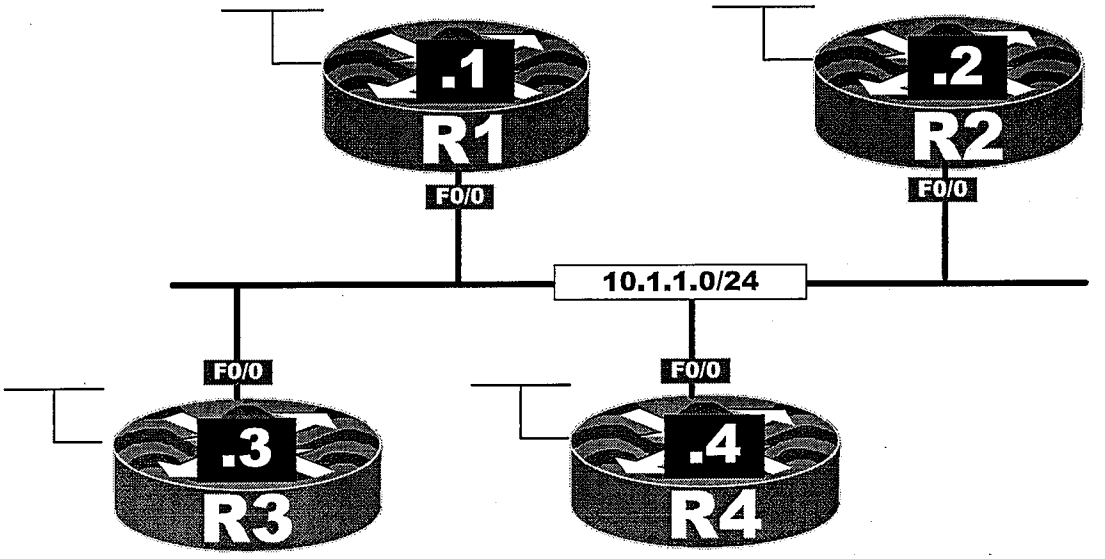
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.7, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 148/152/160 ms
```

Task 8

Erase the startup configuration of the routers, config.text and the VLAN.dat of the switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 3 – OSPF Broadcast Networks



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "OSPF"→ "Lab-3".

Task 1

Configure OSPF area 0 on all routers and run their directly connected interfaces. Ensure that loopback 0 interfaces are advertised with their correct mask. You should configure 0.0.0.1, 0.0.0.2, 0.0.0.3 and 0.0.0.4 as the router IDs of R1, R2, R3 and R4 respectively.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#netw 10.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#netw 10.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0

R4(config-router)#int lo0
R4(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
O     2.0.0.0/8 [110/2] via 10.1.1.2, 00:00:20, FastEthernet0/0
O     3.0.0.0/8 [110/2] via 10.1.1.3, 00:00:20, FastEthernet0/0
O     4.0.0.0/8 [110/2] via 10.1.1.4, 00:00:30, FastEthernet0/0

R1#Show ip ospf database

            OSPF Router with ID (0.0.0.1) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
0.0.0.1         0.0.0.1         90          0x80000004 0x00D02E 2
0.0.0.2         0.0.0.2         91          0x80000004 0x00E119 2
0.0.0.3         0.0.0.3         91          0x80000004 0x00F204 2
0.0.0.4         0.0.0.4         91          0x80000004 0x0004EE 2

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.4        0.0.0.4         90          0x80000001 0x0072A1

R1#Show ip ospf database Network

            OSPF Router with ID (0.0.0.1) (Process ID 1)

                Net Link States (Area 0)

  Routing Bit Set on this LSA in topology Base with MTID 0
  LS age: 189
  Options: (No TOS-capability, DC)
  LS Type: Network Links
  Link State ID: 10.1.1.4 (address of Designated Router)
  Advertising Router: 0.0.0.4
  LS Seq Number: 80000001
  Checksum: 0x72A1
  Length: 40
  Network Mask: /24
        Attached Router: 0.0.0.4
        Attached Router: 0.0.0.1
        Attached Router: 0.0.0.2
        Attached Router: 0.0.0.3
```

We can see that it's the DR that advertises LSA Type 2s, and on this segment R4 is the DR with an IP address of 10.1.1.4/24 which means 10.1.1.0/24 is the network address of this segment, this segment will not be advertised by the other routers, ONLY the DR is responsible for this. Network LSAs or LSA

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Type-2s also reveal the router IDs of the other routers that are attached to this Broadcast Multiaccess network.

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           1   2WAY/DROTHER    00:00:30    10.1.1.2        FastEthernet0/0
0.0.0.3           1   FULL/BDR        00:00:30    10.1.1.3        FastEthernet0/0
0.0.0.4           1   FULL/DR         00:00:39    10.1.1.4        FastEthernet0/0
```

Based on the above output we can see that the local router is in 2WAY state with R2 and in FULL state with DR and BDR. Let's check the routing table of R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

O       2.0.0.0/8 [110/2] via 10.1.1.2, 00:13:30, FastEthernet0/0
O       3.0.0.0/8 [110/2] via 10.1.1.3, 00:13:30, FastEthernet0/0
O       4.0.0.0/8 [110/2] via 10.1.1.4, 00:13:40, FastEthernet0/0
```

We can see the next hop to reach any network is the router that originated that particular network. Let's see how often these routers exchange Hellos:

```
R1#Show ip ospf interface F0/0 | Inc Timer

  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

Let's see the destination address of OSPF hello messages on OSPF Broadcast network types:

```
R1#Debug ip ospf hello
OSPF hello debugging is on

OSPF-1 HELLO Fa0/0: Send hello to 224.0.0.5 area 0 from 10.1.1.1

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

Let's identify the major points of an OSPF Broadcast Network type:

- Ethernet Networks default to OSPF Broadcast Network types.
- The timers are 10/40, meaning that the Hello interval is 10 seconds and the dead interval is set to 40 seconds.
- Next hop does NOT change. In the output of the above "Show ip route ospf" command we can see that the next hop IP address is the IP address of the F0/0 interface of the router that

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

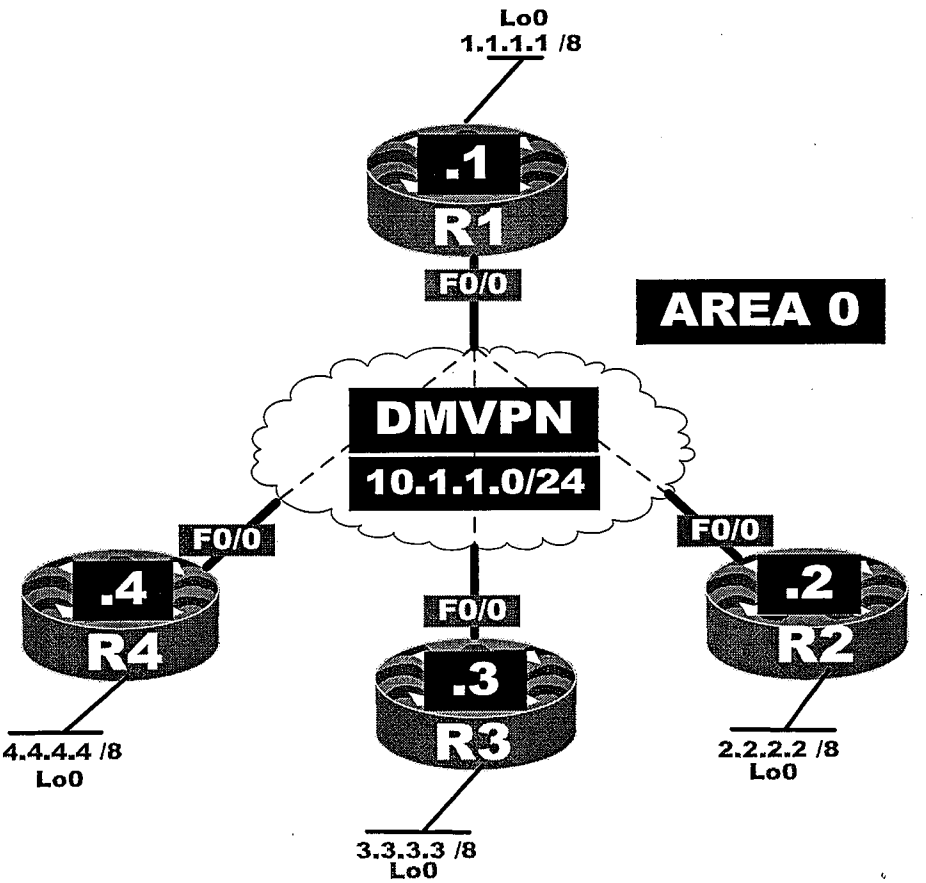
originated the route.

- DR and BDR election will take place in Broadcast Multiaccess networks such as Ethernet.
- In Broadcast network types the hellos are sent to Multicast destination of 224.0.0.5.

Task 2

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "OSPF"→ "Lab-3-Task-2".



Configure OSPF on the tunnel and Loopback 0 interfaces of all routers based on the following policy:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

- R1 is the hub and R2, R3, and R4 are configured as spokes, DO NOT change the topology.
- Configure the tunnel interfaces of all routers to be OSPF broadcast network type.
- The Loopback interfaces should be advertised with their correct mask.
- Configure the router-IDs of the routers to be 0.0.0.1, 0.0.0.2, 0.0.0.3, and 0.0.0.4 for R1, R2, R3 and R4 respectively.
- You should use static maps on the DMVPN network.

On All Routers:

```
Rx(config-router)#int lo0
Rx(config-if)#ip ospf network point-to-point
```

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Network 10.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 1.1.1.1 0.0.0.0 area 0

R1(config)#Int tunne 1234
R1(config-if)#IP ospf network broadcast
```

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Network 2.2.2.2 0.0.0.0 area 0
R2(config-router)#Network 10.1.1.2 0.0.0.0 area 0

R2(config)#Int tunnel 1234
R2(config-if)#IP ospf network broadcast
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Network 3.3.3.3 0.0.0.0 area 0
R3(config-router)#Network 10.1.1.3 0.0.0.0 area 0

R3(config)#Int tunnel 1234
R3(config-if)#IP ospf network broadcast
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#Network 4.4.4.4 0.0.0.0 area 0
R4(config-router)#Network 10.1.1.4 0.0.0.0 area 0

R4(config)#Int tunnel 1234
R4(config-if)#IP ospf network broadcast
```

To verify the configuration:

On R1:

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

We can see that the routers did not establish an OSPF adjacency. We know that in OSPF Broadcast Network types, the Hellos are sent to destination Multicast address of 224.0.0.5, let's check and see if our Network allows Multicasts traffic through:

On R1:

```
R1#Show run int tunnel 1234 | B interface

interface Tunnel1234
 ip address 10.1.1.1 255.255.255.0
 no ip redirects
 ip nhrp map 10.1.1.4 192.1.4.4
 ip nhrp map 10.1.1.3 192.1.3.3
 ip nhrp map 10.1.1.2 192.1.2.2
 ip nhrp network-id 111
 ip ospf network broadcast
 tunnel source FastEthernet0/0
 tunnel mode gre multipoint
end
```

Based on the above output, we can see that Multicast is NOT mapped, let's check R2:

On R2:

```
R2#Show run int tunnel 1234 | B interface

interface Tunnel1234
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 ip address 10.1.1.2 255.255.255.0
 no ip redirects
 ip nhrp map 10.1.1.1 192.1.1.1
 ip nhrp network-id 222
 ip ospf network broadcast
 tunnel source FastEthernet0/0
 tunnel mode gre multipoint
end
```

R2 is configured the same. Let's map Multicast on R1 and R2 and see the result before we move on to the other routers:

On R2:

```
R2(config)#Int tunnel 1234
R2(config-if)#ip nhrp map multicast 192.1.1.1
```

On R1:

```
R1(config)#Int tunnel 1234
R1(config-if)#ip nhrp map multicast 192.1.2.2
```

If both ends are not configured to map Multicast for each other's tunnel IP addresses, the OSPF adjacency will be established and torn down and you will get the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Tunnel1234 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Tunnel1234 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

Let's verify the adjacency and see if these two routers are exchanging routes:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           1   FULL/DR         00:00:33    10.1.1.2        Tunnel1234

R1#Show ip route ospf | B Gate
Gateway of last resort is not set

O     2.0.0.0/8 [110/1001] via 10.1.1.2, 00:01:08, Tunnel1234
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Let's configure the spoke routers to map Multicast traffic:

On R1:

```
R1(config)#Int tunnel 1234
R1(config-if)#ip nhrp map multicast 192.1.3.3
R1(config-if)#ip nhrp map multicast 192.1.4.4
```

On R3:

```
R3(config)#int tunnel 1234
R3(config-if)#ip nhrp map multicast 192.1.1.1
```

On R4:

```
R4(config)#int tunnel 1234
R4(config-if)#ip nhrp map multicast 192.1.1.1
```

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           1   FULL/DROTHER    00:00:36    10.1.1.2        Tunnel1234
0.0.0.3           1   FULL/DROTHER    00:00:33    10.1.1.3        Tunnel1234
0.0.0.4           1   FULL/DR         00:00:35    10.1.1.4        Tunnel1234
```

We can see that R4 which happens to be one of the spokes is the DR, we should always configure the hub router as the DR, let's configure this and verify:

On R2, R3, and R4:

```
Rx(config)#Int tunnel 1234
Rx(config-if)#ip ospf priority 0
```

On All Routers:

```
R2#Clear ip ospf process
Reset ALL OSPF processes? [no]: Yes
```

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
0.0.0.2           0   FULL/DROTHER    00:00:31    10.1.1.2        Tunnel1234
0.0.0.3           0   FULL/DROTHER    00:00:32    10.1.1.3        Tunnel1234
0.0.0.4           0   FULL/DROTHER    00:00:32    10.1.1.4        Tunnel1234

R1#Show ip route ospf | B Gate
Gateway of last resort is not set

O     2.0.0.0/8 [110/1001] via 10.1.1.2, 00:01:58, Tunnel1234
O     3.0.0.0/8 [110/1001] via 10.1.1.3, 00:01:58, Tunnel1234
O     4.0.0.0/8 [110/1001] via 10.1.1.4, 00:01:58, Tunnel1234
```

We know that on Broadcast network types the next-hop IP address is set based on the originating router, from R1's perspective this is not a problem, but let's check R2 or the other spokes:

On R2:

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

O     1.0.0.0/8 [110/1001] via 10.1.1.1, 00:03:18, Tunnel1234
O     3.0.0.0/8 [110/1001] via 10.1.1.3, 00:03:08, Tunnel1234
O     4.0.0.0/8 [110/1001] via 10.1.1.4, 00:03:18, Tunnel1234
```

Does R2 have reachability to the advertised networks? Let's verify:

```
R2#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

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 4.4.4.4

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

R2 has reachability to 1.1.1.1 prefix only, does R2 have reachability to the next-hop IP address of the other spoke routers? Let's verify:

```
R2#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2#Ping 10.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

No, it doesn't.
NBMA networks can be configured in two ways: Point-to-point and/or Multipoint. If the network is configured in a Point-to-point manner, the routers have Unicast, Multicast and/or Broadcast capability, and since there can only be another node/router on the other end of the link/tunnel, as long as the destination network is in the routing table we should be able to reach the destination. Whereas, if the network is configured as Multipoint, we have unicast reachability but Broadcast/Multicast capability is only available if it's provided.
Since in a multipoint networks there can potentially be more than one router on the other end of the tunnel, the local router MUST have a mapping to the next-hop IP address/es or else, NLRI can not be achieved.
Since the spoke routers don't have mapping for each other's tunnel IP address, they can not reach the advertised networks, let's provide this reachability and verify:

On R2:

```
R2(config)#Int tunnel 1234
R2(config-if)#ip nhrp map 10.1.1.3 192.1.1.1
```

On R3:

```
R3(config)#Int tunnel 1234
R3(config-if)#ip nhrp map 10.1.1.2 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R2:

```
R2#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 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

This is great, let's configure full mesh logical mapping between the spoke routers:

On R2:

```
R2(config)#Int tunnel 1234
R2(config-if)#ip nhrp map 10.1.1.4 192.1.1.1
```

On R3:

```
R3(config)#Int tunnel 1234
R3(config-if)#ip nhrp map 10.1.1.4 192.1.1.1
```

On R4:

```
R4(config)#Int tunnel 1234
R4(config-if)#ip nhrp map 10.1.1.2 192.1.1.1
R4(config-if)#ip nhrp map 10.1.1.3 192.1.1.1
```

To verify the mappings:

On R2:

```
R2#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 00:50:36, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.3/32 via 10.1.1.3
   Tunnel1234 created 00:04:57, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
10.1.1.4/32 via 10.1.1.4
   Tunnel1234 created 00:01:41, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

On R3:

```
R3#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 00:50:50, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1234 created 00:05:11, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
10.1.1.4/32 via 10.1.1.4
   Tunnel1234 created 00:02:28, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

On R4:

```
R4#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 00:50:54, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1234 created 00:01:33, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
10.1.1.3/32 via 10.1.1.3
   Tunnel1234 created 00:01:25, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

To test the configuration:

On R2:

```
R2#Ping 3.3.3.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Traceroute 3.3.3.3 numeric

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.1 4 msec 4 msec 0 msec
  2 10.1.1.3 0 msec *  0 msec
```

On R3:

```
R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Since in broadcast network types the next-hop IP address is not changed, and if the network is NBMA in nature and it's configured in Multipoint manner, we have to remember the following:

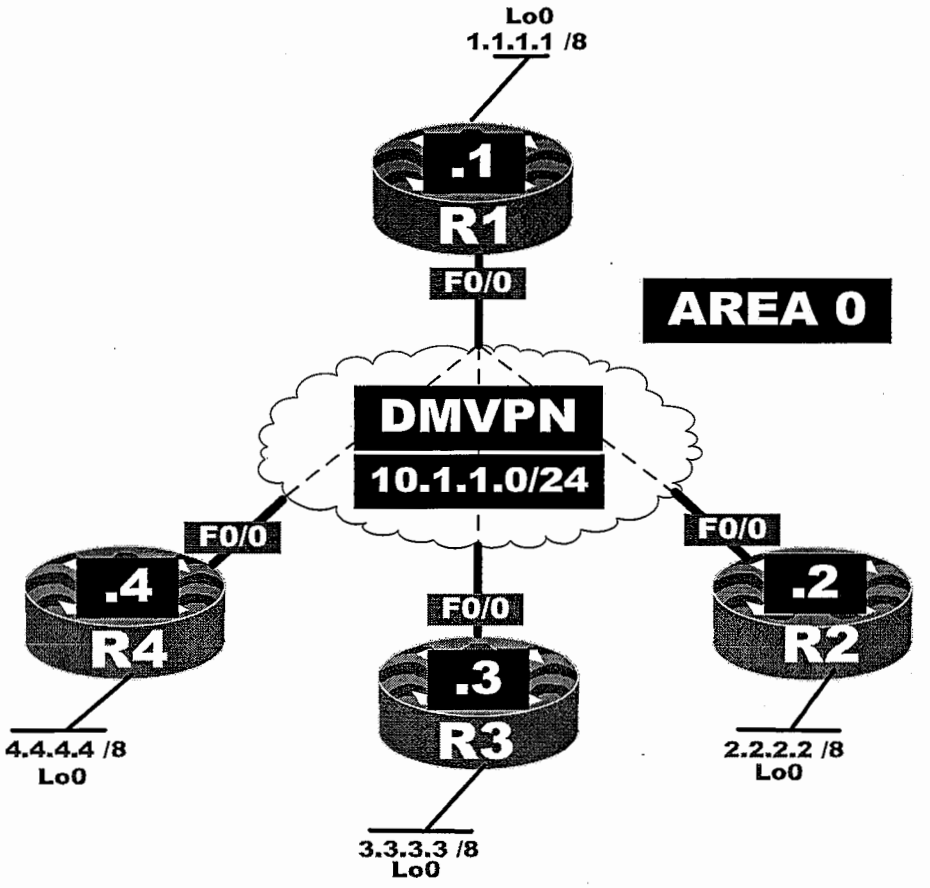
- **Multicast capability must be provided.**
- **Spokes must have mapping to the next-hop IP address to have reachability to the networks that other spoke/s are advertising.**

Task 3

Erase the startup configuration of the routers, the config.text and the VLAN.dat of the switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 4 – OSPF Non-Broadcast Networks



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "OSPF"→ "Lab-4".

Task 1

Configure OSPF area 0 on all routers and their directly connected interface/s. You should configure the tunnel interfaces as OSPF Non-broadcast Network Type. Use the following table for router IDs:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R1 – 0.0.0.1	R2 – 0.0.0.2
R3 – 0.0.0.3	R4 – 0.0.0.4

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Network 10.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 1.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Network 10.1.1.2 0.0.0.0 area 0
R2(config-router)#Network 2.2.2.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Network 3.3.3.3 0.0.0.0 area 0
R3(config-router)#Network 10.1.1.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#Network 4.4.4.4 0.0.0.0 area 0
R4(config-router)#Network 10.1.1.4 0.0.0.0 area 0
```

On All Routers:

```
Rx(config)#int tunnel 1234
Rx(config-if)#IP ospf network non-broadcast
```

To verify the configuration:

On R1:

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


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Show ip ospf int br

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               1.1.1.1/8          1     LOOP  0/0
Tu1234       1     0               10.1.1.1/24        1000  DR    0/0
```

Let's check OSPF's configuration on the tunnel 1234 interface:

```
R1#Show ip ospf int tunnel 1234

Tunnel1234 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 0.0.0.1, Network Type NON_BROADCAST, Cost: 1000
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           1000      no          no            Base
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 0.0.0.1, Interface address 10.1.1.1
  No backup designated router on this network
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:14
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

We can see that the network type is set to "Non-Broadcast" this means that one way to get the local router to establish an adjacency with another OSPF speaking router is to use the "Neighbor" command and change the Multicast destination of 224.0.0.5 to unicast destinations of 10.1.1.2, 10.1.1.3, and 10.1.1.4. So we can see that the rules did not change at all, the same rules apply to all NBMA networks.

Let's configure the "Neighbor" commands on the routers:

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#Neighbor 10.1.1.2
R1(config-router)#Neighbor 10.1.1.3
R1(config-router)#Neighbor 10.1.1.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

You should see the following console messages stating that the adjacencies are established with OSPF 0.0.0.2, 0.0.0.3, and 0.0.0.4 neighbors.

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Tunnel1234 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Tunnel1234 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Tunnel1234 from LOADING to FULL,
Loading Done
```

This worked but let's review some important points about OSPF Non-Broadcast Networks:

- Neighbor command must be used because Non-broadcast network types don't have Multicast capability.
- The Hello intervals are 30 seconds, and the dead interval is set to 120 seconds.
- Requires "DR and "BDR" election.
- In this network types the next-hop IP address is based on the router that originated the route.

Let's go through every item and verify:

If the routers don't have Multicast capability, then the "Neighbor" command must be configured, in a hub and spoke networks, the hub MUST be the "DR", and the spokes are configured with a priority of "0" so they don't participate in the "DR/BDR" election.

On R2:

```
R2#Show ip ospf int tunn 1234 | Inc Priority

  Transmit Delay is 1 sec, State WAITING, Priority 1
```

On R3:

```
R3#Show ip ospf int tunn 1234 | Inc Priority

  Transmit Delay is 1 sec, State BDR, Priority 1
```

On R4:

```
R4#Show ip ospf int tunn 1234 | Inc Priority

  Transmit Delay is 1 sec, State BDR, Priority 1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Let's change the priority of the tunnel interfaces on all spokes to zero and clear the OSPF process to implement the changes.

On R2, R3, and R4:

```
Rx(config)#Int tunnel 1234
Rx(config-if)#IP ospf priority 0

R2#Clear ip ospf process
Reset ALL OSPF processes? [no]: Yes
```

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           0   FULL/DROTHER    00:01:46    10.1.1.2        Tunnel123
0.0.0.3           0   FULL/DROTHER    00:01:46    10.1.1.3        Tunnel1234
0.0.0.4           0   FULL/DROTHER    00:01:46    10.1.1.4        Tunnel1234
```

On All Routers:

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        2.2.2.2 [110/1001] via 10.1.1.2, 00:03:13, Tunnel1234
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/1001] via 10.1.1.3, 00:03:13, Tunnel1234
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/1001] via 10.1.1.4, 00:03:13, Tunnel1234
```

Let's check the routing table of the spokes:

On R2:

```
R2#Show ip route ospf | 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
O        1.1.1.1 [110/1001] via 10.1.1.1, 00:04:16, Tunnel1234
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/1001] via 10.1.1.3, 00:04:06, Tunnel1234
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/1001] via 10.1.1.4, 00:04:06, Tunnel1234
```

On R3:

```
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 [110/1001] via 10.1.1.1, 00:05:03, Tunnel1234
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/1001] via 10.1.1.2, 00:05:03, Tunnel1234
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/1001] via 10.1.1.4, 00:04:53, Tunnel1234
```

On R4:

```
R4#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 [110/1001] via 10.1.1.1, 00:05:12, Tunnel1234
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/1001] via 10.1.1.2, 00:05:12, Tunnel1234
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/1001] via 10.1.1.3, 00:05:02, Tunnel1234
```

As we can see in OSPF "Non-Broadcast" network types the next hop is based on the router's IP address that originated the route. Do we have NLRI to those addresses? Let's verify:

```
R4#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R4#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:
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

R4#Ping 10.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

As we can see the spokes do not have reachability to the routes that are advertised by the other spokes, let's check their NHRP mapping (Does that remind you of frame-relay?):

```
R4#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 01:13:32, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
```

The only mapping is for the hub router and that's why the ping to 1.1.1.1 prefix was successful. On R4, we should configure a mapping for R2 and another one for R3, and for the return traffic, a mapping from R2 and R3 to R4 is required. Let's configure and verify:

On R4:

```
R4(config)#Int tunnel 1234
R4(config-if)#ip nhrp map 10.1.1.2 192.1.1.1
R4(config-if)#ip nhrp map 10.1.1.3 192.1.1.1
```

NOTE: R4 is mapping the next-hop IP address (The tunnel IP address) to the NBMA-IP address of the hub so the topology does not change.

```
R4#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 00:48:05, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1234 created 00:01:34, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
10.1.1.3/32 via 10.1.1.3
   Tunnel1234 created 00:01:27, never expire
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

On R2:

```
R2(config)#int tunnel 1234
R2(config-if)#ip nhrp map 10.1.1.3 192.1.1.1
R2(config-if)#ip nhrp map 10.1.1.4 192.1.1.1
```

To verify the configuration:

```
R2#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 00:50:07, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.3/32 via 10.1.1.3
   Tunnel1234 created 00:00:31, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
10.1.1.4/32 via 10.1.1.4
   Tunnel1234 created 00:00:24, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

On R3:

```
R3(config)#int tunnel 1234
R3(config-if)#ip nhrp map 10.1.1.2 192.1.1.1
R3(config-if)#ip nhrp map 10.1.1.4 192.1.1.1
```

To verify the configuration:

```
R3#Show ip nhrp

10.1.1.1/32 via 10.1.1.1
   Tunnel1234 created 00:51:19, never expire
   Type: static, Flags: used
   NBMA address: 192.1.1.1
10.1.1.2/32 via 10.1.1.2
   Tunnel1234 created 00:00:28, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
10.1.1.4/32 via 10.1.1.4
   Tunnel1234 created 00:00:21, never expire
   Type: static, Flags:
   NBMA address: 192.1.1.1
```

Let's very the routing table and NLRI on the spoke routers:

On R2:

```
R2#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 [110/1001] via 10.1.1.1, 00:22:19, Tunnel1234
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/1001] via 10.1.1.3, 00:22:09, Tunnel1234
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/1001] via 10.1.1.4, 00:22:09, Tunnel1234

R2#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

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R2#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R3:

```
R3#Show ip route ospf | 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
O        1.1.1.1 [110/1001] via 10.1.1.1, 00:32:39, Tunnel1234
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/1001] via 10.1.1.2, 00:32:39, Tunnel1234
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/1001] via 10.1.1.4, 00:32:29, Tunnel1234

R3#Ping 10.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

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 = 1/3/4 ms

R3#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

On R4:

```
R4#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 [110/1001] via 10.1.1.1, 00:33:16, Tunnel1234
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/1001] via 10.1.1.2, 00:33:16, Tunnel1234
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/1001] via 10.1.1.3, 00:33:06, Tunnel1234

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:
!!!!!
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R4#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 = 1/3/4 ms

R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Let's sum up the important points about non-broadcast networks:

- **It requires DR/BDR.**
- **No Broadcast/Multicast capabilities, ONLY Unicast is available.**
- **Next-hop IP address does NOT change.**
- **Hello Intervals are 30 seconds and the dead interval is set to 120 seconds.**

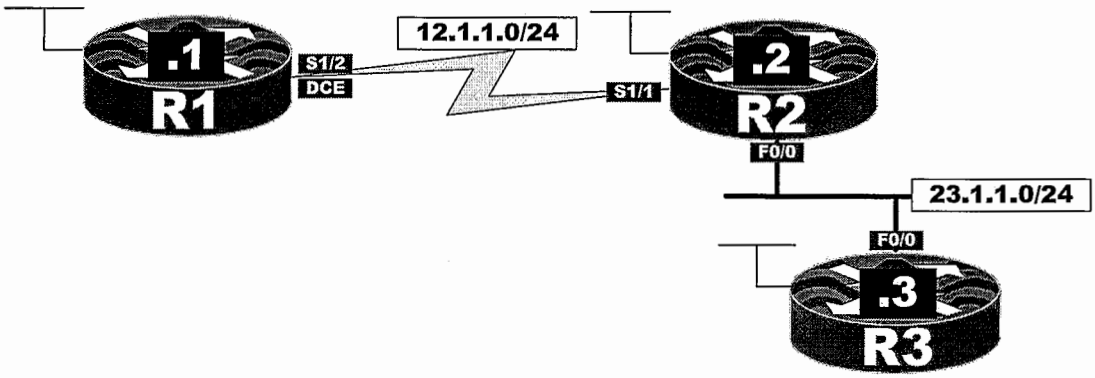
Task 2

Erase the startup configuration of the routers, the config.text and the VLAN.dat of the switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 5 – OSPF Point-to-Point Networks



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "OSPF"→ "Lab-5".

Task 1

Configure OSPF on all routers and run their directly connected interfaces in Area 0 based on the following policy:

- The loopback 0 interface of these routers should be advertised with their correct mask
- Use 0.0.0.1, 0.0.0.2, and 0.0.0.3 as the router-id of R1, R2 and R3 respectively.
- There should NOT be any DR/BDR election on any of the links.
- DO NOT configure "Point-to-Multipoint" or "Point-to-Multipoint Non-broadcast" on any of the links.

On R1:

```
R1(config-if)#router ospf 1
R1(config-router)#router-id 0.0.0.1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 12.1.1.1 0.0.0.0 area 0

R1(config-router)#int lo0
R1(config-if)#ip ospf network point-to-point
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 12.1.1.2 0.0.0.0 area 0
R2(config-router)#netw 23.1.1.2 0.0.0.0 area 0

R2(config-router)#int lo0
R2(config-if)#ip ospf network point-to-point
```

On R3:

```
R3(config)#router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 23.1.1.3 0.0.0.0 area 0

R3(config-router)#int lo0
R3(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

O     2.0.0.0/8 [110/782] via 12.1.1.2, 00:00:17, Serial1/2
O     3.0.0.0/8 [110/783] via 12.1.1.2, 00:00:17, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/782] via 12.1.1.2, 00:00:17, Serial1/2
```

On R2:

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


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O     1.0.0.0/8 [110/782] via 12.1.1.1, 00:01:10, Serial1/1
O     3.0.0.0/8 [110/2] via 23.1.1.3, 00:02:02, FastEthernet0/0
```

On R3:

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

O     1.0.0.0/8 [110/783] via 23.1.1.2, 00:01:39, FastEthernet0/0
O     2.0.0.0/8 [110/2] via 23.1.1.2, 00:02:41, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/782] via 23.1.1.2, 00:01:49, FastEthernet0/0

R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           1   FULL/BDR        00:00:33    23.1.1.2        FastEthernet0/0
```

The task states that there should not be any DR/BDR election, therefore, the network type of the F0/0 interface on R2 and R3 should be changed to "Point-to-Point".

On R2 and R3:

```
Rx(config)#int f0/0
Rx(config-if)#ip ospf network point-to-point
```

To verify the configuration:

On R3:

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

O     1.0.0.0/8 [110/783] via 23.1.1.2, 00:00:09, FastEthernet0/0
O     2.0.0.0/8 [110/2] via 23.1.1.2, 00:00:09, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/782] via 23.1.1.2, 00:00:09, FastEthernet0/0
```

Let's verify some of the parameters on the interfaces that have point-to-point network type:

On R1:

```
R1#Show ip ospf int s1/2 | Inc Network|Hello
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Internet Address 12.1.1.1/24, Area 0, Attached via Network Statement
Process ID 1, Router ID 0.0.0.1, Network Type POINT_TO_POINT, Cost: 781
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:01
```

We see that on OSPF Point-to-point network types there are no DR/BDR election, and the hello and dead interval is set to 10 and 40 respectively. We also saw that we did not experience any problems forming adjacency using Multicast.

The important points in OSPF Point-To-Point network types:

- **Hellos are exchanged every 10 seconds, and the dead interval is set to 40 seconds.**
- **No DR/BDR election.**
- **The routers use 224.0.0.5 to form an adjacency and send Hello messages.**
- **The next hop is the router that advertised the route, this is the neighboring router.**

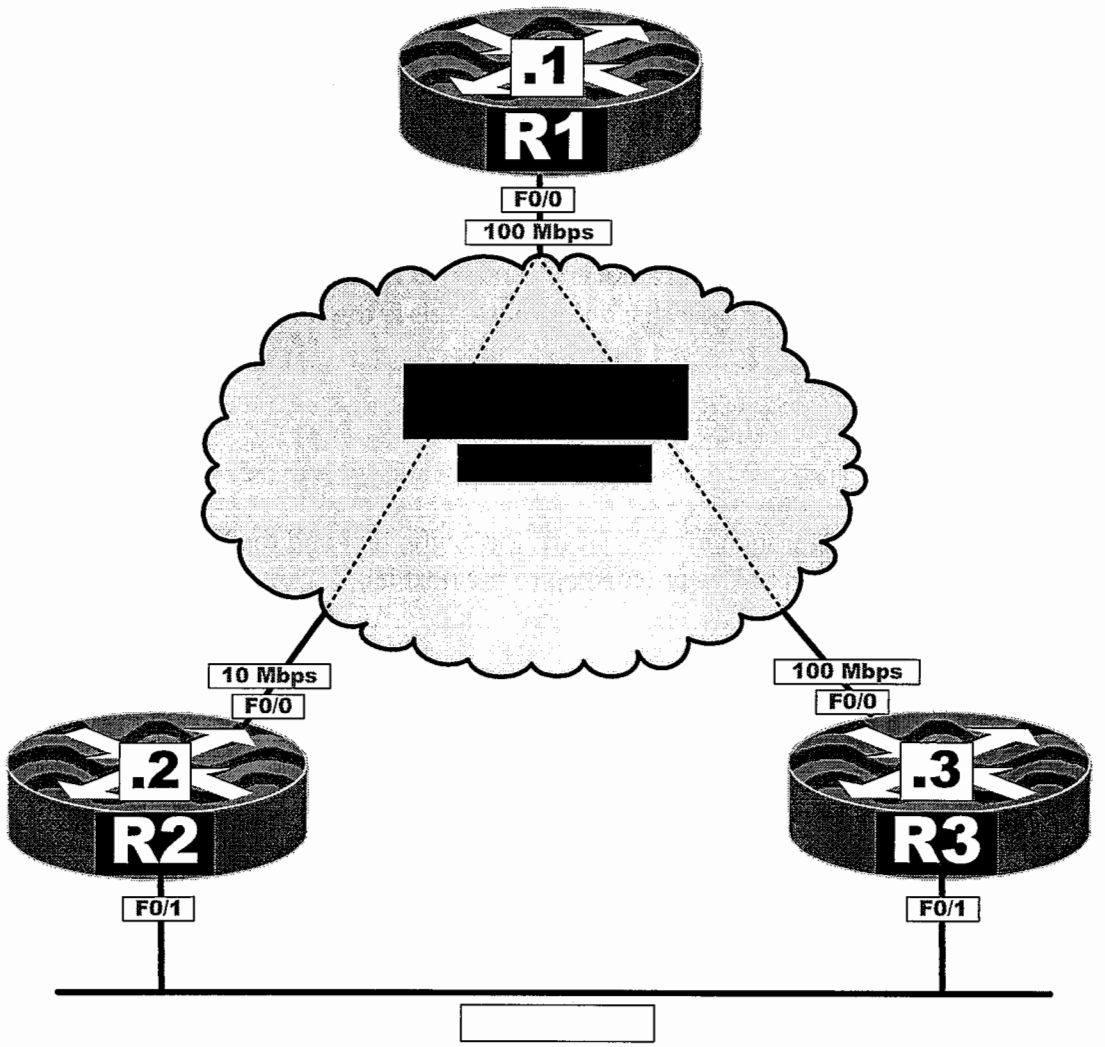
Task 2

Erase the startup configuration of the routers, the config.text and the VLAN.dat of the switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 6 – OSPF "Point-to-Multipoint" & "Point-to-Multipoint Non-Broadcast" Networks



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→ "Advanced-init"→ "OSPF"→ "Lab-6".

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 1

Configure OSPF area 0 on all links in the above topology. If this configuration is performed successfully, the routers in this topology should have full NLRI to every network in this topology. The tunnel interface of these routers should be configured as Point-to-Multipoint OSPF Network type.

On All routers:

```
Rx(config)#int tunnel 123
Rx(config-if)#IP ospf network point-to-multipoint
```

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Network 10.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Network 10.1.1.2 0.0.0.0 area 0
R2(config-router)#Network 23.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Network 10.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 23.1.1.3 0.0.0.0 area 0
```

To verify the configuration:

On R1:

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

The routers did not form an OSPF adjacency, let's look at the important points in OSPF Point-to-Multipoint Network types:

- **The routers use 224.0.0.5**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

- No requirement for DR/BDR
- Next-hop is the advertising router
- Host routes are advertised automatically for NLRI in partial mesh topologies.

Let's see if the tunnel interfaces have Multicast capability:

On R1:

```
R1#Show run int tunnel 123 | B interface

interface Tunnel123
 ip address 10.1.1.1 255.255.255.0
 no ip redirects
 ip nhrp map 10.1.1.3 192.1.3.3
 ip nhrp map 10.1.1.2 192.1.2.2
 ip nhrp network-id 111
 ip ospf network point-to-multipoint
 tunnel source FastEthernet0/0
 tunnel mode gre multipoint
end
```

On R2:

```
R2#Show run int tunnel 123 | B interface

interface Tunnel123
 ip address 10.1.1.2 255.255.255.0
 no ip redirects
 ip nhrp map 10.1.1.1 192.1.1.1
 ip nhrp network-id 222
 ip ospf network point-to-multipoint
 tunnel source FastEthernet0/0
 tunnel mode gre multipoint
end
```

On R3:

```
R3#Show run int tunnel 123 | B inter

interface Tunnel123
 ip address 10.1.1.3 255.255.255.0
 no ip redirects
 ip nhrp map 10.1.1.1 192.1.1.1
 ip nhrp network-id 333
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
 ip ospf network point-to-multipoint
 tunnel source FastEthernet0/0
 tunnel mode gre multipoint
end
```

We can see the problem, the routers don't have mapping for Multicast, let's configure them and verify:

On R1:

```
R1(config)#int tunnel 123
R1(config-if)#ip nhrp map Multicast 192.1.2.2
R1(config-if)#ip nhrp map Multicast 192.1.3.3
```

On R2:

```
R2(config)#int tunnel 123
R2(config-if)#ip nhrp map Multicast 192.1.1.1
```

On R3:

```
R3(config)#int tunnel 123
R3(config-if)#ip nhrp map Multicast 192.1.1.1
```

To verify the configuration:

On R1:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           0   FULL/  -        00:01:57    10.1.1.2        Tunnel123
0.0.0.3           0   FULL/  -        00:01:57    10.1.1.3        Tunnel123

R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        10.1.1.2/32 [110/1000] via 10.1.1.2, 00:02:06, Tunnel123
O        10.1.1.3/32 [110/1000] via 10.1.1.3, 00:02:06, Tunnel123
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/1001] via 10.1.1.3, 00:02:06, Tunnel123
                  [110/1001] via 10.1.1.2, 00:02:06, Tunnel123
```

Let's summarize some of the important points in OSPF Point-to-Multipoint Network types:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

1. The routers use 224.0.0.5
2. No requirement for DR/BDR
3. Next-hop is the advertising router
4. Host routes are advertised automatically for NLRI in partial mesh topologies.

We have fixed the first point by mapping Multicast.
Let's verify the second point:

```
R1#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.2           0   FULL/  -        00:01:54    10.1.1.2        Tunnel123
0.0.0.3           0   FULL/  -        00:01:58    10.1.1.3        Tunnel123
```

NOTE: The routers are in full state and there are no DR/BDR.

Let's verify the third and the fourth point:

On R2:

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

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        10.1.1.1/32 [110/1000] via 10.1.1.1, 00:06:48, Tunnel123
O        10.1.1.3/32 [110/1] via 23.1.1.3, 00:20:03, FastEthernet0/1
```

We can see that the nexthop IP address is the IP address of the router that originated the route, we can also see the host routes that are advertised for reachability.

Task 2

Since R2's connection to the cloud is 10 Mbps and R3's connection is 100, R1 should NOT perform equal cost load sharing, R1 should go through R3 to reach network 23.1.1.0/24. DO NOT configure PBR, the "IP ospf cost" command to accomplish this task.

As we can see R2's F0/0 is configured with a 10 Mbps, whereas, the F0/0 interfaces of R1 and R3 re configured with 100 Mbps. Since OSPF is running on the tunnel interfaces of these routers OSPF will not see the actual cost of the F0/0 interfaces therefore, a suboptimnal routing can result.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

In order to accomplish this task, the OSPF network type can be changed, if the network type is changed to "Point-to-Multipoint Non-broadcast", neighbor commands MUST be configured on the hub router, in this network type the neighbor command can be configured with a cost which can make the connection through a given neighbor more attractive. Let's configure and verify. Before this task is configured we should verify the routing table of R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        10.1.1.2/32 [110/1000] via 10.1.1.2, 00:02:06, Tunnel123
O        10.1.1.3/32 [110/1000] via 10.1.1.3, 00:02:06, Tunnel123
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/1001] via 10.1.1.3, 00:02:06, Tunnel123
                  [110/1001] via 10.1.1.2, 00:02:06, Tunnel123
```

On All Routers:

```
Rx(config)#Int tunnel 123
Rx(config-if)#IP ospf network point-to-multipoint non-broadcast

R1(config)#Router ospf 1
R1(config-router)#Neighbor 10.1.1.2 cost 20
R1(config-router)#Neighbor 10.1.1.3 cost 1
```

You should see the following console messages stating that the neighbor adjacency has been established:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Tunnel123 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Tunnel123 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

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O        10.1.1.2/32 [110/2] via 10.1.1.3, 00:02:55, Tunnel123
O        10.1.1.3/32 [110/1] via 10.1.1.3, 00:02:55, Tunnel123
      23.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O        23.1.1.0 [110/2] via 10.1.1.3, 00:02:55, Tunnel123
```

Let's summarize the important points in OSPF point-to-multipoint network types:

- **No DR/BDR requirement**
- **Hello and dead interval is 30/120**
- **Nexthop is the neighboring router**
- **No Multicast capability, therefore, neighbor commands must be configured to establish OSPF adjacencies.**

Task 3

Erase the startup configuration of the routers, the config.text and the VLAN.dat of the switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Advanced
CCIE Routing & Switching
v5.0

www.MicronicsTraining.com

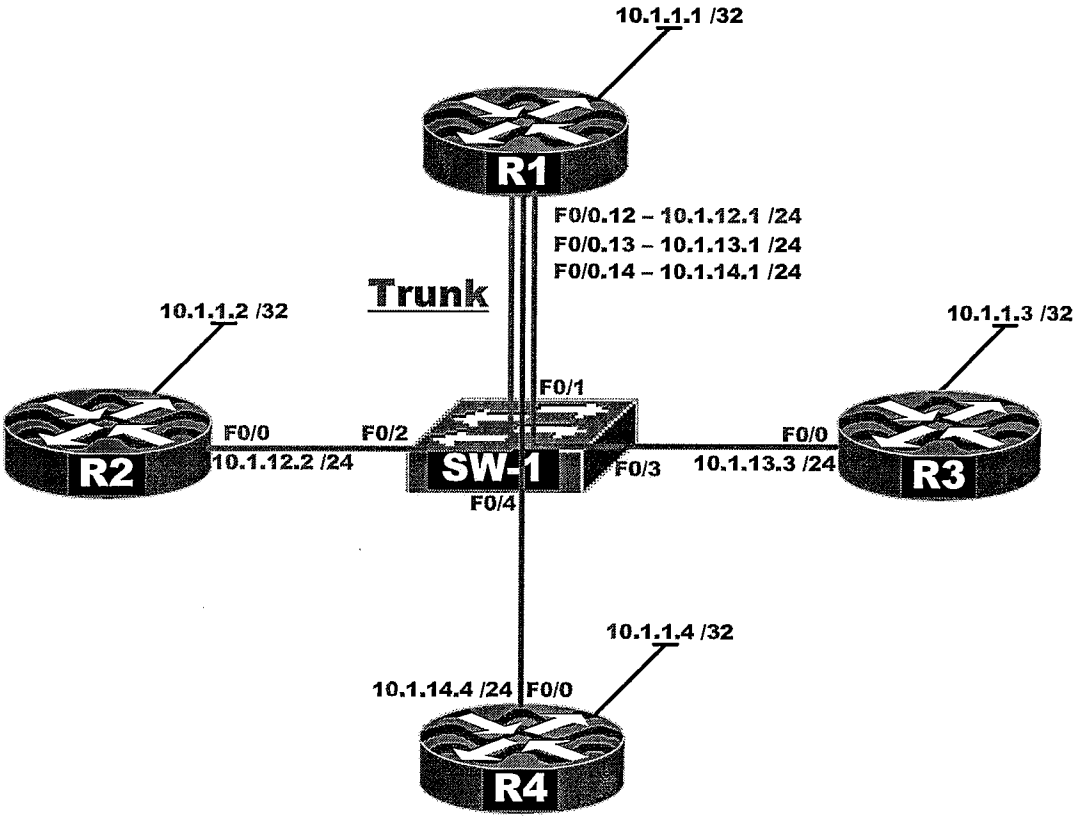
Narbik Kocharians
CCSI, CCIE #12410
R&S, Security, SP

Redistribution

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1 – Basics of Redistribution



Task 1

Configure the routers and the switches based on the above diagram.

On SW1

```
SW1(config)#Int F0/1
SW1(config-if)#Swi trunk encap dot1q
SW1(config-if)#Swi mode trunk
SW1(config-if)#No shut

SW1(config)#Int F0/2
SW1(config-if)#Swi mode access
SW1(config-if)#Swi access vlan 12
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW1(config-if)#No shut

SW1(config)#Int F0/3
SW1(config-if)#Swi mode access
SW1(config-if)#Swi access vlan 13
SW1(config-if)#No shut

SW1(config)#Int F0/4
SW1(config-if)#Swi mode access
SW1(config-if)#Swi access vlan 14
SW1(config-if)#No shut
```

On R1

```
R1(config)#Int F0/0
R1(config-if)#No shut

R1(config-if)#Int F0/0.12
R1(config-subif)#Encap dot1q 12
R1(config-subif)#IP address 10.1.12.1 255.255.255.0

R1(config)#Int F0/0.13
R1(config-subif)#Encap dot1q 13
R1(config-subif)#IP address 10.1.13.1 255.255.255.0

R1(config)#Int F0/0.14
R1(config-subif)#Encap dot1q 14
R1(config-subif)#IP address 10.1.14.1 255.255.255.0

R1(config)#Int Lo0
R1(config-if)#IP address 10.1.1.1 255.255.255.255
```

On R2

```
R2(config)#Int F0/0
R2(config-if)#IP address 10.1.12.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int Lo0
R2(config-if)#IP address 10.1.1.2 255.255.255.255
```

On R3

```
R3(config)#Int F0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-if)#IP address 10.1.13.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Int Lo0
R3(config-if)#IP address 10.1.1.3 255.255.255.255
```

On R4

```
R4(config)#Int F0/0
R4(config-if)#IP address 10.1.14.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int Lo0
R4(config-if)#IP address 10.1.1.4 255.255.255.255
```

To verify the configuration:

On R1

```
R1#Ping 10.1.12.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 10.1.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Ping 10.1.14.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 2

Configure the link between R1 and R3 to be in OSPF area 0. Configure Lo0 interface of R3 in area 2. Do not use the network command to accomplish this task. The OSPF router IDs should be configured as 0.0.0.1 and 0.0.0.3 for R1 and R3 respectively.

On R1

```
R1(config)#int f0/0.13
R1(config-subif)#ip ospf 1 area 0

R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
```

On R3

```
R3(config)#int f0/0
R3(config-if)#ip ospf 1 area 0

R3(config)#int lo0
R3(config-if)#ip ospf 1 area 2
```

When running OSPF on a given interface by using the interface configuration command, the IOS starts the OSPF process automatically:

```
R3#Sh run | s router ospf 1

router ospf 1
 log-adjacency-changes

R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
```

To verify the configuration:

R1 should see R3's Lo0 as "O IA" route in its routing table

On R1

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O IA      10.1.1.3/32 [110/2] via 10.1.13.3, 00:01:33, FastEthernet0/0.13
```

Task 3

Add the following Loopback interfaces to R3

- Lo31: 192.168.31.3/24
- Lo32: 192.168.32.3/24
- Lo33: 192.168.33.3/24

Redistribute Lo31 **and** Lo33. Do not redistribute Lo32 into OSPF. Do not use "access-list" or "prefix-list" to accomplish the task.

The first step in resolving this task is to configure the loopback interfaces on R3, as follows:

On R3

```
R3(config)#int lo31
R3(config-if)#ip addr 192.168.31.3 255.255.255.0

R3(config-if)#int lo32
R3(config-if)#ip addr 192.168.32.3 255.255.255.0

R3(config-if)#int lo33
R3(config-if)#ip addr 192.168.33.3 255.255.255.0
```

Since using "access-lists" or "prefix-lists" is not allowed, a route-map is configured and the required interfaces are matched using the "match interface" option in the route-map configuration mode.

NOTE: The task states that Lo32 should not be redistributed, therefore, the route-map could be configured to deny Lo32 and permit the rest of the networks, or the route-map could simply permit the Lo31 and Lo33 interfaces and deny Lo32 from being redistributed.

Once the route-map is configured, it's referenced in the "redistribute connected" command in the router configuration mode.

Let's see our options as far as configuring a route-map:

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Option 1:

In the following option only "Lo32" is denied, if "route-map TST permit 90" is not configured, the rest of the interfaces will also be denied from being redistributed, because if the logic of the route-map is not completed, the IOS will try to protect you and in that process it will drop the rest of the interfaces/networks. Therefore, the "route-map TST permit 90" is configured to permit the rest of the interfaces.

```
R3(config)#route-map TST deny 10
R3(config-route-map)#match interface Lo32
R3(config)#route-map TST permit 90
```

Option 2:

The following option is more specific and it should be used in this configuration. This option is more specific because it will not redistribute future directly connected routes in this process.

```
R3(config)#route-map TST permit 10
R3(config-route-map)#match interface Lo31 Lo33
```

In the final step, the connected interfaces that are referenced in the "route-map" are redistributed into OSPF routing protocol.

```
R3(config-route-map)#router ospf 1
R3(config-router)#redistribute connected subnets route-map TST
```

The "subnets" keyword is required or else only the classful networks (Networks that are NOT subnetted) will be redistributed.

To verify the configuration:

On R3:

```
R3#Show ip ospf da external | Inc 192

                Type-5 AS External Link States
  Link State ID: 192.168.31.0 (External Network Number )
  Link State ID: 192.168.33.0 (External Network Number )
```

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
O IA     10.1.1.3/32 [110/2] via 10.1.13.3, 00:10:39, FastEthernet0/0.13
O E2 192.168.31.0/24 [110/20] via 10.1.13.3, 00:01:45, FastEthernet0/0.13
O E2 192.168.33.0/24 [110/20] via 10.1.13.3, 00:01:45, FastEthernet0/0.13
```

Task 4

Configure R1 and R2 based on the following policy:

- Configure RIPv2 between R1 and R2. Disable auto-summary.
- R1 should be configured to redistribute OSPF routes into the RIP routing domain.
- R1 should filter R3's Lo33; Do not use "**distribute-list**" or "**offset-list**" to accomplish this task.

The first step is to configure RIPv2 on both R1 and R2:

On R1

```
R1(config)#Router rip
R1(config-router)#version 2
R1(config-router)#No auto-summary
R1(config-router)#Network 10.0.0.0
R1(config-router)#Passive-interface F0/0.13
R1(config-router)#Passive-interface F0/0.14
```

Note: In the above configuration the "Passive-interface" commands are required because the task requires RIPv2 to run between R1 and R2; the "passive-interface" commands turn the specified sub-interfaces into receive-only mode.

On R2:

```
R2(config)#Router rip
R2(config-router)#version 2
R2(config-router)#No auto-summary
R2(config-router)#Net 10.0.0.0
```

To verify the configuration:

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

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

      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R        10.1.13.0/24 [120/1] via 10.1.12.1, 00:00:20, FastEthernet0/0
R        10.1.14.0/24 [120/1] via 10.1.12.1, 00:00:20, FastEthernet0/0
```

In the second step OSPF is redistributed into RIP on R1:

On R1:

```
R1(config)#Router rip
R1(config-router)#redistribute ospf 1
```

To verify the configuration:

On R2:

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

      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
R        10.1.13.0/24 [120/1] via 10.1.12.1, 00:00:12, FastEthernet0/0
R        10.1.14.0/24 [120/1] via 10.1.12.1, 00:00:12, FastEthernet0/0
```

NOTE: Lo0, Lo31 and Lo33 that were redistributed from OSPF routing domain are NOT in the routing table of R2, only R1's interfaces are redistributed.

One of the biggest problems of redistribution is that each routing protocol has its own metric:

- **RIP : hop count**
- **EIGRP : composite of bandwidth, delay, reliability, load and MTU**
- **OSPF : Cost which is based on bandwidth**

When redistributing from OSPF into RIP routing protocol, what should be the metric?
Well, there are many choices, and some of them are as follows:

- **The metric can be configured such that it applies to all existing and future redistributed routes.**
- **The metric can be set separately on each configured redistribute command.**
- **The metric can be set based on usage of a "route-map".**

One notable exception is directly connected routes, which RIP applies a default metric of 0.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To correct the problem, a metric of 3 is assigned to the OSPF routes that are redistributed into RIPv2:

On R1

```
R1(config)#router rip
R1(config-router)#redistribute ospf 1 metric 3
```

To verify the configuration:

On R2

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

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
R        10.1.1.3/32 [120/3] via 10.1.12.1, 00:00:11, FastEthernet0/0
R        10.1.13.0/24 [120/1] via 10.1.12.1, 00:00:22, FastEthernet0/0
R        10.1.14.0/24 [120/1] via 10.1.12.1, 00:00:22, FastEthernet0/0
R     192.168.31.0/24 [120/3] via 10.1.12.1, 00:00:11, FastEthernet0/0
R     192.168.33.0/24 [120/3] via 10.1.12.1, 00:00:11, FastEthernet0/0
```

The output of the above show command verifies that all OSPF routes have been redistributed into RIPv2 on R2; the last step of this task requires filtering of R3's Lo33 (192.168.33.0/24). Remember that the use of "distribute-list" or "offset-list" is not allowed, therefore, a route-map is used to set the metric of interface Lo33 to infinity; this will cause R1 to poison that route.

The steps required to configure this task are as follows:

- A "prefix-list" is configured to identify the IP address of R3's Lo33.
- A "route-map" is configured to reference the "prefix-list" from the previous step, and sets the metric to 16 (Infinity).
- The redistribution is reconfigured to reference the route map.

On R1

The first Step:

```
R1(config)#IP Prefix-list Lo33 permit 192.168.33.0/24
```

The second step:

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match ip address prefix-list Lo33
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-route-map)#set metric 16
R1(config-route-map)#route-map TST permit 90
```

The third and the final step:

```
R1(config-route-map)#Router rip
R1(config-router)#redistribute ospf 1 metric 3 route-map TST
```

To verify the configuration:

On R2:

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

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
R        10.1.1.3/32 [120/3] via 10.1.12.1, 00:00:21, FastEthernet0/0
R        10.1.13.0/24 [120/1] via 10.1.12.1, 00:00:21, FastEthernet0/0
R        10.1.14.0/24 [120/1] via 10.1.12.1, 00:00:21, FastEthernet0/0
R     192.168.31.0/24 [120/3] via 10.1.12.1, 00:00:21, FastEthernet0/0
```

To see the routes advertised to R2, let's enable rip debugging on R1 and R2:

On R1:

```
R1#Debug  ip rip

RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0.12 (10.1.12.1)
RIP: build update entries

        10.1.1.0/24 via 0.0.0.0, metric 1, tag 0
        10.1.1.3/32 via 0.0.0.0, metric 3, tag 0
        10.1.13.0/24 via 0.0.0.0, metric 1, tag 0
        10.1.14.0/24 via 0.0.0.0, metric 1, tag 0
        192.168.31.0/24 via 0.0.0.0, metric 3, tag 0
        192.168.33.0/24 via 0.0.0.0, metric 16, tag 0
```

In the output of the above debug we can clearly see that network 192.168.33.0/24 is advertised with a hop count of 16, and the output of the following debug reveals that R2 discards the route because a hop count of 16 is inaccessible in RIP routing protocol.

On R2

```
RIP: received v2 update from 10.1.12.1 on FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
     10.1.1.0/24 via 0.0.0.0 in 1 hops
     10.1.1.3/32 via 0.0.0.0 in 3 hops
     10.1.13.0/24 via 0.0.0.0 in 1 hops
     10.1.14.0/24 via 0.0.0.0 in 1 hops
     192.168.31.0/24 via 0.0.0.0 in 3 hops
     192.168.33.0/24 via 0.0.0.0 in 16 hops  (inaccessible)
```

To disable debug on R1 and R2:

On R1 and R2

```
R1#u all
```

Task 5

Add the following loopback interfaces to R2 and R3:

- R2: Lo23: 192.168.23.2 /24
- R3: Lo23: 192.168.23.3 /24

Advertise R2's Lo23 into RIPv2 routing protocol; R3's Lo23 should be configured in OSPF area 2 and advertised with its correct mask.

On R2:

```
R2(config)#Int Lo23
R2(config-if)#IP address 192.168.23.2 255.255.255.0

R2(config-if)#Router rip
R2(config-router)#Netw 192.168.23.0
```

On R3:

```
R3(config)#Int lo23
R3(config-if)#IP address 192.168.23.3 255.255.255.0
R3(config-if)#IP Ospf network point-to-point

R3(config-if)#Router ospf 1
R3(config-router)#Netw 192.168.23.3 0.0.0.0 area 2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip route | Inc 192.168.23

     192.168.23.0/24 is subnetted, 1 subnets
O IA    192.168.23.0 [110/2] via 10.1.13.3, 00:01:08, FastEthernet0/0.13
```

The route does NOT appear as a RIP route; did RIP advertise the route? If so, what happened?
Let's check R1's RIP database:

```
R1#Sh ip rip database 192.168.23.0 255.255.255.0

192.168.23.0/24    redistributed
    [3] via 10.1.13.3, from 0.0.0.3,
```

Note: The route is in RIP's database, but it is in the database as a "redistributed" route. It is in the database as "redistributed" because of redistribution of OSPF into RIP in task 4.

So where is the update from R2?
Maybe R2 is not sending that update at all? Let's debug RIP's updates and verify:

```
R1#Debug ip rip
RIP protocol debugging is on

RIP: sending v2 update to 224.0.0.9 via FastEthernet0/0.12 (10.1.12.1)
RIP: build update entries
        10.1.1.0/24 via 0.0.0.0, metric 1, tag 0
        10.1.1.3/32 via 0.0.0.0, metric 3, tag 0
        10.1.13.0/24 via 0.0.0.0, metric 1, tag 0
        10.1.14.0/24 via 0.0.0.0, metric 1, tag 0
        192.168.23.0/24 via 0.0.0.0, metric 3, tag 0
        192.168.31.0/24 via 0.0.0.0, metric 3, tag 0
        192.168.33.0/24 via 0.0.0.0, metric 16, tag 0
```

R1 received an update for network 192.168.23.0 /24 from R2, but it does not install it into its RIP database, RIP database contains only redistributed routes and learned RIP routes which are actually installed in the router's routing table.

Since the "administrative distance" of RIP is higher than OSPF's administrative distance, IOS prefers the OSPF route and rejects the RIP route. To verify this fact:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3

```
R3(config)#Int Lo23
R3(config-if)#Shut
```

On R1

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

R     192.168.23.0/24 [120/1] via 10.1.12.2, 00:00:20, FastEthernet0/0.12

R1#Show ip rip database 192.168.23.0 255.255.255.0

192.168.23.0/24
    [1] via 10.1.12.2, 00:00:11, FastEthernet0/0.12
```

The output of the above show commands reveal that network 192.168.23.0 /24 is a RIP route.

To verify and reveal the comparison of administrative distance of RIP versus OSPF, the "Debug ip routing" is enabled on R1 and then the Lo23 interface of R3 is enabled (No shut):

To disable the existing debug:

```
R1#u all
```

To enable debugging for IP routing:

```
R1#Debug ip routing
IP routing debugging is on
```

On R3

```
R3(config)#interface loopback 23
R3(config-if)#No shut
```

Let's check the debug output on R1:

On R1

```
RT: updating ospf 192.168.23.0/24 (0x0):
    via 10.1.13.3 Fa0/0.13
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
RT: closer admin distance for 192.168.23.0, flushing 1 routes
RT: add 192.168.23.0/24 via 10.1.13.3, ospf metric [110/2]
```

The first line of the above output reveals the comparison of the administrative distance of RIP and OSPF, and the "add 192.168.23.0/24 via 10.1.13.3, ospf metric [110/2]" is what is injected in the routing table of this router, to verify that information:

```
R1#Show ip route ospf | Inc 192.168.23.0

O IA 192.168.23.0/24 [110/2] via 10.1.13.3, 00:06:03, FastEthernet0/0.13
```

Let's disable the debugging:

```
R1#u all
```

Task 6

Configure EIGRP AS 100 on R1's F0/0.14 and R4's F0/0 interfaces. R4 should redistribute its Loopback0 interface into this routing domain.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 10.1.14.1 0.0.0.0
```

On R4:

```
R4(config)#Router eigrp 100
R4(config-router)#Network 10.1.14.4 0.0.0.0
```

You should see the following console message:

```
%DUAL-5-NBRCHANGE: IP-EIGRP(0) 100: Neighbor 10.1.14.1 (FastEthernet0/0)
is up: new adjacency
```

To verify the configuration:

On R4:

```
R4#Show ip eigrp neighbors
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.14.1               Fa0/0             11 00:13:27    1   200  0  5
```

The first step:

A route-map is configured and Lo0 interface is referenced:

```
R4(config)#Route-map TST permit 10
R4(config-route-map)#Match interface lo0
```

The second step:

The "redistribute connected" command is configured to reference the route-map, when redistributing routes into EIGRP, the metric should also be configured, the metric is in the following order:

Bandwidth, delay, relay, load and MTU.

The only exception is when redistributing connected, static routes, or another EIGRP routing domain.

Note: EIGRP does not use the "MTU" as a parameter for metric calculation, but still a required value that must be configured.

```
R4(config)#Router eigrp 100
R4(config-router)#Redistribute connected route-map TST
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

   10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
D EX 10.1.1.4/32 [170/156160] via 10.1.14.4, 00:02:46, FastEthernet0/0.14
```

Task 7

Redistribute RIP into EIGRP. Ensure that R4 installs a route for network 192.168.23.0/24 into its routing table.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Redistribute rip metric 1 1 1 1 1
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D EX     10.1.1.1/32 [170/2560002816] via 10.1.14.1, 00:00:32, FastEthernet0/0
D EX     10.1.1.2/32 [170/2560002816] via 10.1.14.1, 00:00:32, FastEthernet0/0
D EX     10.1.12.0/24
           [170/2560002816] via 10.1.14.1, 00:00:32, FastEthernet0/0
D EX     10.1.13.0/24
           [170/2560002816] via 10.1.14.1, 00:00:32, FastEthernet0/0
```

Network 192.168.23.0 /24 is NOT redistributed.

```
R4#Show ip route 192.168.23.0
% Network not in table
```

********** Important to remember **********

IOS will redistribute the 192.168.23.0 /24 network only if this network is in the routing table of R2 as a RIP route. Let's say the routes from routing protocol A are to be redistributed into routing protocol B; the IOS will ONLY redistribute the routes that are in its routing table as A into routing protocol B's routing table.

****

To verify the routing table of R1 for network 192.168.23.0 /24:

On R1

```
R1#Show ip rip database 192.168.23.0 255.255.255.0

192.168.23.0/24    redistributed
    [3] via 10.1.13.3, from 0.0.0.3,

R1#Show ip route 192.168.23.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Routing entry for 192.168.23.0/24
  Known via "ospf 1", distance 110, metric 2, type inter area
  Redistributing via rip
  Advertised by rip metric 3 route-map TST
  Last update from 10.1.13.3 on FastEthernet0/0.13, 00:32:59 ago
  Routing Descriptor Blocks:
  * 10.1.13.3, from 0.0.0.3, 00:32:59 ago, via FastEthernet0/0.13
    Route metric is 2, traffic share count is 1
```

The output of the above show commands reveal that the route is present in RIP's database, but the route in R1 is known via OSPF, but we are redistributing RIP into EIGRP.

To accomplish this task, the administrative distance of 192.168.23.0 /24 is raised in OSPF by a value that is higher than OSPF's Administrative distance:

On R1:

The desired AD value

It's important to note that this is the OSPF RID of R3

```
R1(config)#Router ospf 1
R1(config-router)#Distance 121 0.0.0.3 0.0.0.0 1
```

This is the inverse mask of R3's RID

```
R1(config)#Access-list 1 permit 192.168.23.0 0.0.0.255
```

To verify the configuration:

On R1:

Note: The route is no longer in the routing table as an OSPF route:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
O IA     10.1.1.3/32 [121/2] via 10.1.13.3, 00:02:54, FastEthernet0/0.13
O E2 192.168.31.0/24 [121/20] via 10.1.13.3, 00:02:54, FastEthernet0/0.13
O E2 192.168.33.0/24 [121/20] via 10.1.13.3, 00:02:54, FastEthernet0/0.13
```

Because of the adjusted AD, the route is now a RIP route:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
R        10.1.1.2/32 [120/1] via 10.1.12.2, 00:00:22, FastEthernet0/0.12
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R      192.168.23.0/24 [120/1] via 10.1.12.2, 00:00:22, FastEthernet0/0.12
```

Note: The output of the above show command reveals that the configured AD is applied to network 192.168.23.0 /24 received from R3.

To verify the configuration:

On R4:

```
R4#Show ip route 192.168.23.0

Routing entry for 192.168.23.0/24
  Known via "eigrp 100", distance 170, metric 2560002816, type external
  Redistributing via eigrp 100
  Last update from 10.1.14.1 on FastEthernet0/0, 00:05:10 ago
  Routing Descriptor Blocks:
  * 10.1.14.1, from 10.1.14.1, 00:05:10 ago, via FastEthernet0/0
      Route metric is 2560002816, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```

Task 8

Configure the following Loopback interface on R2:

Lo40: 172.16.0.2/24

R2 should be configured to advertise this loopback interface in RIPv2.
Ensure that the **only** 172.16.X.X route present in R1's routing table is 172.16.0.0/24.
R3 should see this route as 172.16.0.0/16.

Step one:

Configuring loopback 40 on R2 and advertise this network in RIPv2:

On R2:

```
R2(config)#int lo40
R2(config-if)#ip addr 172.16.0.1 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-if)#Router rip
R2(config-router)#Netw 172.16.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 10 subnets, 2 masks
R        10.1.1.2/32 [120/1] via 10.1.12.2, 00:00:24, FastEthernet0/0.12
      172.16.0.0/24 is subnetted, 1 subnets
R        172.16.0.0 [120/1] via 10.1.12.2, 00:00:24, FastEthernet0/0.12
R     192.168.23.0/24 [120/1] via 10.1.12.2, 00:00:24, FastEthernet0/0.12
```

R1 has the route based on the requirement of the task. To perform the last step of this task, R1 should be configured to redistribute RIPv2 into OSPF.

On R1:

By default, OSPF assigns a metric of 20 and a metric-type of 2 to all redistributed routes; therefore, there is no need to assign the cost from OSPF's perspective when redistributing routes.

```
R1(config)#Router ospf 1
R1(config-router)#Redistribute rip subnets
```

To verify the configuration:

On R3:

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

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O E2     10.1.1.1/32 [110/20] via 10.1.13.1, 00:00:22, FastEthernet0/0
O E2     10.1.1.2/32 [110/20] via 10.1.13.1, 00:00:22, FastEthernet0/0
O E2     10.1.12.0/24 [110/20] via 10.1.13.1, 00:00:22, FastEthernet0/0
O E2     10.1.14.0/24 [110/20] via 10.1.13.1, 00:00:22, FastEthernet0/0
      172.16.0.0/24 is subnetted, 1 subnets
O E2     172.16.0.0 [110/20] via 10.1.13.1, 00:00:22, FastEthernet0/0
```

Note: This task specified that network 172.16.0.0 should appear in the routing table of R3 as a "/16", to accomplish this task, this network is summarized, as follows:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#Summary-address 172.16.0.0 255.255.0.0
```

To verify the configuration:

On R3:

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

      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O E2     10.1.1.1/32 [110/20] via 10.1.13.1, 00:01:57, FastEthernet0/0
O E2     10.1.1.2/32 [110/20] via 10.1.13.1, 00:01:57, FastEthernet0/0
O E2     10.1.12.0/24 [110/20] via 10.1.13.1, 00:01:57, FastEthernet0/0
O E2     10.1.14.0/24 [110/20] via 10.1.13.1, 00:01:57, FastEthernet0/0
O E2  172.16.0.0/16 [110/20] via 10.1.13.1, 00:00:19, FastEthernet0/0
```

Note: This fulfills the requirement of the last task but remember this task specified that R1 should only have 172.16.0.0 /24 in its routing table.

To verify the configuration:

On R1:

```
R1#Show ip route | Inc 172.16

     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
R       172.16.0.0/24 [120/1] via 10.1.12.1, 00:00:15, FastEthernet0/0.21
O       172.16.0.0/16 is a summary, 00:09:41, Null0
```

R1 has two 172.16.X.X routes in its routing table, 172.16.0.0/24 which was learned from RIP and 172.16.0.0/16 which is the direct result of summarization. Remember that when we summarize in OSPF the IOS injects a discard route to avoid forwarding loops, the injected discard route can be for internal OSPF routes that were summarized or external OSPF routes that were summarized.

If the discard route is internal and it needs to be removed, then the "No discard-route internal" command can be used, but in this case it's for external routes, therefore, it can be removed using the following configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Router ospf 1
R1(config-router)#No discard-route external
```

We could have used the "Discard-route external 255" command to accomplish this task, remember any route that has an administrative distance of 255 will NOT be injected into the routing table.

To verify the configuration:

On R1:

```
R1#Show ip route | I 172.16

      172.16.0.0/24 is subnetted, 1 subnets
R        172.16.0.0 [120/1] via 10.1.12.2, 00:00:06, FastEthernet0/0.12
```

On R3:

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

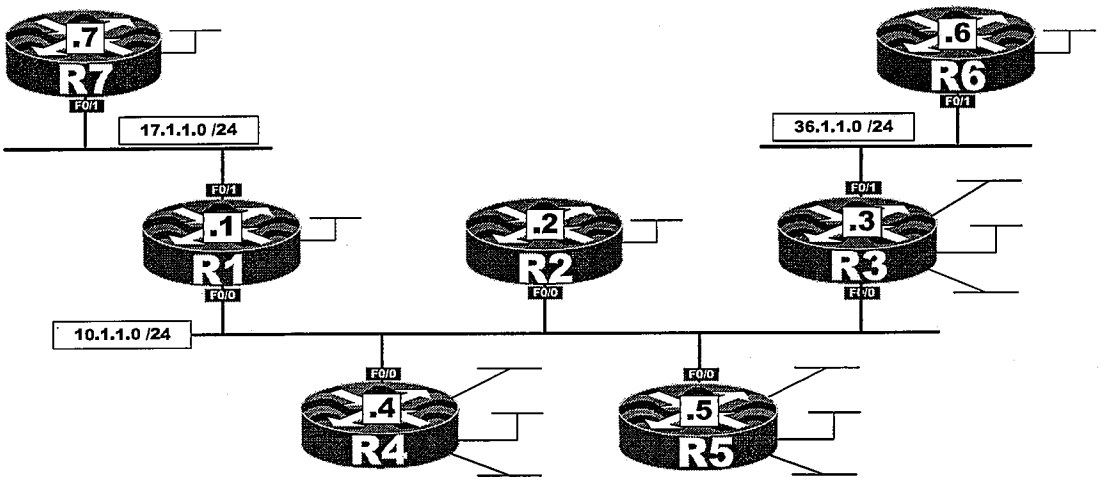
      10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O E2     10.1.1.1/32 [110/20] via 10.1.13.1, 00:09:34, FastEthernet0/0
O E2     10.1.1.2/32 [110/20] via 10.1.13.1, 00:09:34, FastEthernet0/0
O E2     10.1.12.0/24 [110/20] via 10.1.13.1, 00:09:34, FastEthernet0/0
O E2     10.1.14.0/24 [110/20] via 10.1.13.1, 00:09:34, FastEthernet0/0
O E2  172.16.0.0/16 [110/20] via 10.1.13.1, 00:07:56, FastEthernet0/0
```

Task 9

Erase the startup configuration and reload the routers before proceeding to the next task.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 2 – Basics of Redistribution-II



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "Redistribution"→ "Lab-2".

Task 1

Configure OSPF area 0 on the following routers and their interfaces:

- R3 – F0/0, Lo1
- R4, and R5 – F0/0, Lo0, Lo1 and Lo2
- Loopback interfaces should be advertised with their correct mask.

Assign the following OSPF costs to the loopback interfaces of R4 and R5:

- R4 & R5's – Lo0 should have a cost of 10, Lo1 should have a cost of 20, and Lo2 should have a Cost of 30.

When configuring OSPF in this task, OSPF Process ID of 1 should be used

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config)#Router ospf 1
R3(config-router)#Netw 10.1.1.3 0.0.0.0 area 0
R3(config-router)#Netw 30.3.3.3 0.0.0.0 area 0

R3(config)#Int Lo1
R3(config-if)#IP OSPF Network Point-to-point
```

On R4:

```
R4(config-router)#int lo0
R4(config-if)#IP OSPF Network Point-to-point
R4(config-if)#ip ospf cost 10

R4(config-router)#int lo1
R4(config-if)#IP OSPF Network Point-to-point
R4(config-if)#ip ospf cost 20

R4(config-router)#int lo2
R4(config-if)#IP OSPF Network Point-to-point
R4(config-if)#ip ospf cost 30

R4(config)#Router ospf 1
R4(config-router)#Netw 10.1.1.4 0.0.0.0 area 0
R4(config-router)#Netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#Netw 40.4.4.4 0.0.0.0 area 0
R4(config-router)#Netw 44.4.4.4 0.0.0.0 area 0
```

On R5:

```
R5(config-router)#int lo0
R5(config-if)#IP OSPF Network Point-to-point
R5(config-if)#ip ospf cost 10

R5(config-router)#int lo1
R5(config-if)#IP OSPF Network Point-to-point
R5(config-if)#ip ospf cost 20

R5(config-router)#int lo2
R5(config-if)#IP OSPF Network Point-to-point
R5(config-if)#ip ospf cost 30

R5(config)#Router ospf 1
R5(config-router)#Netw 10.1.1.5 0.0.0.0 area 0
R5(config-router)#Netw 5.5.5.5 0.0.0.0 area 0
R5(config-router)#Netw 50.5.5.5 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5(config-router)#Netw 55.5.5.5 0.0.0.0 area 0
```

To verify the configuration:

On R3:

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

      4.0.0.0/24 is subnetted, 1 subnets
O        4.4.4.0 [110/11] via 10.1.1.4, 00:02:38, FastEthernet0/0
      5.0.0.0/24 is subnetted, 1 subnets
O        5.5.5.0 [110/11] via 10.1.1.5, 00:00:52, FastEthernet0/0
      40.0.0.0/24 is subnetted, 1 subnets
O        40.4.4.0 [110/21] via 10.1.1.4, 00:02:28, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/31] via 10.1.1.4, 00:02:28, FastEthernet0/0
      50.0.0.0/24 is subnetted, 1 subnets
O        50.5.5.0 [110/21] via 10.1.1.5, 00:00:42, FastEthernet0/0
      55.0.0.0/24 is subnetted, 1 subnets
O        55.5.5.0 [110/31] via 10.1.1.5, 00:00:32, FastEthernet0/0
```

Task 2

Configure EIGRP AS 100 on the following routers and their interfaces:

- R1, R2 and R3's F0/0 and the loopback0 interfaces.
- The F0/1 interfaces of R1 and R7
- Lo100, Lo101 and Lo102 interfaces of R7

On R3:

```
R3(config)#Router eigrp 100
R3(config-router)#Netw 10.1.1.3 0.0.0.0
R3(config-router)#Netw 3.3.3.3 0.0.0.0
```

On R2:

```
R2(config)#Router eigrp 100
R2(config-router)#Netw 10.1.1.2 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Netw 2.2.2.2 0.0.0.0
```

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Netw 10.1.1.1 0.0.0.0
R1(config-router)#Netw 17.1.1.1 0.0.0.0
R1(config-router)#Netw 1.1.1.1 0.0.0.0
```

On R7:

```
R7(config)#Router eigrp 100
R7(config-router)#Netw 17.1.1.7 0.0.0.0
R7(config-router)#Netw 100.1.1.7 0.0.0.0
R7(config-router)#Netw 101.1.1.7 0.0.0.0
R7(config-router)#Netw 102.1.1.7 0.0.0.0
```

To verify the configuration:

On R3:

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

      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/156160] via 10.1.1.1, 00:08:10, FastEthernet0/0
      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/156160] via 10.1.1.2, 00:08:48, FastEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D        17.1.1.0 [90/30720] via 10.1.1.1, 00:03:14, FastEthernet0/0
      100.0.0.0/24 is subnetted, 1 subnets
D        100.1.1.0 [90/158720] via 10.1.1.1, 00:00:42, FastEthernet0/0
      101.0.0.0/24 is subnetted, 1 subnets
D        101.1.1.0 [90/158720] via 10.1.1.1, 00:00:28, FastEthernet0/0
      102.0.0.0/24 is subnetted, 1 subnets
D        102.1.1.0 [90/158720] via 10.1.1.1, 00:00:20, FastEthernet0/0
```

Task 3

Configure OSPF Area 0 with process ID of 36 on the following routers and their interfaces:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

- The F0/1 interfaces of R3 and R6
- The Lo0 interface of R6
- The Lo2 interface of R3

These loopback interfaces should be advertised with their correct mask

On R3:

```
R3(config-if)#int lo2
R3(config-if)#IP OSPF Network Point-to-point


R3(config)#Router ospf 36
R3(config-router)#Netw 36.1.1.3 0.0.0.0 area 0
R3(config-router)#Netw 33.3.3.3 0.0.0.0 area 0
```

On R6:

```
R6(config)#int lo0
R6(config-if)#IP OSPF Network Point-to-point


R6(config)#Router ospf 36
R6(config-router)#Netw 36.1.1.6 0.0.0.0 area 0
R6(config-router)#Netw 6.6.6.6 0.0.0.0 area 0
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf 36 | inc O

O        6.6.6.0 [110/2] via 10.1.36.6, 00:00:59, FastEthernet0/1
```

Task 4

Configure R3 to redistribute EIGRP 100 into OSPF 1 such that networks 1.1.1.0 /24 and 2.2.2.0 /24 carry a route tag of 111 and 222 respectively, the rest of the routes should have a route tag of 333.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3(config)#Access-list 1 permit 1.1.1.0 0.0.0.255
R3(config)#Access-list 2 permit 2.2.2.0 0.0.0.255

R3(config)#Route-map TST permit 10
R3(config-route-map)#match ip addr 1
R3(config-route-map)#set tag 111

R3(config)#Route-map TST permit 20
R3(config-route-map)#match ip addr 2
R3(config-route-map)#set tag 222

R3(config)#Route-map TST permit 30
R3(config-route-map)#set tag 333

R3(config)#Router ospf 1
R3(config-router)#Redistribute eigrp 100 route-map TST subnets
```

To verify the configuration:

On R3:

```
R3#Show ip ospf 1 database | B Type-5
                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
1.1.1.0         0.0.0.3         38          0x80000001 0x001D01 111
2.2.2.0         0.0.0.3         38          0x80000001 0x00DCCD 222
3.3.3.0         0.0.0.3         38          0x80000001 0x00E95C 333
17.1.1.0        0.0.0.3         38          0x80000001 0x00F638 333
100.1.1.0       0.0.0.3         38          0x80000001 0x00R80 333
101.1.1.0       0.0.0.3         38          0x80000001 0x00AE2C 333
102.1.1.0       0.0.0.3         38          0x80000001 0x00A138 333
```

On R4:

```
R4#Show ip route ospf | Inc E2

       E1 - OSPF external type 1, E2 - OSPF external type 2
O E2     1.1.1.0 [110/20] via 10.1.1.1, 00:03:22, FastEthernet0/0
O E2     2.2.2.0 [110/20] via 10.1.1.2, 00:03:22, FastEthernet0/0
O E2     3.3.3.0 [110/20] via 10.1.1.3, 00:03:22, FastEthernet0/0
O E2     17.1.1.0 [110/20] via 10.1.1.1, 00:03:22, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
O E2     100.1.1.0 [110/20] via 10.1.1.1, 00:03:22, FastEthernet0/0
O E2     101.1.1.0 [110/20] via 10.1.1.1, 00:03:22, FastEthernet0/0
O E2     102.1.1.0 [110/20] via 10.1.1.1, 00:03:22, FastEthernet0/0
```

On R4:

```
R4#Show ip ospf database | B Type-5
                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
1.1.1.0         0.0.0.3         256         0x80000001 0x001D01 111
2.2.2.0         0.0.0.3         256         0x80000001 0x00DCCD 222
3.3.3.0         0.0.0.3         256         0x80000001 0x00E95C 333
17.1.1.0        0.0.0.3         256         0x80000001 0x00F638 333
100.1.1.0       0.0.0.3         256         0x80000001 0x00R80 333
101.1.1.0       0.0.0.3         256         0x80000001 0x00AE2C 333
102.1.1.0       0.0.0.3         256         0x80000001 0x00A138 333
```

Task 5

Configure R4 to filter all routes that are tagged with 111, you should **not** use an access-list or a prefix-list to accomplish this task.

On R4:

```
R4(config)#Route-map TST deny 10
R4(config-route-map)#match tag 111

R4(config)#Route-map TST permit 20

R4(config)#Router ospf 1
R4(config-router)#Distribute-list route-map TST in
```

To verify the configuration:

On R4:

```
R4#Show ip route ospf | Inc E2

       E1 - OSPF external type 1, E2 - OSPF external type 2
O E2     2.2.2.0 [110/20] via 10.1.1.2, 00:00:18, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O E2     3.3.3.0 [110/20] via 10.1.1.3, 00:00:18, FastEthernet0/0
O E2     17.1.1.0 [110/20] via 10.1.1.1, 00:00:18, FastEthernet0/0
O E2     100.1.1.0 [110/20] via 10.1.1.1, 00:00:18, FastEthernet0/0
O E2     101.1.1.0 [110/20] via 10.1.1.1, 00:00:18, FastEthernet0/0
O E2     102.1.1.0 [110/20] via 10.1.1.1, 00:00:18, FastEthernet0/0
```

Route tags are very handy tool, once the routes or a given route is tagged with a value, that route can be manipulated based on the tag that it carries.

Note: Network 1.1.1.0 /24 is no longer in the routing table of R1, but you should remember that it is still in the database of this router, the output of the following show command reveals that information:

```
R4#Show ip ospf database external 1.1.1.0

            OSPF Router with ID (0.0.0.4) (Process ID 1)

                Type-5 AS External Link States

  Routing Bit Set on this LSA in topology Base with MTID 0
  LS age: 536
  Options: (No TOS-capability, DC)
  LS Type: AS External Link
  Link State ID: 1.1.1.0 (External Network Number )
  Advertising Router: 0.0.0.3
  LS Seq Number: 80000001
  Checksum: 0x1D01
  Length: 36
  Network Mask: /24
        Metric Type: 2 (Larger than any link state path)
        MTID: 0
        Metric: 20
        Forward Address: 10.1.1.1
        External Route Tag: 111
```

Task 6

Configure R5 to filter all routes that carry a route tag of 222, you should **not** use an access-list or a prefix-list to accomplish this task.

On R5:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5(config)#Route-map TST deny 10
R5(config-route-map)#match tag 222

R5(config)#Route-map TST permit 90

R5(config)#Router ospf 1
R5(config-router)#distribute-list route-map TST in
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | Inc E2

       E1 - OSPF external type 1, E2 - OSPF external type 2
O E2     1.1.1.0 [110/20] via 10.1.1.1, 00:00:32, FastEthernet0/0
O E2     3.3.3.0 [110/20] via 10.1.1.3, 00:00:32, FastEthernet0/0
O E2     17.1.1.0 [110/20] via 10.1.1.1, 00:00:32, FastEthernet0/0
O E2     100.1.1.0 [110/20] via 10.1.1.1, 00:00:32, FastEthernet0/0
O E2     101.1.1.0 [110/20] via 10.1.1.1, 00:00:32, FastEthernet0/0
O E2     102.1.1.0 [110/20] via 10.1.1.1, 00:00:32, FastEthernet0/0
```

Task 7

Configure R3 to redistribute OSPF 1 into EIGRP 100 such that network 4.4.4.0 /24 and 5.5.5.0 /24 are tagged with 444 and 555 respectively.

On R3:

```
R3(config)#Access-list 4 permit 4.4.4.0 0.0.0.255

R3(config)#Access-list 5 permit 5.5.5.0 0.0.0.255

R3(config)#Route-map TST1 permit 10
R3(config-route-map)#match ip addr 4
R3(config-route-map)#set tag 444

R3(config)#Route-map TST1 permit 20
R3(config-route-map)#match ip addr 5
R3(config-route-map)#set tag 555
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config)#Route-map TST1 permit 30

R3(config)#Router eigrp 100
R3(config-router)#redistribute ospf 1 route-map TST1 metric 1 1 1 1 1
```

To verify the configuration:

On R3:

```
R3#Sh ip eigrp topology | inc 444|555

P 5.5.5.0/24, 1 successors, FD is 2560000256, tag is 555
P 4.4.4.0/24, 1 successors, FD is 2560000256, tag is 444
```

Task 8

Configure R2 to filter network 4.4.4.0 /24; do not use an access-list or a prefix-list to accomplish this task.

On R2:

```
R2#Sh ip eigrp topology | Inc 444

P 4.4.4.0/24, 1 successors, FD is 2560002816, tag is 444
```

Note: R2 sees the tag that was assigned in the previous task.

```
R2(config)#Route-map TST deny 10
R2(config-route-map)#match tag 444

R2(config)#Route-map TST permit 20

R2(config)#Router eigrp 100
R2(config-router)#distribute-list route-map TST in
```

You should see the following console messages verifying that some route configuration has changed:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.3 (FastEthernet0/0) is
resync: route configuration changed

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.1 (FastEthernet0/0) is
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
resync: route configuration changed
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc EX

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
D EX     5.5.5.0 [170/2560002816] via 10.1.1.3, 00:05:17, FastEthernet0/0
D EX    30.3.3.0 [170/2560002816] via 10.1.1.3, 00:05:17, FastEthernet0/0
D EX    40.4.4.0 [170/2560002816] via 10.1.1.3, 00:05:17, FastEthernet0/0
D EX    44.4.4.0 [170/2560002816] via 10.1.1.3, 00:05:17, FastEthernet0/0
D EX    50.5.5.0 [170/2560002816] via 10.1.1.3, 00:05:17, FastEthernet0/0
D EX    55.5.5.0 [170/2560002816] via 10.1.1.3, 00:05:17, FastEthernet0/0
```

Task 9

Configure R5 to filter all routes that originated by R4. **Do not** use a prefix-list or route tags to accomplish this task.

On R5:

Note: R5 receives the following three routes from R4:

```
R5#Show ip route | Inc 10.1.1.4

O        4.4.4.0 [110/11] via 10.1.1.4, 00:10:30, FastEthernet0/0
O        40.4.4.0 [110/21] via 10.1.1.4, 00:10:30, FastEthernet0/0
O        44.4.4.0 [110/31] via 10.1.1.4, 00:10:30, FastEthernet0/0
```

To configure the task:

The "match ip route-source" can be used to accomplish this task; in ospf the route-source is not the next hop ip address, it is the router id of the router that advertised the routes; to see the router id of R4:

```
R4#Sh ip ospf | inc ID

 Routing Process "ospf 1" with ID 0.0.0.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config)#Access-list 4 permit host 0.0.0.4
```

In task 6 a route-map was configured, let's see that route-map:

```
R5#Sh run | S route-map

 distribute-list route-map TST in
route-map TST deny 10
 match tag 222
route-map TST permit 90
```

In the above output we can see that route-map "TST" denies any route/s that carry a route tag of 222, and this route-map is applied inbound on R5. To resolve this task, we can add the following line to the existing route-map:

```
R5(config)#Route-map TST deny 20
R5(config-route-map)#Match ip route-source 4
```

The above configuration denies any prefix originated from a route source with a router-id that is referenced in access-list "4" and "access-list 4" is permiting "0.0.0.4". Since the above configuration is added to the existing route-map, as a result the new route-map should look like the following:

On R5:

```
R5#Show run | S route-map

 distribute-list route-map TST in
route-map TST deny 10
 match tag 222
route-map TST deny 20
 match ip route-source 4
route-map TST permit 90
```

Note: The reason we have the sequence numbers in a route-map is so we can add configuration lines in the beginning, middle, or at the end of a route-map, in this case route-map sequence number 20 is added between route-map 10 and 90.

Since the route-map is already applied, no further configuration is necessary.

To verify the configuration:

On R5:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R5#Show ip route | Inc 10.1.1.4
R5#
```

NOTE: The networks that were advertised by R4 are filtered.

Task 10

Configure R3 to redistribute "OSPF 1" and "Eigrp 100" into "OSPF 36" routing domain using the following policy:

- The routes from "OSPF 1" routing domain with an OSPF cost of 11 should be redistributed as "E2" OSPF routes, these routes should have a tag of 11
- The routes from "OSPF 1" routing domain with an OSPF cost of 21 should be redistributed as "E2" OSPF routes, these routes should have a tag of 21
- The routes from "OSPF 1" routing domain with an OSPF cost of 31 should be redistributed as "E2" OSPF routes, these routes should have a tag of 31
- Eigrp 100 routes should be redistributed as "E1" routes with a route tag of 99.
- **Do not** use an access-list or prefix-list to accomplish this task

Let's see the OSPF routes and their cost:

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

      4.0.0.0/24 is subnetted, 1 subnets
O        4.4.4.0 [110/11] via 10.1.1.4, 00:09:07, FastEthernet0/0
      5.0.0.0/24 is subnetted, 1 subnets
O        5.5.5.0 [110/11] via 10.1.1.5, 00:09:07, FastEthernet0/0
      6.0.0.0/24 is subnetted, 1 subnets
O        6.6.6.0 [110/2] via 36.1.1.6, 00:09:07, FastEthernet0/1
      40.0.0.0/24 is subnetted, 1 subnets
O        40.4.4.0 [110/21] via 10.1.1.4, 00:09:07, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/31] via 10.1.1.4, 00:09:07, FastEthernet0/0
      50.0.0.0/24 is subnetted, 1 subnets
O        50.5.5.0 [110/21] via 10.1.1.5, 00:09:07, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      55.0.0.0/24 is subnetted, 1 subnets
O        55.5.5.0 [110/31] via 10.1.1.5, 00:09:07, FastEthernet0/0
```

The "Source-protocol" is another useful and effective tool that we can use to achieve certain results. The "Source-protocol" specifies a match clause that matches external routes from sources that match the source protocol.

On R3:

```
R3(config)#Route-map Task10 permit 10
R3(config-route-map)#match source-protocol OSPF 1
R3(config-route-map)#match metric 11
R3(config-route-map)#set tag 11

R3(config)#Route-map Task10 permit 20
R3(config-route-map)#match source-protocol OSPF 1
R3(config-route-map)#match metric 21
R3(config-route-map)#set tag 21


R3(config)#Route-map Task10 permit 30
R3(config-route-map)#match source-protocol ospf 1
R3(config-route-map)#match metric 31
R3(config-route-map)#set tag 31
R3(config)#Route-map Task10 permit 40

R3(config)#Route-map eigrp permit 10
R3(config-route-map)#match source-protocol eigrp 100
R3(config-route-map)#set tag 99
R3(config-route-map)#set metric-type type-1

R3(config)#Router ospf 36
R3(config-router)#Redistribute ospf 1 route-map Task10 subnets
R3(config-router)#Redistribute eigrp 100 subnets route-map eigrp
```

To verify the configuration:

On R6:

```
R6#Show ip route ospf | Inc E1

       E1 - OSPF external type 1, E2 - OSPF external type 2
O E1     1.1.1.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
O E1     2.2.2.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
O E1     3.3.3.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
O E1     10.1.1.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
O E1     17.1.1.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
O E1     100.1.1.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
O E1     101.1.1.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
O E1     102.1.1.0 [110/21] via 36.1.1.3, 00:00:22, FastEthernet0/1
```

```
R6#Show ip route ospf | Inc E2

      E1 - OSPF external type 1, E2 - OSPF external type 2
O E2     4.4.4.0 [110/11] via 36.1.1.3, 00:01:06, FastEthernet0/1
O E2     5.5.5.0 [110/11] via 36.1.1.3, 00:01:06, FastEthernet0/1
O E2     30.3.3.0 [110/1] via 36.1.1.3, 00:01:06, FastEthernet0/1
O E2     40.4.4.0 [110/21] via 36.1.1.3, 00:01:06, FastEthernet0/1
O E2     44.4.4.0 [110/31] via 36.1.1.3, 00:01:06, FastEthernet0/1
O E2     50.5.5.0 [110/21] via 36.1.1.3, 00:01:06, FastEthernet0/1
O E2     55.5.5.0 [110/31] via 36.1.1.3, 00:01:06, FastEthernet0/1

R6#Show ip ospf database | Inc 11_

4.4.4.0         0.0.0.33        122         0x80000001 0x00FC74 11
5.5.5.0         0.0.0.33        122         0x80000001 0x00D895 11

R6#Show ip ospf database | inc 21_

40.4.4.0        0.0.0.33        176         0x80000001 0x0040F8 21
50.5.5.0        0.0.0.33        176         0x80000001 0x00A686 21

R6#Show ip ospf database | inc 31_

44.4.4.0        0.0.0.33        220         0x80000001 0x0025FB 31
55.5.5.0        0.0.0.33        220         0x80000001 0x007E95 31

R6#Show ip ospf database | inc 99_

1.1.1.0         0.0.0.33        242         0x80000001 0x007623 99
2.2.2.0         0.0.0.33        242         0x80000001 0x005244 99
3.3.3.0         0.0.0.33        242         0x80000001 0x002E65 99
10.1.1.0        0.0.0.33        242         0x80000002 0x00FE90 99
17.1.1.0        0.0.0.33        242         0x80000001 0x00A5E3 99
100.1.1.0       0.0.0.33        242         0x80000001 0x006ACB 99
101.1.1.0       0.0.0.33        242         0x80000001 0x005DD7 99
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
102.1.1.0        0.0.0.33        242          0x80000001 0x0050E3 99
```

Task 11

Configure R3 such that Eigrp routes that have a composite metric of 156160 to 158720 are **not** redistributed into "OSPF 36" routing domain. None of the previous configurations should be removed or overridden to accomplish this task.

This task is asking to filter the routes that have a composite metric of 156160 to 158720, the formula to calculate the metric and the deviation value is as follows:

The two numbers are added and then divided by 2:

156160 + 158720 = 314880
314880 / 2 = 157440 → This gives us the metric value

To calculate the deviation value:

Subtract the start of the range number (the lower value or the *"from"* value, in this case 156160) from the *"to"* value, in this case 158720 and then divide the result by two:

158720 – 156160 = 2560
2560 / 2 = 1280

Therefore, the "match metric" statement in the route-map should have the following values:

Match metric 157440 +-1280

To test the values:

157440 - 1280 = 156160 → This is the lowest value within the range.
157440 + 1280 = 158720 → This is the highest value within the range.

Let's see the existing EIGRP routes:

On R3:

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


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      1.0.0.0/24 is subnetted, 1 subnets
D        1.1.1.0 [90/156160] via 10.1.1.1, 01:37:26, FastEthernet0/0
      2.0.0.0/24 is subnetted, 1 subnets
D        2.2.2.0 [90/156160] via 10.1.1.2, 01:37:29, FastEthernet0/0
      17.0.0.0/24 is subnetted, 1 subnets
D        17.1.1.0 [90/30720] via 10.1.1.1, 01:37:26, FastEthernet0/0
      100.0.0.0/24 is subnetted, 1 subnets
D        100.1.1.0 [90/158720] via 10.1.1.1, 01:37:26, FastEthernet0/0
      101.0.0.0/24 is subnetted, 1 subnets
D        101.1.1.0 [90/158720] via 10.1.1.1, 01:37:26, FastEthernet0/0
      102.0.0.0/24 is subnetted, 1 subnets
D        102.1.1.0 [90/158720] via 10.1.1.1, 01:37:26, FastEthernet0/0
```

A route-map is configured to deny any route/s that have a composite metric of 156160 to 158720 inclusively:

```
R3(config)#Route-map Task-11 deny 10
R3(config-route-map)#match metric 157440 +- 1280

R3(config)#Route-map Task-11 permit 90

R3(config-route-map)#Router eigrp 100
R3(config-router)#distribute-list route-map Task-11 in
```

To verify the configuration:

On R3:

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

      17.0.0.0/24 is subnetted, 1 subnets
D        17.1.1.0 [90/30720] via 10.1.1.1, 01:41:44, FastEthernet0/0
```

Note: If it is not in the routing table of R3 as an EIGRP route, it will not be redistributed.

On R6:

```
R6#Show ip ospf database | inc 99_

3.3.3.0         0.0.0.33        1104        0x80000003 0x002A67 99
10.1.1.0        0.0.0.33        1104        0x80000003 0x00FC91 99
17.1.1.0        0.0.0.33        1104        0x80000003 0x00A1E5 99
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

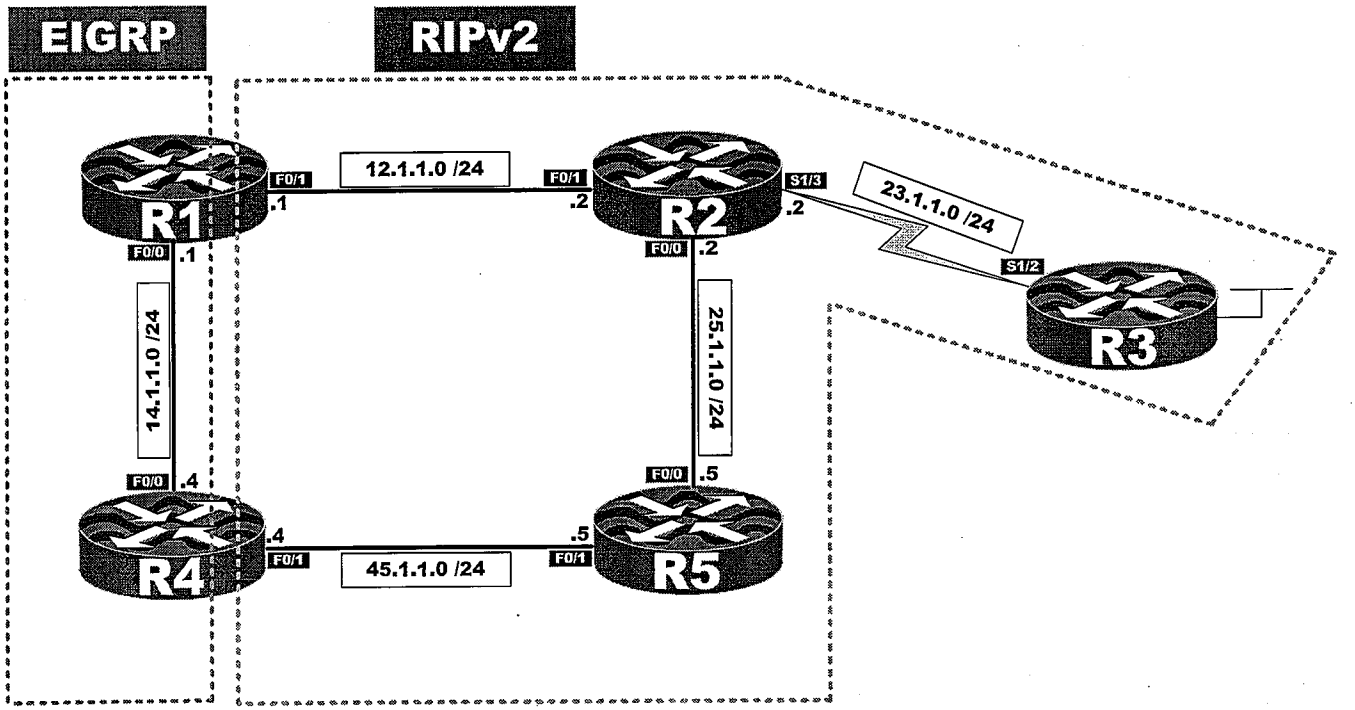
NOTE: The above Networks 3.3.3.0, 17.1.1.0 and 10.1.1.0 are NOT within the stated range.

Task 12

Erase the startup configuration and reload the routers before proceeding to the next task.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 3
Redistribution – RIPv2 & Eigrp



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "Redistribution"→ "Lab-3".

Task 1

Configure RIPv2 on R2, R3 and R5, these routers should advertise their directly connected interfaces in this routing domain. You should disable the auto-summary.

On R2:

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 12.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Network 23.0.0.0
R2(config-router)#Network 25.0.0.0
```

On R3:

```
R3(config)#Router rip
R3(config-router)#No au
R3(config-router)#ver 2
R3(config-router)#Network 3.0.0.0
R3(config-router)#Network 23.0.0.0
```

On R5:

```
R5(config)#Router rip
R5(config-router)#No au
R5(config-router)#ver 2
R5(config-router)#Network 25.0.0.0
R5(config-router)#Network 45.0.0.0
```

To verify the configuration:

On R5:

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

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/2] via 25.1.1.2, 00:00:20, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 25.1.1.2, 00:00:21, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 25.1.1.2, 00:00:21, FastEthernet0/0
```

Task 2

Configure RIPv2 on the F0/1 interfaces of R1 and R4.

On R1:

```
R1(config)#Router rip
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#Network 12.0.0.0
```

On R4:

```
R4(config)#Router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#Network 45.0.0.0
```

To verify the configuration:

On R4:

```
R4#Show ip route rip | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/3] via 45.1.1.5, 00:00:05, FastEthernet0/1
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/2] via 45.1.1.5, 00:00:05, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/2] via 45.1.1.5, 00:00:05, FastEthernet0/1
      25.0.0.0/24 is subnetted, 1 subnets
R        25.1.1.0 [120/1] via 45.1.1.5, 00:00:05, FastEthernet0/1
```

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/2] via 12.1.1.2, 00:00:19, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 12.1.1.2, 00:00:19, FastEthernet0/1
      25.0.0.0/24 is subnetted, 1 subnets
R        25.1.1.0 [120/1] via 12.1.1.2, 00:00:19, FastEthernet0/1
      45.0.0.0/24 is subnetted, 1 subnets
R        45.1.1.0 [120/2] via 12.1.1.2, 00:00:19, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 3

R3's Loopback0 interface should be advertised with a hop count of 5.

On R3:

```
R3(config)#Access-list 3 permit 3.3.3.0 0.0.0.255

R3(config)#Router rip
R3(config-router)#Offset-list 3 out 4
```

To verify the configuration:

On R2:

```
R2#Sh ip rou | Inc 3.3.3.0

R        3.3.3.0 [120/5] via 23.1.1.3, 00:00:23, Serial0/0.23
```

On R1:

```
R1#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/6] via 12.1.1.2, 00:00:03, FastEthernet0/1
```

On R5:

```
R5#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/6] via 25.1.1.2, 00:00:23, FastEthernet0/0
```

On R4:

```
R4#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/7] via 45.1.1.5, 00:00:09, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

Configure Eigrp AS 100 on the F0/0 interfaces of R1 and R4 and set the administrative distance of all external routes to 110.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Network 14.1.1.1 0.0.0.0
R1(config-router)#distance eigrp 90 110
```

On R4:

```
R4(config)#Router eigrp 100
R4(config-router)#Network 14.1.1.4 0.0.0.0
R4(config-router)#distance eigrp 90 110
```

To verify the configuration:

On R1:

```
R4#Show ip eigrp neighbor

IP-EIGRP neighbors for process 100
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   14.1.1.1                Fa0/0             10 03:05:54    1   200  0  41
```

Task 5

Perform a mutual redistribution between RIPv2 and Eigrp on Routers R1 and R4. Eigrp routes should be redistributed into RIPv2 routing domain with a hop count of 1, and RIPv2 routes should be redistributed into EIGRP using the "1 1 1 1 1" metric. You should perform a "**Clear ip route ***" once this redistribution is completed.

This configuration should be performed on R1 first and then on R4. Ensure that every router has rechability to 3.3.3.3 IP address, use Ping to verify.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Router rip
R1(config-router)#redistribute eigrp 100 metric 1

R1(config)#Router eigrp 100
R1(config-router)#redistribute rip metric 1 1 1 1 1
```

On R4:

```
R4(config)#Router rip
R4(config-router)#redistribute eigrp 100 metric 1

R4(config)#Router eigrp 100
R4(config-router)#redistribute rip metric 1 1 1 1 1
```

On All Routers:

```
Rx#Clear ip route *
```

To verify the configuration:

On R1:

```
R1#Show ip route | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/3] via 12.1.1.2, 00:00:16, FastEthernet0/1
      12.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        12.1.1.0/24 is directly connected, FastEthernet0/1
L        12.1.1.1/32 is directly connected, FastEthernet0/1
      14.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        14.1.1.0/24 is directly connected, FastEthernet0/0
L        14.1.1.1/32 is directly connected, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 12.1.1.2, 00:00:22, FastEthernet0/1
      25.0.0.0/24 is subnetted, 1 subnets
R        25.1.1.0 [120/1] via 12.1.1.2, 00:00:22, FastEthernet0/1
      45.0.0.0/24 is subnetted, 1 subnets
D EX     45.1.1.0 [110/2560002816] via 14.1.1.4, 00:00:24, FastEthernet0/0
```

On R2:

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


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/2] via 25.1.1.5, 00:00:23, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/1] via 12.1.1.1, 00:00:25, FastEthernet0/1
      45.0.0.0/24 is subnetted, 1 subnets
R        45.1.1.0 [120/1] via 25.1.1.5, 00:00:23, FastEthernet0/0
                  [120/1] via 12.1.1.1, 00:00:25, FastEthernet0/1
```

On R3:

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

      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 23.1.1.2, 00:00:19, Serial1/2
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/2] via 23.1.1.2, 00:00:19, Serial1/2
      25.0.0.0/24 is subnetted, 1 subnets
R        25.1.1.0 [120/1] via 23.1.1.2, 00:00:19, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
R        45.1.1.0 [120/2] via 23.1.1.2, 00:00:19, Serial1/2
```

On R4:

```
R4#Show ip route | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 1 subnets
D EX     3.3.3.0 [110/2560002816] via 14.1.1.1, 00:03:09, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
D EX     12.1.1.0 [110/2560002816] via 14.1.1.1, 00:03:11, FastEthernet0/0
      14.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        14.1.1.0/24 is directly connected, FastEthernet0/0
L        14.1.1.4/32 is directly connected, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
D EX     23.1.1.0 [110/2560002816] via 14.1.1.1, 00:03:11, FastEthernet0/0
      25.0.0.0/24 is subnetted, 1 subnets
D EX     25.1.1.0 [110/2560002816] via 14.1.1.1, 00:03:11, FastEthernet0/0
      45.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        45.1.1.0/24 is directly connected, FastEthernet0/1
L        45.1.1.4/32 is directly connected, FastEthernet0/1
```

On R5:

```
R5#Show ip route rip | B Gate
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/1] via 45.1.1.4, 00:00:26, FastEthernet0/1
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 45.1.1.4, 00:00:26, FastEthernet0/1
                  [120/1] via 25.1.1.2, 00:00:23, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/1] via 45.1.1.4, 00:00:26, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 45.1.1.4, 00:00:26, FastEthernet0/1
                  [120/1] via 25.1.1.2, 00:00:23, FastEthernet0/0
```

The output of the above show commands reveal that every router has every route in their routing table, but this only means that the control plane worked properly, let's verify and test the data plane by pinging 3.3.3.3 from R1:

Let's test reachability by pinging 3.3.3.3 from R4.

On R4:

```
R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the routing table of R4 and see the next hop IP address to reach 3.3.3.3:

```
R4#Show ip route | Inc 3.3.3.0

D EX    . 3.3.3.0 [110/2560002816] via 14.1.1.1, 02:50:34, FastEthernet0/0
```

It looks like R1 is the next hop to reach 3.3.3.3, let's go to R1 and do the same:

On R1:

```
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/3] via 12.1.1.2, 00:00:18, FastEthernet0/1
```

R1 is using R2 as the next hop router to reach network 3.3.3.0/24. Let's go to R2.

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2#Sh ip route | Inc 3.3.3.0

R        3.3.3.0 [120/2] via 25.1.1.5, 00:00:02, FastEthernet0/0
```

Hmmm, why is R2 going to R5 and NOT R3?

Well, R3 advertised network 3.3.3.0/24 with a hop count of 5, but R5 is advertising the same network with a hop count of 2, what is going on? Let's go to R5:

On R5:

```
R5#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/1] via 45.1.1.4, 00:00:22, FastEthernet0/1
```

This looks like a perfect loop, because R5 is using R4 as its next hop IP address to reach network 3.3.3.0/24. Let's use a "Traceroute" and verify:

On R4:

```
R4#Traceroute 3.3.3.3 Numeric

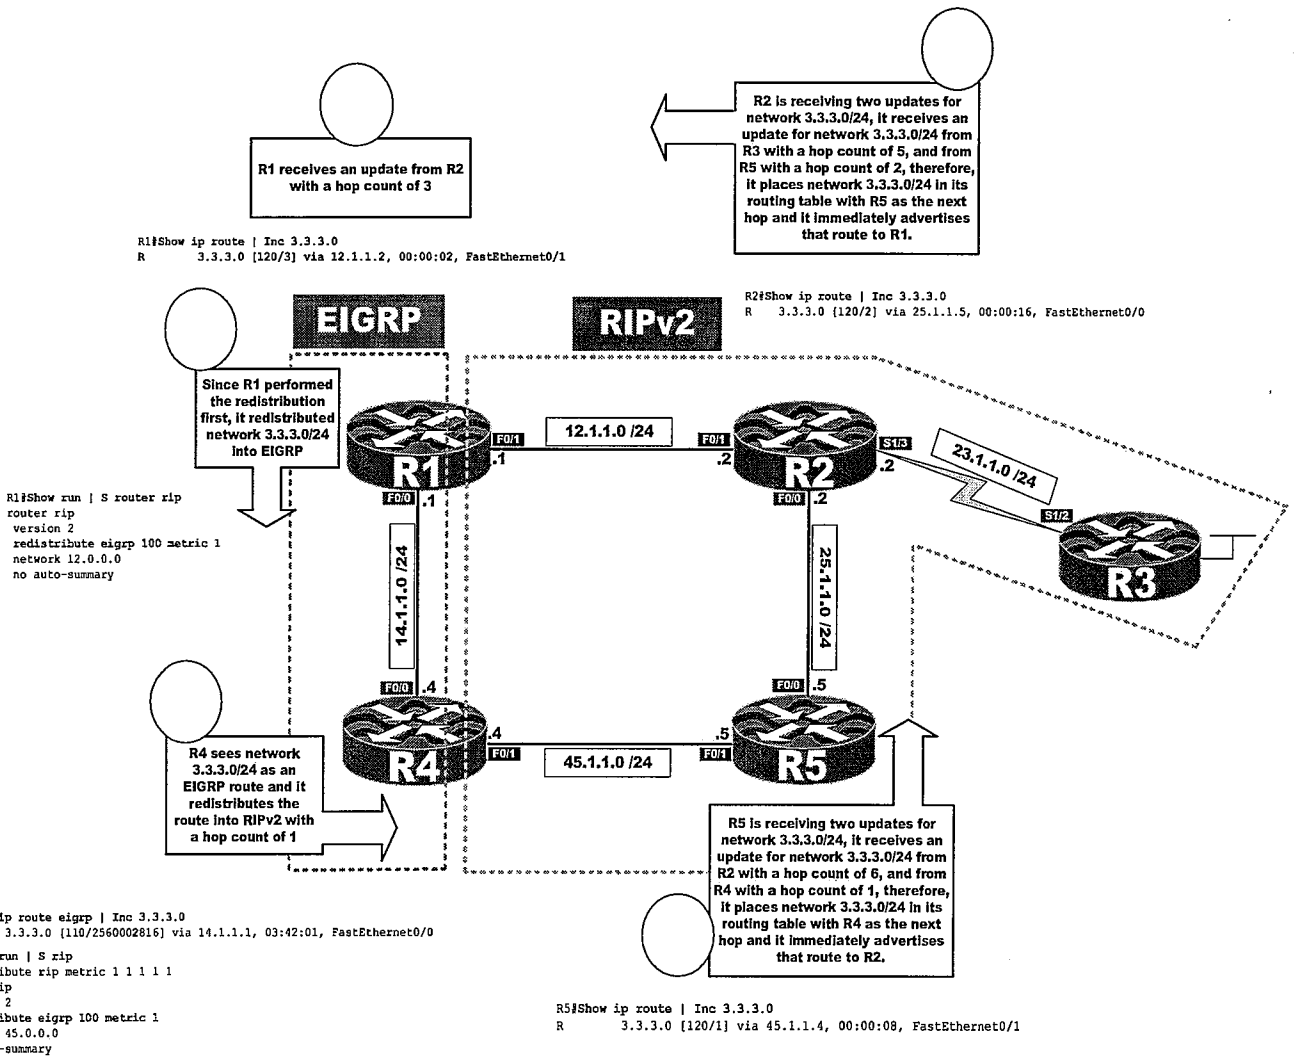
Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.1 0 msec 0 msec 0 msec
  2 12.1.1.2 4 msec 0 msec 4 msec
  3 25.1.1.5 0 msec 4 msec 0 msec
  4 45.1.1.4 4 msec 0 msec 0 msec
  5 14.1.1.1 4 msec 0 msec 4 msec
  6 12.1.1.2 0 msec 4 msec 0 msec
  7 25.1.1.5 4 msec 0 msec 0 msec
  8 45.1.1.4 0 msec 4 msec 0 msec
  9 14.1.1.1 4 msec 0 msec 4 msec
 10 12.1.1.2 0 msec 0 msec 4 msec
 11 25.1.1.5 0 msec 4 msec 0 msec
 12 45.1.1.4 4 msec 0 msec 4 msec
 13 14.1.1.1 0 msec 4 msec 0 msec
 14 12.1.1.2 0 msec 4 msec 0 msec
 15 25.1.1.5 4 msec 0 msec 0 msec
 16 45.1.1.4 4 msec 0 msec 4 msec
 17 14.1.1.1 0 msec 4 msec 4 msec
 18 12.1.1.2 0 msec 4 msec 0 msec
 19 25.1.1.5 0 msec 4 msec 0 msec
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
20 45.1.1.4 4 msec 4 msec 0 msec
21 14.1.1.1 4 msec 0 msec 0 msec
22 12.1.1.2 4 msec 0 msec 0 msec
23 25.1.1.5 4 msec 0 msec 0 msec
24 45.1.1.4 4 msec 0 msec 4 msec
25 14.1.1.1 4 msec 0 msec 4 msec
26 12.1.1.2 4 msec 0 msec 4 msec
27 25.1.1.5 4 msec 0 msec 0 msec
28 45.1.1.4 4 msec 4 msec 0 msec
29 14.1.1.1 4 msec 4 msec 0 msec
30 12.1.1.2 4 msec 4 msec 4 msec
```

Let's visually look at this:



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Now that we know the reason, how do we fix this problem?

One way to resolve this problem is to stop R1 and R4 from advertising network 3.3.3.0 /24 back into RIP routing domain.

Let's try this solution and verify:

On R1 and R4

```
Rx(config)#Access-list 1 deny 3.3.3.0 0.0.0.255
Rx(config)#Access-list 1 permit any

Rx(config)#Router rip
Rx(config-router)#Distribute-list 1 out F0/1
```

Let's Clear ip route * on all routers and verify the result:

On All Routers:

```
Rx#Clear ip route *
```

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/49/52 ms

R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 0 msec 4 msec 0 msec
  2 23.1.1.3 16 msec *  12 msec
```

We can see that R1 goes through R2 and then R3 to reach 3.3.3.3, but why are we getting the asterisk (*) in the second line of the "Traceroute"?
The reason we get this kind of timeout is because of the IOS default value for ICMP rate-limit. Let's verify this parameter:

```
R1#Show ip icmp rate-limit
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
                              DF bit unreachables       All other unreachables
Interval (millisecond)        500                       500


Interface                     # DF bit unreachables     # All other unreachables
---------                     ---------------------     ------------------------
FastEthernet0/0               0                         0
FastEthernet0/1               0                         0
```

The output of the above command reveals that the ICMP timeout is 500 ms, this value can be increased, decreased or negated all together. Let's negate the timeout:

On All Routers:

```
R1(config)#No ip icmp rate-limit unreachable
```

On R1:

```
R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 0 msec 4 msec 0 msec
  2 23.1.1.3 16 msec 16 msec 16 msec
```

On R2:

```
R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/48/52 ms

R2#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 23.1.1.3 16 msec 24 msec 32 msec
```

On R4:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms

R4#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.1 0 msec 4 msec 0 msec
  2 12.1.1.2 4 msec 0 msec 4 msec
  3 23.1.1.3 36 msec 32 msec 24 msec
```

On R5:

```
R5#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms

R5#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 25.1.1.2 4 msec 0 msec 4 msec
  2 23.1.1.3 24 msec 32 msec 32 msec
```

Task 6

Remove the previous solution and use another solution to fix the routing loops, do not use filtering to accomplish this task.

In this topology, we can configure the two Eigrp neighbors such that they discard any external routes from each other, this can be accomplished by configuring their Eigrp router-ids to match.

Let's remove the "access-list" and the "distribute-list" from R1 and R4 and verify:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1 and R4:

```
Rx(config)#No access-list 1

Rx(config)#Router rip
Rx(config-router)#No distribute-list 1 out FastEthernet0/1
```

Let's check the configuration of R1 and R4:

On R1:

```
R1#Show run | S router rip

router rip
 version 2
 redistribute eigrp 100 metric 1
 network 12.0.0.0
 no auto-summary

R1#Show run | S router eigrp 100
router eigrp 100
 redistribute rip metric 1 1 1 1 1
 network 14.1.1.1 0.0.0.0
 distance eigrp 90 110
```

On R4:

```
R4#Show run | S router rip

router rip
 version 2
 redistribute eigrp 100 metric 1
 network 45.0.0.0
 no auto-summary

R4#Show run | S router eigrp 100

router eigrp 100
 redistribute rip metric 1 1 1 1 1
 network 14.1.1.4 0.0.0.0
 distance eigrp 90 110
```

Let's clear the routing table on all routers:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On All Routers:

```
Rx#Clear ip route *
```

We should see the routing loop back again.

On R1:

```
R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3

  1 14.1.1.4 0 msec 4 msec 0 msec
  2 45.1.1.5 0 msec 0 msec 0 msec
  3 25.1.1.2 4 msec 0 msec 0 msec
  4 12.1.1.1 0 msec 0 msec 0 msec
  5 14.1.1.4 0 msec 4 msec 0 msec
  6 45.1.1.5 0 msec 0 msec 0 msec
  7 25.1.1.2 4 msec 0 msec 0 msec
  8 12.1.1.1 0 msec 0 msec 0 msec
  9 14.1.1.4 4 msec 0 msec 0 msec
 10 45.1.1.5 0 msec 4 msec 0 msec
 11 25.1.1.2 0 msec 0 msec 0 msec
 12 12.1.1.1 4 msec 0 msec 0 msec
 13 14.1.1.4 0 msec 0 msec 4 msec
 14 45.1.1.5 0 msec 0 msec 4 msec
 15 25.1.1.2 0 msec 0 msec 0 msec
 16 12.1.1.1 0 msec 0 msec 0 msec
 17 14.1.1.4 0 msec 0 msec 4 msec
 18 45.1.1.5 0 msec 0 msec 4 msec
 19 25.1.1.2 0 msec 0 msec 0 msec
 20 12.1.1.1 4 msec 0 msec 0 msec
 21 14.1.1.4 4 msec 0 msec 0 msec
 22 45.1.1.5 0 msec 4 msec 0 msec
 23 25.1.1.2 0 msec 0 msec 0 msec
 24 12.1.1.1 0 msec 4 msec 0 msec
 25 14.1.1.4 0 msec 4 msec 0 msec
 26 45.1.1.5 0 msec 4 msec 0 msec
 27 25.1.1.2 0 msec 4 msec 0 msec
 28 12.1.1.1 0 msec 4 msec 0 msec
 29 14.1.1.4 0 msec 4 msec 0 msec
 30 45.1.1.5 4 msec 0 msec 0 msec
```

Yes.....the routing loop is back again.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's configure the Eigrp routers with identical router-ids and clear the routing table of all routers using the "Clear ip route *" command:

On R1 and R4:

```
Rx(config)#Router eigrp 100
Rx(config-router)#Eigrp router-id 1.1.1.1
```

On All Routers:

```
Rx#Clear ip route *
```

To verify and test the configuration:

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/52/60 ms

R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 4 msec 0 msec 4 msec
  2 23.1.1.3 16 msec 32 msec 32 msec

R1#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/6] via 12.1.1.2, 00:00:10, FastEthernet0/1
```

Task 7

Remove the solution from the previous task and use another solution to fix the routing loop.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1 and R4:

```
Rx(config)#Router eigrp 100
Rx(config-router)#No eigrp router-id 1.1.1.1
```

On All Routers:

```
Rx#Clear ip eigrp neighbor
Rx#Clear ip route *
```

To verify the configuration:

On R1:

```
R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3

  1 14.1.1.4 0 msec 0 msec 0 msec
  2 45.1.1.5 0 msec 4 msec 0 msec
  3 25.1.1.2 0 msec 0 msec 0 msec
  4 12.1.1.1 0 msec 4 msec 0 msec
  5 14.1.1.4 0 msec 0 msec 0 msec
  6 45.1.1.5 0 msec 4 msec 0 msec
  7 25.1.1.2 0 msec 0 msec 0 msec
  8 12.1.1.1 4 msec 0 msec 0 msec
  9 14.1.1.4 0 msec 0 msec 0 msec
 10 45.1.1.5 4 msec 0 msec 0 msec
 11 25.1.1.2 0 msec 4 msec 0 msec
 12 12.1.1.1 0 msec 0 msec 0 msec
 13 14.1.1.4 4 msec 0 msec 0 msec
 14 45.1.1.5 0 msec 4 msec 0 msec
 15 25.1.1.2 0 msec 0 msec 4 msec
 16 12.1.1.1 0 msec 0 msec 0 msec
 17 14.1.1.4 4 msec 0 msec 0 msec
 18 45.1.1.5 0 msec 4 msec 0 msec
 19 25.1.1.2 0 msec 4 msec 0 msec
 20 12.1.1.1 0 msec 0 msec 4 msec
 21 14.1.1.4 0 msec 0 msec 0 msec
 22 45.1.1.5 0 msec 0 msec 4 msec
 23 25.1.1.2 0 msec 0 msec 0 msec
 24 12.1.1.1 4 msec 0 msec 0 msec
 25 14.1.1.4 4 msec 0 msec 0 msec
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 26 45.1.1.5 4 msec 0 msec 0 msec
 27 25.1.1.2 4 msec 0 msec 0 msec
 28 12.1.1.1 4 msec 0 msec 0 msec
 29 14.1.1.4 4 msec 0 msec 0 msec
 30 45.1.1.5 0 msec 4 msec 0 msec
```

Another way to perform this redistribution is to ONLY allow the required routes to be redistributed, in this case, on R1 and R4 we ONLY need to redistribute the Eigrp route/s into RIP and NO other route/s.

On R1 and R4:

```
Rx(config)#IP Prefix-list NET deny 12.1.1.0/24
Rx(config)#IP Prefix-list NET deny 25.1.1.0/24
Rx(config)#IP Prefix-list NET deny 45.1.1.0/24
Rx(config)#IP Prefix-list NET deny 3.3.3.0/24
Rx(config)#IP Prefix-list NET permit 0.0.0.0/0 le 32

Rx(config)#Route-map tst permit 10
Rx(config-route-map)#Match ip addr prefix NET

Rx(config)#Router rip
Rx(config-router)#redistribute eigrp 100 route-map tst metric 1
```

On All Routers:

```
Rx#Clear ip route *
```

You may have to save and reload R1 and R4 on this IOS to see the result.

To verify and test the configuration:

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/52/60 ms
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 8

Remove the solution from the previous task and use route tags to fix the routing loop.

On R1 and R4:

```
Rx(config)#No IP Prefix-list NET

Rx(config)#No Route-map tst

Rx(config)#Router rip
Rx(config-router)#No redistribute eigrp 100
Rx(config-router)#redistribute eigrp 100 metric 1
```

On All Routers:

```
Rx#Clear ip route *
```

To verify and test the configuration:

On R1:

```
R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3

  1 14.1.1.4 0 msec 0 msec 0 msec
  2 45.1.1.5 4 msec 0 msec 0 msec
  3 25.1.1.2 0 msec 0 msec 4 msec
  4 12.1.1.1 0 msec 0 msec 0 msec
  5 14.1.1.4 0 msec 0 msec 4 msec
  6 45.1.1.5 0 msec 0 msec 0 msec
  7 25.1.1.2 0 msec 4 msec 0 msec
  8 12.1.1.1 0 msec 0 msec 0 msec
  9 14.1.1.4 0 msec 4 msec 0 msec
 10 45.1.1.5 0 msec 0 msec 4 msec
 11 25.1.1.2 0 msec 0 msec 0 msec
 12 12.1.1.1 0 msec 0 msec 0 msec
 13 14.1.1.4 0 msec 0 msec 4 msec
 14 45.1.1.5 0 msec 0 msec 0 msec
 15 25.1.1.2 4 msec 0 msec 0 msec
 16 12.1.1.1 0 msec 4 msec 0 msec
 17 14.1.1.4 0 msec 0 msec 4 msec
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 18 45.1.1.5 0 msec 0 msec 4 msec
 19 25.1.1.2 0 msec 0 msec 0 msec
 20 12.1.1.1 4 msec 0 msec 0 msec
 21 14.1.1.4 0 msec 4 msec 0 msec
 22 45.1.1.5 0 msec 4 msec 0 msec
 23 25.1.1.2 0 msec 4 msec 0 msec
 24 12.1.1.1 0 msec 4 msec 0 msec
 25 14.1.1.4 0 msec 4 msec 0 msec
 26 45.1.1.5 0 msec 4 msec 0 msec
 27 25.1.1.2 0 msec 4 msec 0 msec
 28 12.1.1.1 0 msec 4 msec 0 msec
 29 14.1.1.4 0 msec 4 msec 0 msec
 30 45.1.1.5 0 msec 4 msec 0 msec
```

Let's configure this task using tags:

To avoid feed back and routing loops when performing multipoint mutual redistribution, the routes originating in local routing domain are denied back into the local routing domain, let's say a mutual redistribution is being performed between protocols "A" and "B". The routes from protocol "A" are tagged with a value, let's say 111, and the routes from protocol "B" are tagged with a value, let's say 222. When the redistribution is performed in a mutual manner, the routing table of protocols "A" and "B" will consist of routes that have a tag value of 111, and 222. Therefore, when redistributing protocol "B" into protocol "A", any route that carries a tag 111 is denied back in, and when redistributing protocol "A" into protocol "B", any route that carries a tag of 222 is denied back in. Let's configure this task and verify the result.

On R1 and R4:

The best way to understand the following route-maps is to read them from buttom to top. The "RIP→EIGRP" route-map, sequence number 90, all RIP routes are tagged with a value of 120. In the "EIGRP→RIP route-map" sequence 10, these routes are denied. The same philosophy is used going the other direction.

```
Rx(config)#Route-map RIP->EIGRP deny 10
Rx(config-route-map)#match tag 90
Rx(config)#Route-map RIP->EIGRP permit 90
Rx(config-route-map)#set tag 120

Rx(config)#Route-map EIGRP->RIP DENY 10
Rx(config-route-map)#match tag 120
Rx(config)#Route-map EIGRP->RIP permit 90
Rx(config-route-map)#set tag 90
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The route-map "EIGRP->RIP" is referenced when redistributing EIGRP into RIP, and the route-map "RIP->EIGRP" is referenced when RIP routes are redistributed into EIGRP:

```
Rx(config)#Router rip
Rx(config-router)#redistribute eigrp 100 route-map EIGRP->RIP metric 1

Rx(config)#Router eigrp 100
Rx(config-router)#redistribute rip route-map RIP->EIGRP metric 1 1 1 1 1
```

To verify the configuration:

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/60 ms

R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 4 msec 0 msec 4 msec
  2 23.1.1.3 16 msec 12 msec 12 msec
```

Task 9

Remove the previous solution and use another method to fix the routing loop, do not use prefix-list, access-list, or a route-map to fix the routing loop.

On R1 and R4

```
Rx(config)#No Route-map RIP->EIGRP deny 10
Rx(config)#No Route-map EIGRP->RIP DENY 10

Rx(config)#Router rip
Rx(config-router)#No redistribute eigrp 100 route-map EIGRP->RIP
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Rx(config)#Router eigrp 100
Rx(config-router)#No redistribute rip route-map RIP->EIGRP
```

To see the running configuration of router R1 and R4:

On R1:

```
R1#Show run | S router rip
router rip
 version 2
 redistribute eigrp 100 metric 1
 network 12.0.0.0
 no auto-summary

R1#Show run | S router eigrp 100
router eigrp 100
 redistribute rip metric 1 1 1 1 1
 network 14.1.1.1 0.0.0.0
 no auto-summary
```

On R4:

```
R4#Show run | S router rip

router rip
 version 2
 redistribute eigrp 100 metric 1
 network 45.0.0.0
 no auto-summary

R4#Show run | S router eigrp 100

router eigrp 100
 redistribute rip metric 1 1 1 1 1
 network 14.1.1.4 0.0.0.0
 no auto-summary
```

To verify the configuration:

On R1:

```
R1#Traceroute 3.3.3.3

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.4 0 msec 0 msec 0 msec
  2 45.1.1.5 4 msec 0 msec 0 msec
  3 25.1.1.2 0 msec 0 msec 4 msec
  4 12.1.1.1 0 msec 0 msec 0 msec
  5 14.1.1.4 0 msec 0 msec 0 msec
  6 45.1.1.5 4 msec 0 msec 0 msec
  7 25.1.1.2 0 msec 4 msec 0 msec
  8 12.1.1.1 0 msec 0 msec 0 msec
  9 14.1.1.4 0 msec 4 msec 0 msec
 10 45.1.1.5 0 msec 0 msec 4 msec
 11 25.1.1.2 0 msec 0 msec 0 msec
 12 12.1.1.1 0 msec 4 msec 0 msec
 13 14.1.1.4 0 msec 0 msec 0 msec
 14 45.1.1.5 4 msec 0 msec 0 msec
 15 25.1.1.2 4 msec 0 msec 0 msec
 16 12.1.1.1 0 msec 4 msec 0 msec
 17 14.1.1.4 0 msec 0 msec 4 msec
 18 45.1.1.5 0 msec 0 msec 4 msec
 19 25.1.1.2 0 msec 0 msec 0 msec
 20 12.1.1.1 4 msec 0 msec 0 msec
 21 14.1.1.4 0 msec 4 msec 0 msec
 22 45.1.1.5 0 msec 4 msec 0 msec
 23 25.1.1.2 0 msec 4 msec 0 msec
 24 12.1.1.1 0 msec 4 msec 0 msec
 25 14.1.1.4 0 msec 0 msec 4 msec
 26 45.1.1.5 0 msec 4 msec 0 msec
 27 25.1.1.2 0 msec 4 msec 0 msec
 28 12.1.1.1 0 msec 4 msec 0 msec
 29 14.1.1.4 0 msec 4 msec 0 msec
 30 45.1.1.5 0 msec 4 msec 0 msec
```

We can use route summarization to fix this problem.

On R1 and R4:

```
Rx(config)#Int F0/0
Rx(config-if)#Ip summary-address eigrp 100 3.0.0.0 255.0.0.0

R1#Clear ip route *
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms

R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)

  1 12.1.1.2 0 msec 0 msec 0 msec
  2 23.1.1.3 32 msec 24 msec 24 msec
```

To see how the summarization works, the routing table of all routers are displayed:

On R1:

```
R1#Show ip route | Inc 3.3.3.0

R        3.3.3.0/24 [120/6] via 12.1.1.2, 00:00:16, FastEthernet0/1
```

On R2:

```
R2#Show ip route | Inc 3.3.3.0

R        3.3.3.0/24 [120/5] via 23.1.1.3, 00:00:00, Serial1/3
```

On R4:

```
R4#Show ip route | Inc 3.3.3.0

R        3.3.3.0/24 [120/7] via 45.1.1.5, 00:00:11, FastEthernet0/1
```

On R5:

```
R5#Show ip route | Inc 3.3.3.0

R        3.3.3.0/24 [120/6] via 25.1.1.2, 00:00:07, FastEthernet0/0
```

The output of the above show commands clearly states that 3.3.3.0/24 is reachable through RIP and NOT

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Eigrp. But how does this work and what if there is another router in the Eigrp AS 100, will this router see the network and have reachability to 3.3.3.3?
Let's add R6 to the Ethernet segment that connects R1 to R4 and run Eigrp on the F0/0 interface of this router, and verify the routing table:

On R6:

```
R6(config)#Int F0/0
R6(config-if)#IP address 14.1.1.6 255.255.255.0
R6(config-if)#No shut

R6(config)#Router eigrp 100
R6(config-router)#Network 14.1.1.6 0.0.0.0
```

On SW1:

```
SW1(config)#Int f0/6
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc v 14
SW1(config-if)#Spannin portf
SW1(config-if)#No shut
```

To verify the configuration:

On R6:

```
R6#Show ip route 3.0.0.0

Routing entry for 3.0.0.0/8
  Known via "eigrp 100", distance 90, metric 2560002816, type internal
  Redistributing via eigrp 100
  Last update from 14.1.1.1 on FastEthernet0/0, 00:00:45 ago
  Routing Descriptor Blocks:
  * 14.1.1.4, from 14.1.1.4, 00:00:45 ago, via FastEthernet0/0
      Route metric is 2560002816, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
    14.1.1.1, from 14.1.1.1, 00:00:45 ago, via FastEthernet0/0
      Route metric is 2560002816, traffic share count is 1
      Total delay is 110 microseconds, minimum bandwidth is 1 Kbit
      Reliability 1/255, minimum MTU 1 bytes
      Loading 1/255, Hops 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To test the configuration:

On R6:

```
R6#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

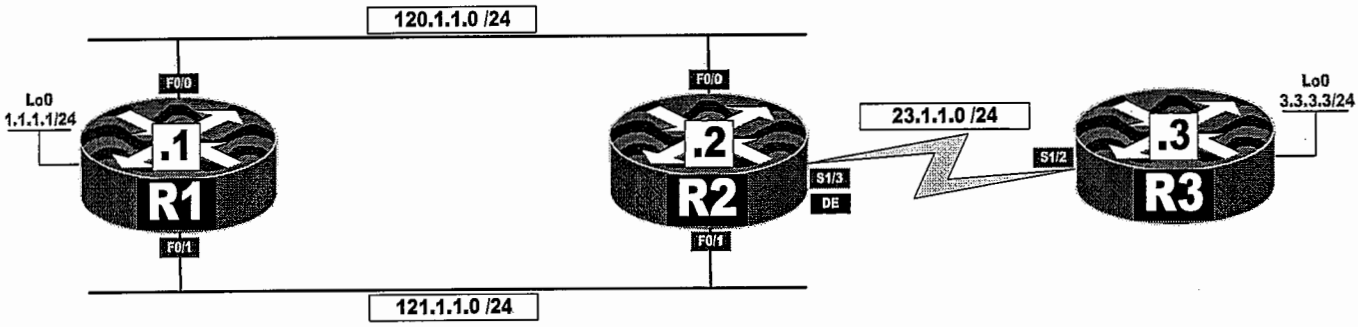
As you can see R6 does NOT have 3.3.3.0/24 in its routing table and instead it has the summary route.

Task 11

Erase the startup configuration of the routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 4
Redistribution – RIPv2 & OSPF



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "Redistribution"→ "Lab-4".

Task 1

Configure OSPF area 0 on the F0/0 interfaces of R1 and R2, R1 should also run OSPF area 0 on its Lo0 interface. This lo0 should be advertised with its correct mask. Configure any router-id for these two routers, DO NOT use an IP address to accomplish this task.

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Net 120.1.1.1 0.0.0.0 area 0
R1(config-router)#Net 1.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#Router ospf 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Net 120.1.1.2 0.0.0.0 area 0
```

To verify the configuration:

On R2:

```
R2#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 [110/2] via 120.1.1.1, 00:00:22, FastEthernet0/0
```

Task 2

Configure RIPv2 on R1, R2 and R3 based on the following requirement:

- R1 – Configure RIPv2 on the F0/1 interface.
- R2 – Configure RIPv2 on the F0/1, and the S1/3 interfaces.
- R3 – Configure RIPv2 on the S1/2 and the Loopback0 interfaces.
- Disable auto-summarization on all three routers.

On R1:

```
R1(config)#Router rip
R1(config-router)#ver 2
R1(config-router)#No au
R1(config-router)#Network 121.1.0.0
```

On R2:

```
R2(config)#Router rip
R2(config-router)#ver 2
R2(config-router)#No au
R2(config-router)#Network 121.0.0.0
R2(config-router)#Network 23.0.0.0
```

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config)#Router rip
R3(config-router)#ver 2
R3(config-router)#No au
R3(config-router)#Network 3.0.0.0
R3(config-router)#Network 23.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:08, FastEthernet0/1
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 121.1.1.2, 00:00:24, FastEthernet0/1
```

On R2:

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

      3.0.0.0/24 is subnetted, 1 subnets
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:06, Serial1/3
```

Task 3

Configure the Update, Invalidation, and Flush timer to be 1/3rd of the default value.

On All Routers:

```
R1#Show ip protocol | Inc seconds

  Sending updates every 30 seconds, next due in 1 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
```

On All Routers:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Rx(config)#Router rip
Rx(config-router)#Timers basic 10 60 60 80
```

To verify the configuration:

On All Routers:

```
R1#Show ip protocol | Inc seconds

  Sending updates every 10 seconds, next due in 3 seconds
  Invalid after 60 seconds, hold down 60, flushed after 80
```

Task 4

Perform a mutual redistribution between OSPF and RIPv2 on R1. R1 should redistribute the ospf routes into RIPv2 with a metric of 1, whereas, RIP routes should be redistributed into OSPF using the default cost and metric type. Ensure full reachability using ping.

On R1

```
R1(config)#Router rip
R1(config-router)#redistribute ospf 1 metric 1

R1(config)#Router ospf 1
R1(config-router)#redistribute rip subnets
```

To verify the configuration:

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's see if network 3.3.3.0/24 is in R1's routing table:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:09, FastEthernet0/1
```

It's there, let's ping again:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

NOTE: If you don't get the same result try again few times and you should see the same result.

Do we have a routing loop? But we redistributed on a single router. Let's verify by entering the "Show ip route | Inc 3.3.3.0" and then pressing the up arrow and then enter in rapid succession to see the problem. Let's do this on R1.

On R1:

```
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:07, FastEthernet0/1
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:07, FastEthernet0/1
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:08, FastEthernet0/1
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:08, FastEthernet0/1
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:08, FastEthernet0/1
R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:09, FastEthernet0/1

R1#Show ip route | Inc 3.3.3.0

R1#Show ip route | Inc 3.3.3.0
R1#Show ip route | Inc 3.3.3.0
R1#Show ip route | Inc 3.3.3.0

R1#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/2] via 121.1.1.2, 00:00:00, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The route shows up as a RIP route and then it disappears and then comes back as a RIP route again, let's try this on R2:

On R2:

```
R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:00, FastEthernet0/0
R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:04, FastEthernet0/0
R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:06, FastEthernet0/0
R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:07, FastEthernet0/0
R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:08, FastEthernet0/0

R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0
R2#Show ip route | Inc 3.3.3.0

R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:00, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:01, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:02, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:03, Serial1/3
R2#Show ip route | Inc 3.3.3.0

R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:04, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:04, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:05, Serial1/3
R2#Show ip route | Inc 3.3.3.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:06, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:07, Serial1/3
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:08, Serial1/3

R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:00, FastEthernet0/0
```

Wow, on this router, it shows up as an OSPF "E2" route and then goes away and comes back as RIP and then the cycle repeats.
Let's enable some debugging to resolve this problem. Let's enable "Debug ip routing" and "Debug ip rip"on R2 and "Debug ip ospf lsa-generation" on R1.

On R1:

```
R1#Debug ip ospf lsa-generation
OSPF summary lsa generation debugging is on
```

On R2:

```
R3#Debug ip routing
IP routing debugging is on

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

Let's analyze the output of the debug commands:

Step #1:

On R2:

Network 3.3.3.0/24 is advertised to R2 by R3 (23.1.1.3):

```
RT: updating rip 3.3.3.0/24 (0x0):
    via 23.1.1.3 Se1/3

RT: add 3.3.3.0/24 via 23.1.1.3, rip metric [120/1]
```

Let's check the routing table of R2:

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:08, Serial1/3
```

Once R2 receives network 3.3.3.0/24, it advertises the network to R1.

Step #2:

On R1:

R1 receives network 3.3.3.0/24 from R1 and redistributes it into OSPF, and it sends an LSA Type-5 to R2:

```
OSPF-1 LSGEN: Generate external LSA 3.3.3.0, mask 255.255.255.0, type 5, age 0,
seq 0x80000013
OSPF-1 LSGEN: MTID      Metric      Metric-type   FA      Tag      Topology Name
OSPF-1 LSGEN: 0         20          2             0.0.0.0 0        Base
```

To see this on R1:

```
R1#Show ip ospf da external | Inc 3.3.3.0
  Link State ID: 3.3.3.0 (External Network Number )
```

Step #3:

R2 receives the OSPF external update from R1, and because OSPF has a lower administrative distance, it flushes the RIP learned route and replaces it with OSPF with a next-hop IP address of 120.1.1.1.

```
RT: rib update return code: 17
RT: updating ospf 3.3.3.0/24 (0x0):
    via 120.1.1.1 Fa0/0

RT: closer admin distance for 3.3.3.0, flushing 1 routes
RT: add 3.3.3.0/24 via 120.1.1.1, ospf metric [110/20]
```

To see this on R2:

```
R2#Show ip route | Inc 3.3.3.0
O E2     3.3.3.0 [110/20] via 120.1.1.1, 00:00:00, FastEthernet0/0
```

Step #4:

Since R2 flushed the RIP learned (3.3.3.0/24) network and replaced it with OSPF, it poisons network 3.3.3.0/24 and tells R1 that network 3.3.3.0/24 is inaccessible meaning that it's 16 hops away from the local router. Once R1 gets that message, it replies back saying that network 3.3.3.0/24 is no longer

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

reachable and poison reverses the network. This is the output from the "Debug ip rip":

The following is being generated by R2 and sent to R1:

```
RIP: sending v2 update to 224.0.0.9 via FastEthernet0/1 (121.1.1.2)
RIP: build update entries
        3.3.3.0/24 via 0.0.0.0, metric 16, tag 0
        23.1.1.0/24 via 0.0.0.0, metric 1, tag 0
```

The following is being generated by R1 and sent to R2:

```
RIP: received v2 update from 121.1.1.1 on FastEthernet0/1
     3.3.3.0/24 via 0.0.0.0 in 16 hops  (inaccessible)
```

Step #5:

Since R1 does not have network 3.3.3.0/24 in its routing table it can no longer redistribute it into OSPF, so it advertises network 3.3.3.0/24 with its maximum metric poisoning the route.

```
OSPF-1 LSGEN: Generate external LSA 3.3.3.0, mask 255.255.255.0, type 5, age
3600, seq 0x80000016
OSPF-1 LSGEN: MTID    Metric     Metric-type    FA        Tag        Topology Name
OSPF-1 LSGEN: 0       16777215   2              0.0.0.0   0            Base
```

This is when R2 removes the route as an OSPF "E2" from its routing table, so because it removed the network as a RIP and OSPF it no longer has network 3.3.3.0/24 in its routing table:

```
R2#Show ip route | Inc 3.3.3.0
R2#
```

But few seconds after that R2 receives a RIP update from R3 and it injects it into its routing table and the cycle repeats.

```
R2#Show ip route | Inc 3.3.3.0
R        3.3.3.0 [120/1] via 23.1.1.3, 00:00:04, Serial1/3
```

To fix this problem, R2 should always trust RIP's information for network 3.3.3.0/24 over OSPF. Let's change the administrative distance of RIP to be lower than OSPF's admin distance, once this configuration is performed, R3 will ignore OSPF's update regarding network 3.3.3.0/24.

On R3

```
R3(config)#Router rip
R3(config-router)#Distance 109
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To test and verify the configuration:

On R1:

```
R1#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
```

On R2:

```
R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

Task 5

Erase the startup configuration of the routers and switches and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Advanced
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

BGP:

Lab-1:
Establishing Neighbor Adjacency

Lab-2:
Route Reflectors

Lab-3:
Conditional Advertisement & BGP Backdoor

Lab-4:
Route Dampening

Lab-5:
BGP Aggregation

Lab-6:
The Community Attribute

Lab-7:
BGP Cost Community

Lab-8:
BGP Load Balancing

Lab-9:
Local-Preference –I

Lab-10:
BGP Local-Preference – II

Lab-11:
The AS-Path Attribute

Lab-12:
The Weight Attribute

Lab-13:
Multi Exit Discriminator Attribute

Lab-14:
BGP Filtering

Lab-15:
Administrative Distance

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab-16:
Advanced BGP Configurations

Lab-17:
BGP Confederation

Lab-18:
BGP Hiding Local AS number

Lab-19:
BGP Allow-AS

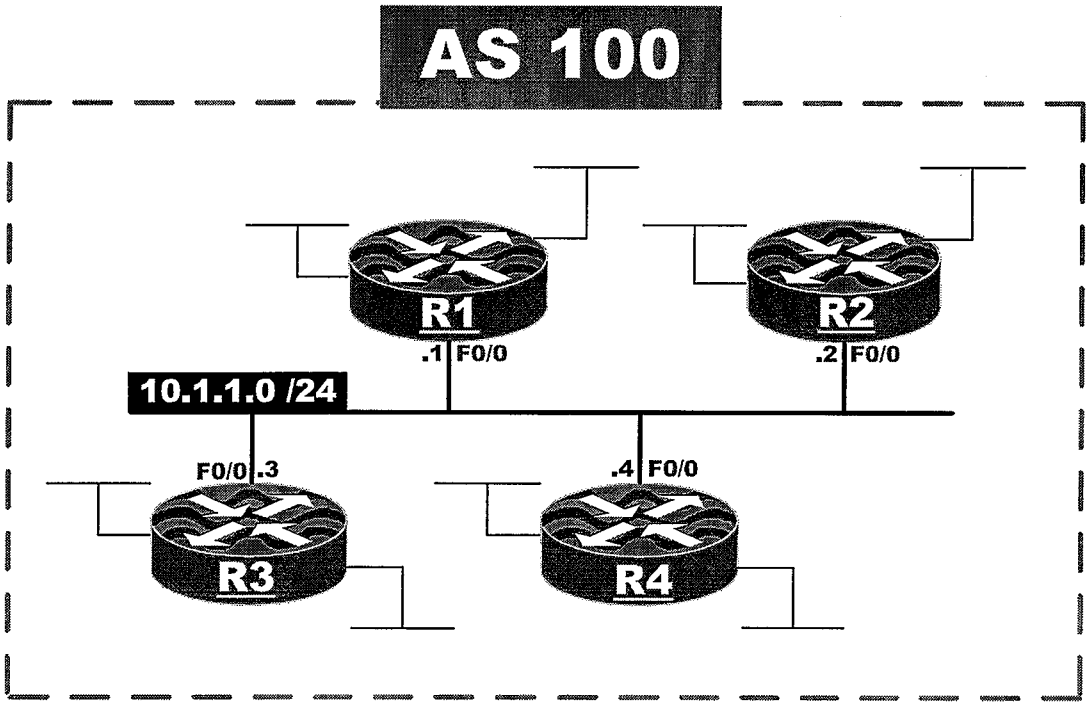
Lab-20:
BGP Automatic-tag, AS-path tag and Table-map

Lab-21:
4 Byte AS Numbers

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1 – Establishing Neighbor Adjacency



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-1"→"Task-1-4" →"Lab1-T1-4".

ALL BGP LABS ARE CONDUCTED ON THE FOLLOWING IOS ON 2811 SERIES ROUTERS:

c2800nm-adventerprisek9-mz.151-4.M7.bin

Task 1

Configure routers R1 – R4 in AS 100, these routers should create an IBGP peer sessions between them in a full mesh manner, ensure that these routers advertise their loopback 0 interface in this AS.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote-as 100
R1(config-router)#neighbor 10.1.1.3 remote-as 100
R1(config-router)#neighbor 10.1.1.4 remote-as 100
R1(config-router)#No syn
R1(config-router)#network 1.0.0.0
```

On R2:

```
R2(config)#router bgp 100
R2(config-router)#neighbor 10.1.1.1 remote-as 100
R2(config-router)#neighbor 10.1.1.3 remote-as 100
R2(config-router)#neighbor 10.1.1.4 remote-as 100
R2(config-router)#No syn
R2(config-router)#network 2.0.0.0
```

You should see the following console message stating that an adjacency is established with neighbor 10.1.1.1:

```
%BGP-5-ADJCHANGE: neighbor 10.1.1.1 Up
```

On R3:

```
R3(config)#router bgp 100
R3(config-router)#neighbor 10.1.1.1 remote-as 100
R3(config-router)#neighbor 10.1.1.2 remote-as 100
R3(config-router)#neighbor 10.1.1.4 remote-as 100
R3(config-router)#No syn
R3(config-router)#network 3.0.0.0
```

You should see the following console messages:

```
%BGP-5-ADJCHANGE: neighbor 10.1.1.1 Up
%BGP-5-ADJCHANGE: neighbor 10.1.1.2 Up
```

On R4:

```
R1(config)#router bgp 100
R4(config-router)#neighbor 10.1.1.1 remote-as 100
R4(config-router)#neighbor 10.1.1.2 remote-as 100
R4(config-router)#neighbor 10.1.1.3 remote-as 100
R4(config-router)#No syn
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config-router)#network 4.0.0.0
```

You should see the following console messages:

```
%BGP-5-ADJCHANGE: neighbor 10.1.1.1 Up
%BGP-5-ADJCHANGE: neighbor 10.1.1.3 Up
%BGP-5-ADJCHANGE: neighbor 10.1.1.2 Up
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	0.0.0.0	0		32768	i
*>i2.0.0.0	10.1.1.2	0	100	0	i
*>i3.0.0.0	10.1.1.3	0	100	0	i
*>i4.0.0.0	10.1.1.4	0	100	0	i

The following explains the output of the above show command:

- The "*" means that the entry in the table is valid.
- The ">" means that this is the best entry for a given prefix.
- The "i" means that the entry was learned via an IBGP, this is the "i" to the right of the greater than sign (>) in the network column. The letter "i" under the path column, specifies the origin of the route.
- The "Network" column identifies the prefix entry for the network/s, if the mask is omitted, the default mask is assumed.
- The "Next Hop" column identifies the next hop's IP address to reach the specified network address, if its set to "0.0.0.0", it identifies a prefix that the local router has advertised.
- The "Metric" column identifies the inter-as metric, or the "MED" attribute which is set to zero by default.
- The "LocPrf" column identifies the "local preference" attribute; this is used in the route selection process carried within the local AS ONLY. With the local-pref attribute, the higher value has more preference. The prefixes that are received from a peer AS are tagged with a local-pref value of 100; this is the default behavior and this value can be changed to influence the best path selection

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

process. The changed value is only advertised to IBGP peers. When the local router advertises a prefix, no local-pref value is seen in the output of the "Show ip bgp" command. The default value of 100 can be changed by the "BGP default local-preference" command.

- The "Weight" column identifies the weight attribute; the prefixes that are received via a neighbor (IBGP or EBGP) will have a weight attribute of 0, but the prefixes that are originated by the local router will have a weight value of 32768. This attribute overrides any other attribute for performing best path determination.

- The "Path" column identifies the path attribute. If the prefixes were originated or learned via an IBGP neighbor, the path column will have the letter "i" without any ASN. If the prefix was learned through another AS, then this column will have the AS number/s followed by the letter "i", the ASNs indicate which AS advertised the prefix. The maximum number of ASes that a prefix can traverse through is 255.

Task 2

Reconfigure the routers based on the following policy:

R2, R3 and R4 should be configured in AS 200, 300 and 400 respectively. Configure a full mesh peer session between these routers.

On All Routers:

```
Rx(config)#No Router bgp 100
```

On R1:

```
R1(config)#router bgp 100
R1(config-router)#neighbor 10.1.1.2 remote-as 200
R1(config-router)#neighbor 10.1.1.3 remote-as 300
R1(config-router)#neighbor 10.1.1.4 remote-as 400
R1(config-router)#No auto
R1(config-router)#network 1.0.0.0
```

On R2:

```
R2(config)#router bgp 200
R2(config-router)#neighbor 10.1.1.1 remote-as 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#neighbor 10.1.1.3 remote-as 300
R2(config-router)#neighbor 10.1.1.4 remote-as 400
R2(config-router)#No auto
R2(config-router)#network 2.0.0.0
```

On R3:

```
R3(config)#router bgp 300
R3(config-router)#neighbor 10.1.1.1 remote-as 100
R3(config-router)#neighbor 10.1.1.2 remote-as 200
R3(config-router)#neighbor 10.1.1.4 remote-as 400
R3(config-router)#No auto
R3(config-router)#network 3.0.0.0
```

On R4:

```
R1(config)#router bgp 400
R4(config-router)#neighbor 10.1.1.1 remote-as 100
R4(config-router)#neighbor 10.1.1.2 remote-as 200
R4(config-router)#neighbor 10.1.1.3 remote-as 300
R4(config-router)#No auto
R4(config-router)#network 4.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Net
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          10.1.1.2                               0 400 200 i
*                   10.1.1.2                               0 300 200 i
*>                  10.1.1.2                 0             0 200 i
*  3.0.0.0          10.1.1.3                               0 400 300 i
*                   10.1.1.3                               0 200 300 i
*>                  10.1.1.3                 0             0 300 i
*  4.0.0.0          10.1.1.4                               0 200 400 i
*                   10.1.1.4                               0 300 400 i
*>                  10.1.1.4                 0             0 400 i
```

Note: The local-preference attribute is not assigned on any of the prefixes; this is because the prefixes are advertised by an EBGP peer. The best selection in the above output is based on the shortest ASN.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The MED value (metric column) is zero for some of the prefixes, and it is not set on others, this is because when the prefix is advertised by the originating AS, the metric is set to "0", but when the same prefix is advertised by another AS, the MED value is removed.

NOTE: The next hop attribute does NOT change even though they are EBGP neighbors, this is because of

the Multi-access broadcast network type (Ethernet) where each router can see the other router/s through the same interface.

Task 3

Reconfigure all the routers in AS 100; use the following policy for their IBGP peer sessions:

- Authentication must be enabled between the peers using "**cisco**" as the password.
- The peer session must be established based on the loopback 0's IP address.
- These routers should **only** advertise their loopback 1 in BGP.
- Provide NLRI to loopback0 interface using RIPv2.
- The peer session between the routers should only be established if they are running BGP version 4.
- Use peer-groups to accomplish this task.

Cisco's implementation of BGP in IOS 12.0(5)T or earlier releases supports BGP versions 2, 3, and 4, with dynamic negotiation down to Version 2. However, in IOS version 12.0(6)T or later, Cisco routers only support version 4 and they do not support dynamic negotiation down to Version 2. The reason you may see the "neighbor version" command configured on some Cisco routers is because may be the router is connecting and establishing a peer session with a Non-Cisco router, or the administrator is not aware of this fact.

On R1:

```
R1(config)#No Router bgp 100

R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Net 192.168.1.0
R1(config-router)#Neighbor TST peer-group
R1(config-router)#Neighbor TST remote-as 100
```


© 2014 Narbik Kocharians, All rights reserved

https://t.me/learningnets

```
R1(config-router)#Neighbor TST update-source lo0
R1(config-router)#Neighbor TST version 4
R1(config-router)#Neighbor TST password cisco

R1(config-router)#Neighbor 2.2.2.2 peer-group TST
R1(config-router)#Neighbor 3.3.3.3 peer-group TST
R1(config-router)#Neighbor 4.4.4.4 peer-group TST

R1(config-router)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Netw 10.0.0.0
R1(config-router)#Netw 1.0.0.0
```

On R2:

```
R2(config)#No Router bgp 200

R2(config)#Router bgp 100
R2(config-router)#No au
R2(config-router)#Netw 192.168.2.0
R2(config-router)#Neighbor TST peer-group
R2(config-router)#Neighbor TST remote-as 100
R2(config-router)#Neighbor TST update-source lo0
R2(config-router)#Neighbor TST version 4
R2(config-router)#Neighbor TST password cisco

R2(config-router)#Neighbor 1.1.1.1 peer-group TST
R2(config-router)#Neighbor 3.3.3.3 peer-group TST
R2(config-router)#Neighbor 4.4.4.4 peer-group TST

R2(config-router)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Netw 10.0.0.0
R2(config-router)#Netw 2.0.0.0
```

On R3:

```
R3(config)#No Router bgp 300

R3(config)#Router bgp 100
R3(config-router)#No au
R3(config-router)#Network 192.168.3.0
R3(config-router)#Neighbor TST peer-group
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config-router)#Neighbor TST remote-as 100
R3(config-router)#Neighbor TST update-source lo0
R3(config-router)#Neighbor TST version 4
R3(config-router)#Neighbor TST password cisco

R3(config-router)#Neighbor 1.1.1.1 peer-group TST
R3(config-router)#Neighbor 2.2.2.2 peer-group TST
R3(config-router)#Neighbor 4.4.4.4 peer-group TST

R3(config-router)#Router rip
R3(config-router)#No au
R3(config-router)#Ver 2
R3(config-router)#Netw 10.0.0.0
R3(config-router)#Netw 3.0.0.0
```

On R4:

```
R4(config)#No Router bgp 400

R4(config)#Router bgp 100
R4(config-router)#No au
R4(config-router)#Network 192.168.4.0
R4(config-router)#Neighbor TST peer-group
R4(config-router)#Neighbor TST remote-as 100
R4(config-router)#Neighbor TST update-source lo0
R4(config-router)#Neighbor TST version 4
R4(config-router)#Neighbor TST password cisco

R4(config-router)#Neighbor 1.1.1.1 peer-group TST
R4(config-router)#Neighbor 2.2.2.2 peer-group TST
R4(config-router)#Neighbor 3.3.3.3 peer-group TST

R4(config-router)#Router rip
R4(config-router)#No au
R4(config-router)#Ver 2
R4(config-router)#Netw 10.0.0.0
R4(config-router)#Netw 4.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      0.0.0.0                  0         32768 i
*>i192.168.2.0      2.2.2.2                  0    100      0 i
*>i192.168.3.0      3.3.3.3                  0    100      0 i
*>i192.168.4.0      4.4.4.4                  0    100      0 i
```

To verify the configuration:

On R1:

```
R1#Show ip bgp peer-group TST

BGP peer-group is TST,  remote AS 100
  BGP version 4
  Default minimum time between advertisement runs is 0 seconds
 For address family: IPv4 Unicast
  BGP neighbor is TST, peer-group internal, members:
  2.2.2.2 3.3.3.3 4.4.4.4
  Index 0, Offset 0, Mask 0x0
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
```

Note: The output of the "Show ip bgp peer-group TST" reveals the ip address of the members of the peer-group. One of the major benefits of a peer-group is that it can reduce the administrative overhead by cutting down redundant configuration on the routers.

Task 4

Remove the BGP configuration from the routers and reconfigure all four routers in AS 100 using **peer-session templates**; you should configure the following two templates to accomplish this task:

- Common Template: This template should contain the "**neighbor version 4**" and "**neighbor password**" command, this template should be applied to all neighbors.
- IBGP Template: This template should contain the "**neighbor update-source**" and "**neighbor remote-as**" commands. This template should be applied to all IBGP neighbors.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

You should advertise Loopback1 interface in BGP and loopback0 should be used as the IP address for establishing the peer sessions. **Do not** remove RIPv2's configuration.

On R1:

```
R1(config)#No Router bgp 100

R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Network 192.168.1.0

R1(config-router)#Template peer-session Common
R1(config-router-stmp)#password cisco
R1(config-router-stmp)#version 4
R1(config-router-stmp)#exit-peer-session

R1(config-router)#Template peer-session IBGP
R1(config-router-stmp)#inherit peer-session Common
R1(config-router-stmp)#update-source lo0
R1(config-router-stmp)#remote-as 100
R1(config-router-stmp)#exit-peer-session

R1(config-router)#Neighbor 2.2.2.2 inherit peer-session IBGP
R1(config-router)#Neighbor 3.3.3.3 inherit peer-session IBGP
R1(config-router)#Neighbor 4.4.4.4 inherit peer-session IBGP
```

On R2:

```
R2(config)#No Router bgp 100

R2(config)#Router bgp 100
R2(config-router)#No au
R2(config-router)#Network 192.168.2.0

R2(config-router)#Template peer-session Common
R2(config-router-stmp)#password cisco
R2(config-router-stmp)#version 4
R2(config-router-stmp)#exit-peer-session

R2(config-router)#Template peer-session IBGP
R2(config-router-stmp)#inherit peer-session Common
R2(config-router-stmp)#update-source lo0
R2(config-router-stmp)#remote-as 100
R2(config-router-stmp)#exit-peer-session
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Neighbor 1.1.1.1 inherit peer-session IBGP
R2(config-router)#Neighbor 3.3.3.3 inherit peer-session IBGP
R2(config-router)#Neighbor 4.4.4.4 inherit peer-session IBGP
```

On R3:

```
R3(config)#No Router bgp 100

R3(config)#Router bgp 100
R3(config-router)#No au
R3(config-router)#Network 192.168.3.0

R3(config-router)#Template peer-session Common
R3(config-router-stmp)#password cisco
R3(config-router-stmp)#version 4
R3(config-router-stmp)#exit-peer-session

R3(config-router)#Template peer-session IBGP
R3(config-router-stmp)#inherit peer-session Common
R3(config-router-stmp)#update-source lo0
R3(config-router-stmp)#remote-as 100
R3(config-router-stmp)#exit-peer-session

R3(config-router)#Neighbor 1.1.1.1 inherit peer-session IBGP
R3(config-router)#Neighbor 2.2.2.2 inherit peer-session IBGP
R3(config-router)#Neighbor 4.4.4.4 inherit peer-session IBGP
```

On R4:

```
R4(config)#No Router bgp 100

R4(config)#Router bgp 100
R4(config-router)#No au
R4(config-router)#Network 192.168.4.0

R4(config-router)#Template peer-session Common
R4(config-router-stmp)#password cisco
R4(config-router-stmp)#version 4
R4(config-router-stmp)#exit-peer-session

R4(config-router)#Template peer-session IBGP
R4(config-router-stmp)#inherit peer-session Common
R4(config-router-stmp)#update-source lo0
R4(config-router-stmp)#remote-as 100
R4(config-router-stmp)#exit-peer-session
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config-router)#Neighbor 1.1.1.1 inherit peer-session IBGP
R4(config-router)#Neighbor 2.2.2.2 inherit peer-session IBGP
R4(config-router)#Neighbor 3.3.3.3 inherit peer-session IBGP
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 192.168.1.0	0.0.0.0	0		32768	i
*>i192.168.2.0	2.2.2.2	0	100	0	i
*>i192.168.3.0	3.3.3.3	0	100	0	i
*>i192.168.4.0	4.4.4.4	0	100	0	i

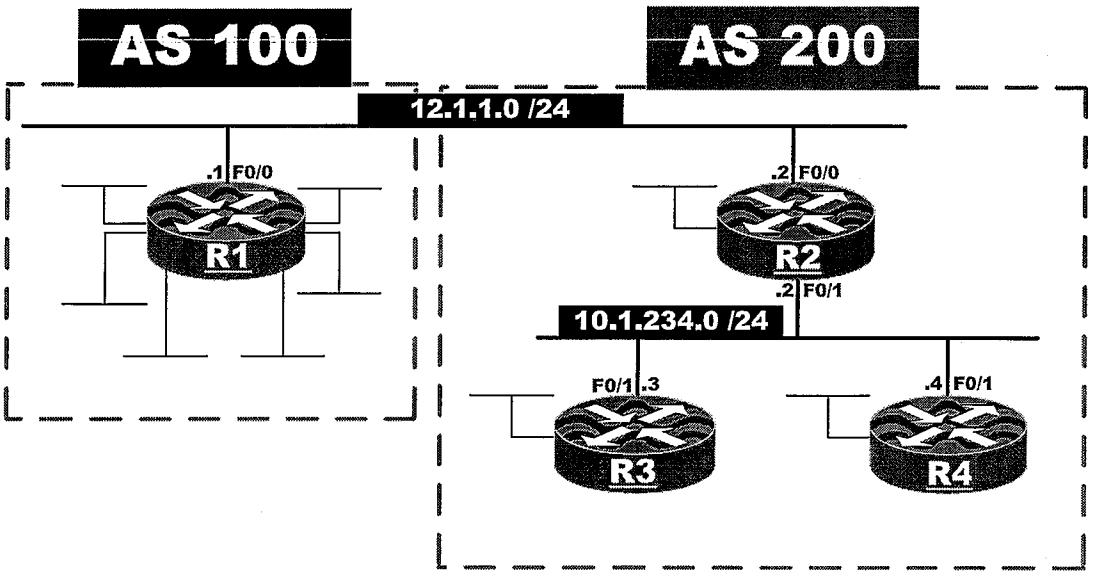
- Peer-session template can be used to apply session specific configuration command/s to a group of neighbors that share a common session configuration.
- Peer-session templates can be reused and they support inheritance of another peer-session template/s, this means that nested peer-sessions can also be configured.
- Peer-session templates support session specific commands only.

Task 5

Reconfigure the routers based on the following IP address space and diagram.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-1"→"Task-5-7" →"Lab1-T5-7".

Task 6

- R1 in AS 100 should establish an EBGP peer session with R2 in AS 200. R1 should advertise all of its loopback interfaces in AS 100.

- R2, R3 and R4 should be configured in AS 200; these routers should establish IBGP peer sessions between them and advertise their loopback 0 interface in AS 200.
- Configure the router-ids of the routers as follows:

  R1 = 10.1.1.1, R2 = 10.2.2.2, R3 = 10.3.3.3 and R4 = 10.4.4.4

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No synchronization
R1(config-router)#Bgp router-id 10.1.1.1
R1(config-router)#Network 1.1.0.0 mask 255.255.255.0
R1(config-router)#Network 1.1.1.0 mask 255.255.255.0
R1(config-router)#Network 1.1.2.0 mask 255.255.255.0
R1(config-router)#Network 1.1.3.0 mask 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-router)#Network 100.1.1.0 mask 255.255.255.0
R1(config-router)#Network 100.2.2.0 mask 255.255.255.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#No auto-summary
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No synchronization
R2(config-router)#bgp router-id 10.2.2.2
R2(config-router)#Network 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 10.1.234.3 remote-as 200
R2(config-router)#Neighbor 10.1.234.4 remote-as 200
R2(config-router)#No auto-summary
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#No synchronization
R3(config-router)#Bgp router-id 10.3.3.3
R3(config-router)#Network 3.0.0.0
R3(config-router)#Neighbor 10.1.234.2 remote-as 200
R3(config-router)#Neighbor 10.1.234.4 remote-as 200
R3(config-router)#No auto-summary
```

On R4:

```
R4(config)#Router bgp 200
R4(config-router)#No synchronization
R4(config-router)#Bgp router-id 10.4.4.4
R4(config-router)#Network 4.0.0.0
R4(config-router)#Neighbor 10.1.234.2 remote-as 200
R4(config-router)#Neighbor 10.1.234.3 remote-as 200
R4(config-router)#No auto-summary
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*> 1.1.0.0/24        0.0.0.0                  0         32768 i
*> 1.1.1.0/24        0.0.0.0                  0         32768 i
*> 1.1.2.0/24        0.0.0.0                  0         32768 i
*> 1.1.3.0/24        0.0.0.0                  0         32768 i
*> 2.0.0.0           12.1.1.2                 0             0 200 i
*> 3.0.0.0           12.1.1.2                               0 200 i
*> 4.0.0.0           12.1.1.2                               0 200 i
*> 100.1.1.0/24      0.0.0.0                  0         32768 i
*> 100.2.2.0/24      0.0.0.0                  0         32768 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.0.0/24       12.1.1.1                 0             0 100 i
*> 1.1.1.0/24       12.1.1.1                 0             0 100 i
*> 1.1.2.0/24       12.1.1.1                 0             0 100 i
*> 1.1.3.0/24       12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          10.1.234.3               0    100      0 i
*>i4.0.0.0          10.1.234.4               0    100      0 i
*> 100.1.1.0/24     12.1.1.1                 0             0 100 i
*> 100.2.2.0/24     12.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.1.0.0/24       12.1.1.1                 0    100      0 100 i
* i1.1.1.0/24       12.1.1.1                 0    100      0 100 i
* i1.1.2.0/24       12.1.1.1                 0    100      0 100 i
* i1.1.3.0/24       12.1.1.1                 0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.234.4               0    100      0 i
* i100.1.1.0/24     12.1.1.1                 0    100      0 100 i
* i100.2.2.0/24     12.1.1.1                 0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
* i1.1.0.0/24       12.1.1.1                 0    100      0 100 i
* i1.1.1.0/24       12.1.1.1                 0    100      0 100 i
* i1.1.2.0/24       12.1.1.1                 0    100      0 100 i
* i1.1.3.0/24       12.1.1.1                 0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*>i3.0.0.0          10.1.234.3               0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i100.1.1.0/24     12.1.1.1                 0    100      0 100 i
* i100.2.2.0/24     12.1.1.1                 0    100      0 100 i
```

Note: R3 and R4 do not have NLRI to the next-hop IP address of 12.1.1.1, therefore, they will not have reachability to these addresses this is why the greater than sign is NOT in the BGP table.
This reachability can be achieved by R2 injecting a default route, or R2 changing the next hop IP address to its internal IP address (10.1.234.2), or having R2 to advertise the external link (12.1.1.0 /24) in the IGP that it's running in its AS. When the last option is configured, don't forget to set the external link as passive-interface or else AS 100 can run the same IGP and know about your internal networks.

Task 7

Configure R2 to provide reachability to R3 and R4 by changing the next hop IP address for all the networks advertised by R1 to the IP address of its F0/1 interface. You should use a template so the future policies can be installed once in that template and have it effect R3 and R4. **DO NOT** use peer-groups to accomplish this task.

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Template peer-policy TST
R2(config-router-ptmp)#next-hop-self
R2(config-router-ptmp)#exit-peer-policy

R2(config-router)#Neighbor 10.1.234.3 inherit peer-policy TST
R2(config-router)#Neighbor 10.1.234.4 inherit peer-policy TST
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp | b Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.0.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.1.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.2.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.3.0/24       10.1.234.2               0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          10.1.234.4               0    100      0 i
*>i100.1.1.0/24     10.1.234.2               0    100      0 100 i
*>i100.2.2.0/24     10.1.234.2               0    100      0 100 i
```

On R4:

```
R4#Sh ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.0.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.1.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.2.0/24       10.1.234.2               0    100      0 100 i
*>i1.1.3.0/24       10.1.234.2               0    100      0 100 i
*>i2.0.0.0          10.1.234.2               0    100      0 i
*>i3.0.0.0          10.1.234.3               0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i100.1.1.0/24     10.1.234.2               0    100      0 100 i
*>i100.2.2.0/24     10.1.234.2               0    100      0 100 i
```

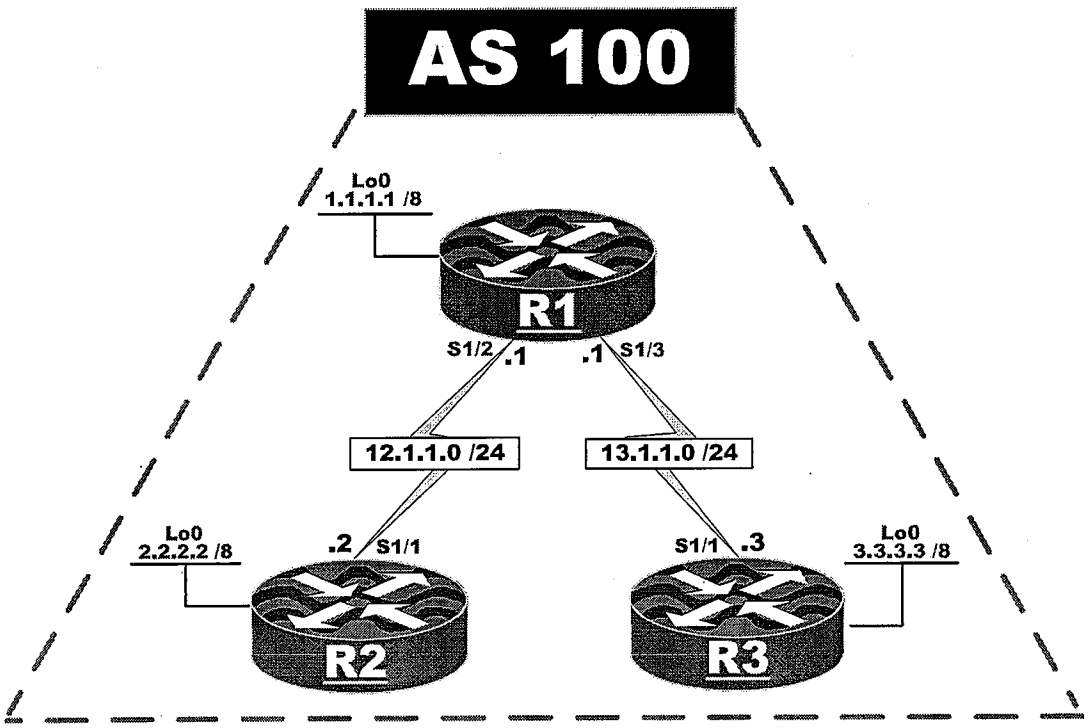
Peer-policy templates are used to build a template of policy information that can be inherited by a given neighbor/s. The peer-policy template cannot be inherited by a peer-session template or a peer-group.

Task 8

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

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 2
Route Reflectors



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-2"→"Task-1-3" →"Lab-2-T1-3.txt".

Task 1

Configure BGP AS 100 on all routers and ensure that the routers can successfully establish an IBGP peer session with each other. These routers should only advertise their Loopback0 interface in BGP.

On R1

```
R1(config)#Router bgp 100
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 100
R1(config-router)#Neighbor 13.1.1.3 remote-as 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-router)#No Auto
```

On R2

```
R2(config)#Router bgp 100
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 13.1.1.3 remote-as 100
R2(config-router)#No Auto
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 12.1.1.1 Up
```

On R3

```
R3(config)#Router bgp 100
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Neighbor 13.1.1.1 remote-as 100
R3(config-router)#Neighbor 12.1.1.2 remote-as 100
R3(config-router)#No Auto
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 13.1.1.1 Up
```

To verify the configuration:

On R1

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          12.1.1.2                 0    100      0 i
*>i3.0.0.0          13.1.1.3                 0    100      0 i

R1#Show ip route bgp | Inc B

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
B       2.0.0.0/8 [200/0] via 12.1.1.2, 00:07:34
B       3.0.0.0/8 [200/0] via 13.1.1.3, 00:05:27
```

On R2

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          12.1.1.1                 0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i

R2#Show ip route bgp | Inc B

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
B       1.0.0.0/8 [200/0] via 12.1.1.1, 00:08:56
```

Network 3.0.0.0 is NOT in the BGP table of R2. Let's check their neighbor adjacency:

```
R2#Show ip bgp summary | B Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.1.1.1        4   100      16      16        4    0    0 00:10:33        1
13.1.1.3        4   100       0       0        1    0    0 never    Idle
```

The neighbor is in idle state.Do we have reachability to this neighbor? Let's verify using ping:

```
R2#Ping 13.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 13.1.1.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Let's check the routing table:

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

B       1.0.0.0/8 [200/0] via 12.1.1.1, 00:17:47
        2.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C          2.0.0.0/8 is directly connected, Loopback0
L          2.2.2.2/32 is directly connected, Loopback0
        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
```

We can NOT see Network 13.1.1.0 /24 in the routing table, this is one reason that an IGP should be configured within an AS, this IGP should provide reachability to the links within the AS. Let's run OSPF on all the links within this topology:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Router ospf 1
R1(config-router)#Router-id 0.0.0.1
R1(config-router)#Network 12.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 13.1.1.1 0.0.0.0 area 0
```

On R2

```
R2(config)#Router ospf 1
R2(config-router)#Router-id 0.0.0.2
R2(config-router)#Network 12.1.1.2 0.0.0.0 area 0
```

On R3

```
R3(config)#Router ospf 1
R3(config-router)#Router-id 0.0.0.3
R3(config-router)#Network 13.1.1.3 0.0.0.0 area 0

R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          13.1.1.1                 0    100      0 i
*>i2.0.0.0          12.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i

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 = 56/57/60 ms
```

Task 3

You received an e-mail from the management stating that within the next 12 months 20 additional routers will be added to this AS. In order to minimize the number of peer sessions within this AS, you decided to implement route reflectors. Configure R1 as a route reflector for this AS.

On R1

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 12.1.1.2 route-reflector-client
R1(config-router)#Neighbor 13.1.1.3 route-reflector-client
```

On R2

```
R2(config)#Router bgp 100
R2(config-router)#No Neighbor 13.1.1.3 remote-as 100
```

On R3

```
R3(config)#Router bgp 100
R3(config-router)#No Neighbor 12.1.1.2 remote-as 100
```

To verify the configuration:

On R3

```
R3#Show ip bgp | b Net

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          13.1.1.1                 0    100      0 i
*>i2.0.0.0          12.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i

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 = 56/58/60 ms
```

In order for all IBGP speakers in an AS to exchange routes with one another, the IBGP speakers must be fully meshed (Every router must establish a peer session to every other router) this is because of a rule called IBGP split-horizon, which states that when a BGP speaker receives and update from an IBGP peer, it will NEVER advertise that update to another IBGP peer. Route-reflectors can be configured to reduce the number of peer sessions that must be established between the routers within a given AS. If a route-reflector is used, all IBGP speakers need not be fully meshed. In this model, the router that is configured to be the route-reflector must have a peer session established to every client, the clients must establish a peer session with the route reflector ONLY. The route reflector will reflect routes learned from one client to the other client/s.

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          12.1.1.2                 0    100      0 i
*>i3.0.0.0          13.1.1.3                 0    100      0 i
```

The output of the following show commands reveal the fact that prefixes 2.0.0.0 and 3.0.0.0/8 were received from a route-Reflector Client:

```
R1#Show ip bgp 2.0.0.0/8

BGP routing table entry for 2.0.0.0/8, version 6
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        2
  Local, (Received from a RR-client)
    10.1.12.2 from 10.1.12.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R1#Show ip bgp 3.0.0.0/8

BGP routing table entry for 3.0.0.0/8, version 9
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     2
  Local, (Received from a RR-client)
    13.1.1.3 from 13.1.1.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

On R2

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          12.1.1.1                 0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          13.1.1.3                 0    100      0 i

R2#Ping 3.3.3.3

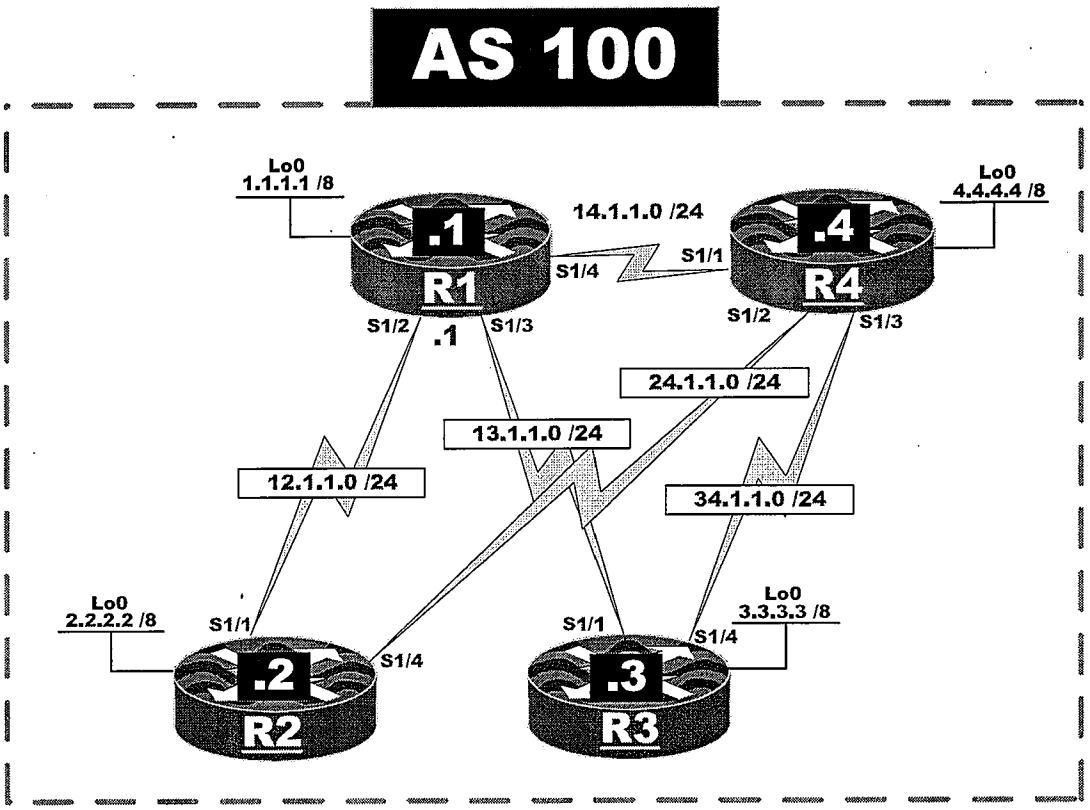
Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms
```

Task 4

After implementing the route reflector, you realized that if the route reflector is down the entire network is dysfunctional; therefore, you decided to add R4 as the second route reflector for redundancy. Ensure that these routers can reach the advertised networks and the redundancy is operational. R4 should advertise its Lo0 interface in BGP.



DO NOT erase the existing configuration, add the following configuration to the existing configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-2"→"Task-4".

R4 is the secondary route-reflector. R4 should be configured based on the following:

- R4 should have a peer session with R1 – the route-reflectors should have full mesh peer sessions between them.
- R4 must have a peer session with R2 and R3.
- R4 must configure R2 and R3 as route-reflector clients.

On R4:

```
R4(config)#Int S1/1
R4(config-if)#IP address 14.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int S1/2
R4(config-if)#IP address 24.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int S1/3
R4(config-if)#IP address 34.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int Lo0
R4(config-if)#IP address 4.4.4.4 255.0.0.0

R4(config)#Router bgp 100
R4(config)#Network 4.0.0.0
R4(config-router)#Neighbor 14.1.1.1 remote-as 100
R4(config-router)#Neighbor 24.1.1.2 remote-as 100
R4(config-router)#Neighbor 34.1.1.3 remote-as 100
R4(config-router)#Neighbor 24.1.1.2 route-reflector client
R4(config-router)#Neighbor 34.1.1.3 route-reflector client

R4(config)#Router ospf 1
R4(config-router)#Router-id 0.0.0.4
R4(config-router)#Network 14.1.1.4 0.0.0.0 area 0
R4(config-router)#Network 24.1.1.4 0.0.0.0 area 0
R4(config-router)#Network 34.1.1.4 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R2:

```
R2(config)#Int S1/4
R2(config-if)#IP addr 24.1.1.2 255.255.255.0
R2(config-if)#Clock rate 64000
R2(config-if)#No shut

R2(config)#Router bgp 100
R2(config-router)#Neighbor 24.1.1.4 remote-as 100

R2(config)#Router ospf 1
R2(config-router)#Network 24.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Int S1/4
R3(config-if)#IP addre 34.1.1.3 255.255.255.0
R3(config-if)#Clock rate 64000
R3(config-if)#No shut

R3(config)#Router bgp 100
R3(config-router)#Neighbor 34.1.1.4 remote-as 100

R3(config)#Router ospf 1
R3(config-router)#Network 34.1.1.3 0.0.0.0 area 0
```

On R1:

```
R1(config)#Int S1/4
R1(config-if)#IP address 14.1.1.1 255.255.255.0
R1(config-if)#Clock rate 64000
R1(config-if)#No shut

R1(config)#Router bgp 100
R1(config-router)#Neighbor 14.1.1.4 remote-as 100

R1(config)#Router ospf 1
R1(config-router)#Network 14.1.1.1 0.0.0.0 area 0
```

Having a single RR can introduce a single point of failure, it's best to have multiple RRs incase the main RR fails, this redundancy is critical when there are many RR clients.

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Show ip bgp | B Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	0.0.0.0	0		32768	i
* i2.0.0.0	24.1.1.2	0	100	0	i
*>i	12.1.1.2	0	100	0	i
* i3.0.0.0	34.1.1.3	0	100	0	i
*>i	13.1.1.3	0	100	0	i
*>i4.0.0.0	14.1.1.4	0	100	0	i

On R2:

```
R2#Show ip bgp | B Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* i1.0.0.0	14.1.1.1	0	100	0	i
*>i	12.1.1.1	0	100	0	i
*> 2.0.0.0	0.0.0.0	0		32768	i
* i3.0.0.0	34.1.1.3	0	100	0	i
*>i	13.1.1.3	0	100	0	i
* i4.0.0.0	14.1.1.4	0	100	0	i
*>i	24.1.1.4	0	100	0	i

On R3:

```
R3#Show ip bgp | B Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* i1.0.0.0	14.1.1.1	0	100	0	i
*>i	13.1.1.1	0	100	0	i
* i2.0.0.0	24.1.1.2	0	100	0	i
*>i	12.1.1.2	0	100	0	i
*> 3.0.0.0	0.0.0.0	0		32768	i
* i4.0.0.0	14.1.1.4	0	100	0	i
*>i	34.1.1.4	0	100	0	i

On R4:

```
R4#Show ip bgp | B Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i1.0.0.0	14.1.1.1	0	100	0	i

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
* i2.0.0.0          12.1.1.2                 0    100      0 i
*>i                 24.1.1.2                 0    100      0 i
* i3.0.0.0          13.1.1.3                 0    100      0 i
*>i                 34.1.1.3                 0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

To test the configuration:

On R1:

```
R1(config)#int s1/2
R1(config-if)#shut
```

On R2:

You should wait to see the following console message before entering the following show command:

```
%LINK-3-UPDOWN: Interface Serial1/1, changed state to down

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from FULL to DOWN,
Neighbor Down: Interface down or detached

%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial1/1, changed state
to down

%BGP-5-ADJCHANGE: neighbor 12.1.1.1 Down BGP Notification sent
%BGP-3-NOTIFICATION: sent to neighbor 12.1.1.1 4/0 (hold time expired) 0
bytes

%BGP_SESSION-5-ADJCHANGE: neighbor 12.1.1.1 IPv4 Unicast topology base
removed from session  BGP Notification sent *  28 msec

R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          14.1.1.1                 0    100      0 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          34.1.1.3                 0    100      0 i
*>i4.0.0.0          24.1.1.4                 0    100      0 i

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, 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 = 112/114/116 ms

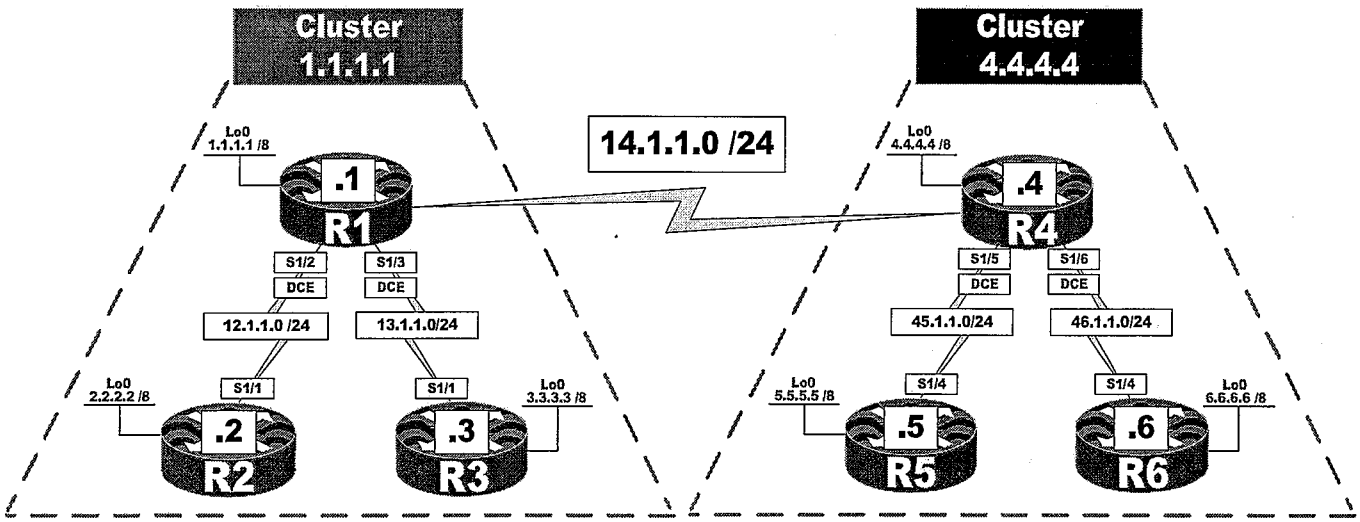
R2#Traceroute 3.3.3.3 numeric

Type escape sequence to abort.

Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 24.1.1.4 12 msec 16 msec 16 msec
  2 34.1.1.3 28 msec *  28 msec
```

Task 5

Erase the startup configuration and reload the routers. Reconfigure the routers based on the following topology. DO NOT configure BGP.



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-2"→"Task-5" →"Lab-2-T5".

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 6

Configure BGP on R1 – R6 based on the following policy:

- R1 should be configured as the route reflector for routers R2 and R3, whereas R4 should be configured to be the route-reflector for routers R5 and R6.
- R1 and R4 should be configured to have an IBGP peer session between them; these routers (R1-R6) should be configured in BGP AS 100.
- NLRI for the links should be provided through OSPF.

On R1

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Network 1.0.0.0
R1(config-router)#Neighbor 14.1.1.4 remote-as 100
R1(config-router)#Neighbor 12.1.1.2 remote-as 100
R1(config-router)#Neighbor 13.1.1.3 remote-as 100
R1(config-router)#Neighbor 12.1.1.2 route-reflector-client
R1(config-router)#Neighbor 13.1.1.3 route-reflector-client

R1(config)#Router ospf 1
R1(config-router)#Router-id 0.0.0.1
R1(config-router)#Network 12.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 14.1.1.1 0.0.0.0 area 0
R1(config-router)#Network 13.1.1.1 0.0.0.0 area 0
```

On R2

```
R2(config)#Router bgp 100
R2(config-router)#No au
R2(config-router)#Network 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100

R2(config)#Router ospf 1
R2(config-router)#Router-id 0.0.0.2
R2(config-router)#Network 12.1.1.2 0.0.0.0 area 0
```

On R3

```
R3(config)#Router bgp 100
R3(config-router)#No au
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-router)#Network 3.0.0.0
R3(config-router)#Neighbor 13.1.1.1 remote-as 100

R3(config)#Router ospf 1
R3(config-router)#Router-id 0.0.0.3
R3(config-router)#Network 13.1.1.3 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*>i2.0.0.0          12.1.1.2                 0    100      0 i
*>i3.0.0.0          13.1.1.3                 0    100      0 i
```

On R2:

```
R2#Show ip route ospf | Inc O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
O        13.1.1.0 [110/1562] via 12.1.1.1, 00:03:29, Serial1/1
O        14.1.1.0 [110/1562] via 12.1.1.1, 00:03:29, Serial1/1

R2#Ping 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#No au
R4(config-router)#Network 4.0.0.0
R4(config-router)#Neighbor 14.1.1.1 remote-as 100
R4(config-router)#Neighbor 45.1.1.5 remote-as 100
R4(config-router)#Neighbor 46.1.1.6 remote-as 100
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config-router)#Neighbor 45.1.1.5 route-reflector-client
R4(config-router)#Neighbor 46.1.1.6 route-reflector-client

R4(config)#Router ospf 1
R4(config-router)#Router-id 0.0.0.4
R4(config-router)#Network 14.1.1.4 0.0.0.0 area 0
R4(config-router)#Network 45.1.1.4 0.0.0.0 area 0
R4(config-router)#Network 46.1.1.4 0.0.0.0 area 0
```

On R5:

```
R5(config)#Router bgp 100
R5(config-router)#No au
R5(config-router)#Network 5.0.0.0
R5(config-router)#Neighbor 45.1.1.4 remote-as 100

R5(config)#Router ospf 1
R5(config-router)#Router-id 0.0.0.5
R5(config-router)#Network 45.1.1.5 0.0.0.0 area 0
```

On R6:

```
R6(config)#Router bgp 100
R6(config-router)#No au
R6(config-router)#Network 6.0.0.0
R6(config-router)#Neighbor 46.1.1.4 remote-as 100

R6(config)#Router ospf 1
R6(config-router)#Router-id 0.0.0.6
R6(config-router)#Network 46.1.1.6 0.0.0.0 area 0
```

To verify the configuration:

On R6:

```
R6#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          14.1.1.1                 0    100      0 i
*>i2.0.0.0          12.1.1.2                 0    100      0 i
*>i3.0.0.0          13.1.1.3                 0    100      0 i
*>i4.0.0.0          46.1.1.4                 0    100      0 i
*>i5.0.0.0          45.1.1.5                 0    100      0 i
*> 6.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

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

B      1.0.0.0/8 [200/0] via 13.1.1.1, 00:21:49
B      2.0.0.0/8 [200/0] via 12.1.1.2, 00:20:49
B      4.0.0.0/8 [200/0] via 14.1.1.4, 00:17:19
B      5.0.0.0/8 [200/0] via 45.1.1.5, 00:13:43
B      6.0.0.0/8 [200/0] via 46.1.1.6, 00:01:47

R3#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 = 84/86/88 ms

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 13.1.1.1 16 msec 12 msec 12 msec
  2 14.1.1.4 32 msec 28 msec 28 msec
  3 46.1.1.6 44 msec *  40 msec
```

Note: AS 100 has two route reflectors, each route-reflector has its own clients, when a given RR receives an update from one of its clients, it advertises that prefix to the other RR/s, the other RR/s in turn advertise that prefix to their clients.

There are some additional optional non-transitive attributes that can be used when RRs are configured and they are: "originator-id", "cluster-id" and "cluster-list".

- **Originator-id: This attribute is created by the RR; this is the router-id of the router that originated the prefix. The originator-id created to avoid routing loops, a RR will NOT advertise a route back to the originator of the prefix, and if the originator of a prefix receives an update with its own router-id, it will ignore that prefix.**

- **Cluster and Cluster-id: A RR/s and its clients are collectively known as a cluster, each cluster must be uniquely identified, and the "cluster-id" is the router-id of the RR unless specifically configured.**

- **Cluster-list: This attribute is analogous to AS-path attribute, and it keeps track of the cluster-ids in the same way that the AS-path attribute keeps track of the AS numbers. When the RR advertises a**

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

prefix to a non-client, it appends the cluster-id to that prefix's cluster-list, if a RR receives an update and sees its own cluster-id in the cluster-list, it will ignore that update.

To see the attributes:

On R1:

```
R1#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 7
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Advertised to update-groups:
        2
  Local
    46.1.1.6 (metric 1) from 14.1.1.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 6.6.6.6, Cluster list: 4.4.4.4
```

Note: Prefix 6.0.0.0 was originated by R6 with a router-id of 6.6.6.6 (The originator of the prefix), and it came from cluster-id of 4.4.4.4 (The cluster-list).

On R3:

```
R3#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 8
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Local
    46.1.1.6 (metric 2343) from 13.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 6.6.6.6, Cluster list: 1.1.1.1, 4.4.4.4
```

Note: This prefix has gone through cluster-ids of 4.4.4.4 first and then it traversed through cluster-id 1.1.1.1 before it was received by the local router.

The originator-id is the router-id of the router that originated that prefix, the output of the following "Show" command reveals the router-id of the router that originated the route.

```
R6#Sh ip bgp | Inc router ID

BGP table version is 8, local router ID is 6.6.6.6
```


© 2014 Narbik Kocharians. All rights reserved

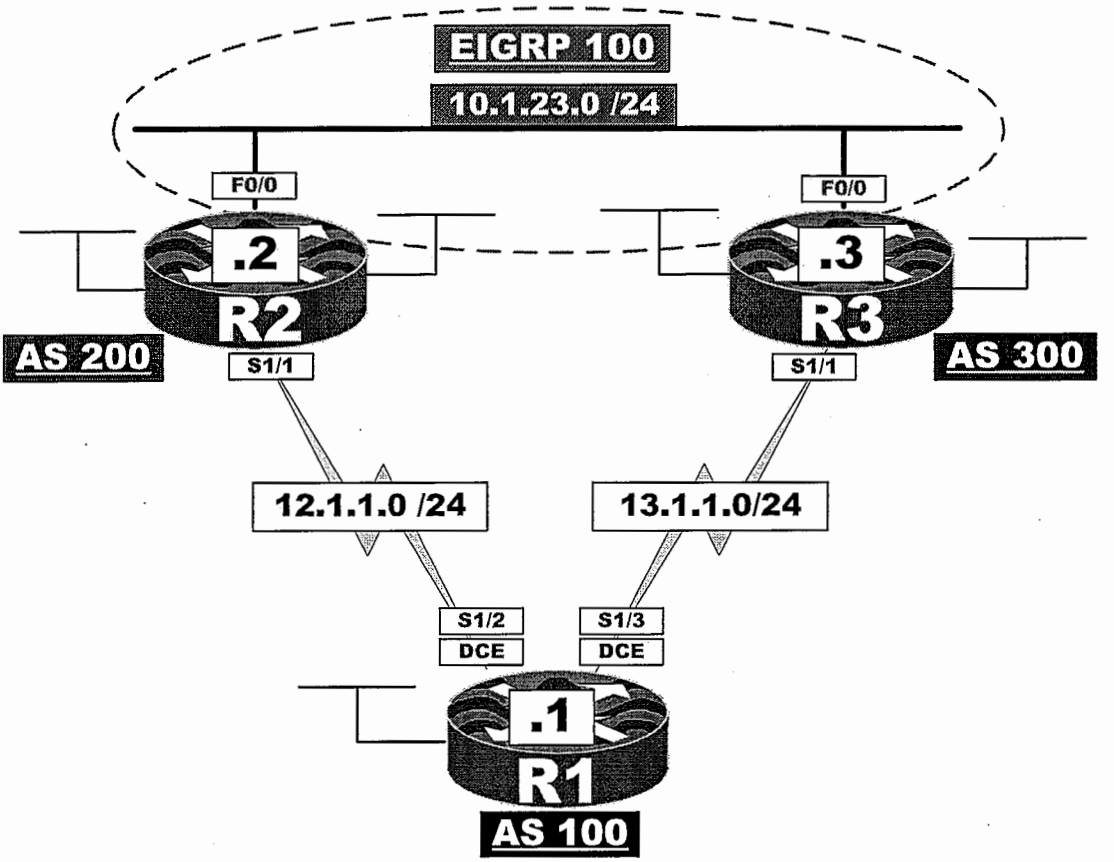
https://t.me/learningnets

Task 7

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 3
Conditional Advertisement & BGP Backdoor



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-3".

Task 1

Configure R1 in AS 100 to establish an EBGP peer session with R2 and R3 in AS 200 and 300 respectively.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No auto
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Neighbor 13.1.1.3 remote-as 300
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Neighbor 13.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.12.2       4   200       3       3        1    0    0 00:00:39        0
10.1.13.3       4   300       3       3        1    0    0 00:00:32        0
```

Task 2

Configure R1, R2 and R3 to advertise their loopback0 interface in BGP.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Network 1.0.0.0
```

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config)#Router bgp 200
R2(config-router)#Network 2.0.0.0
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 3.0.0.0
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 2.0.0.0          13.1.1.1                               0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 3

Configure RIPv2 and Eigrp 100 on the routers as follows:

- Configure RIPv2 on networks 12.1.1.0 /24 and 13.1.1.0 /24; disable auto summarization.
- R2 and R3 should advertise their F0/0, and Loopback 1 interfaces in Eigrp AS 100. Disable auto summarization.

On R1:

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 12.0.0.0
R1(config-router)#Network 13.0.0.0
```

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 12.0.0.0

R2(config)#Router eigrp 100
R2(config-router)#No au
R2(config-router)#Network 10.1.23.2 0.0.0.0
R2(config-router)#Network 10.1.2.2 0.0.0.0
```

On R3:

```
R3(config)#Router rip
R3(config-router)#No au
R3(config-router)#Ver 2
R3(config-router)#Network 13.0.0.0

R3(config)#Router eigrp 100
R3(config-router)#No au
R3(config-router)#Network 10.1.3.3 0.0.0.0
R3(config-router)#Network 10.1.23.3 0.0.0.0
```

To verify the configuration:

On R2:

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

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D        10.1.3.0/24 [90/156160] via 10.1.23.3, 00:00:52, FastEthernet0/0

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

      13.0.0.0/24 is subnetted, 1 subnets
R        13.1.1.0 [120/1] via 12.1.1.1, 00:00:15, Serial1/1
```

Task 4

Since network 10.1.23.0 is NOT advertised in BGP, if the link between R2 and R3 (The F0/0 interface) goes down, the Loopback1 network of these two routers won't have

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

reachability to each other, even though there is a redundant link between these two routers through BGP, therefore, the administrator of R2 and R3 decided that Loopback 1 interfaces of R2 and R3 should be advertised in BGP for redundancy, configure these routers to accommodate this decision.

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Network 10.1.2.0 mask 255.255.255.0
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 10.1.3.0 mask 255.255.255.0
```

To verify the configuration:

On R2:

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

B      1.0.0.0/8 [20/0] via 12.1.1.1, 00:16:27
B      3.0.0.0/8 [20/0] via 12.1.1.1, 00:15:57
       10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B         10.1.3.0/24 [20/0] via 12.1.1.1, 00:00:13
```

On R3:

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

B      1.0.0.0/8 [20/0] via 13.1.1.1, 00:17:06
B      2.0.0.0/8 [20/0] via 13.1.1.1, 00:16:05
       10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B         10.1.2.0/24 [20/0] via 13.1.1.1, 00:01:22
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 5

After implementing the previous task, the administrators realized that the traffic between networks 10.1.2.0 /24 and 10.1.3.0 /24 is taking a sub-optimal path and it is not using the direct path between routers R2 and R3.

Implement a BGP solution to fix this problem; you should **NOT** use the distance, PBR or any global config mode command to accomplish this task.

To see the suboptimal path:

On R3:

```
R3#Traceroute 10.1.2.2

Type escape sequence to abort.
Tracing the route to 10.1.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 13.1.1.1 16 msec 16 msec 12 msec
  2 12.1.1.2 32 msec *  28 msec

R3#Show ip route 10.1.2.2

Routing entry for 10.1.2.0/24
  Known via "bgp 300", distance 20, metric 0
  Tag 100, type external
  Last update from 13.1.1.1 00:07:02 ago
  Routing Descriptor Blocks:
  * 13.1.1.1, from 13.1.1.1, 00:07:02 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 100
      MPLS label: none
```

NOTE: The BGP "Backdoor" option can help us to accomplish this task. The "Backdoor" keyword is added to the network command that is advertised to you, therefore, you should reference the network that is advertised to you and NOT the network that your local router is advertising:

On R2:

```
R2(config)#Router bgp 200
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config-router)#Network 10.1.3.0 mask 255.255.255.0 backdoor
```

To verify the configuration:

On R2:

```
R2#Show ip route 10.1.3.3

Routing entry for 10.1.3.0/24
  Known via "eigrp 100", distance 90, metric 156160, type internal
  Redistributing via eigrp 100
  Last update from 10.1.23.3 on FastEthernet0/0, 00:00:56 ago
  Routing Descriptor Blocks:
  * 10.1.23.3, from 10.1.23.3, 00:00:56 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

R2#Traceroute 10.1.3.3

Type escape sequence to abort.
Tracing the route to 10.1.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.23.3 4 msec *  0 msec
```

We can see that R2 uses its direct connection (F0/0 interface) to reach the Loopback 1 interface of R3. The "Backdoor" keyword increases the administrative distance through EBGP for the advertised network to 200 so the local route can use the IGP and NOT EBGP's advertisement. Let's test the redundancy:

On R2, let's shutdown its F0/0 interface and verify reachability:

On R2:

```
R2(config)#Int F0/0
R2(config-if)#Shut

R2#Show ip route 10.1.3.3

Routing entry for 10.1.3.0/24
  Known via "bgp 200", distance 200, metric 0
  Tag 100, type locally generated
  Last update from 12.1.1.1 00:00:42 ago
  Routing Descriptor Blocks:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  * 12.1.1.1, from 12.1.1.1, 00:00:42 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 100
      MPLS label: none


R2#Traceroute 10.1.3.3

Type escape sequence to abort.
Tracing the route to 10.1.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.1 16 msec 16 msec 12 msec
  2 13.1.1.3 32 msec *  28 msec
```

Let's enable the F0/0 interface of R2 and configure the same on R3:

On R2:

```
R2(config)#Int F0/0
R2(config-if)#No shut

R2#Show ip route 10.1.3.3

Routing entry for 10.1.3.0/24
  Known via "eigrp 100", distance 90, metric 156160, type internal
  Redistributing via eigrp 100
  Last update from 10.1.23.3 on FastEthernet0/0, 00:00:33 ago
  Routing Descriptor Blocks:
  * 10.1.23.3, from 10.1.23.3, 00:00:33 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
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 10.1.2.0 mask 255.255.255.0 backdoor
```

To verify the configuration:

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3#Sh ip rou eigrp | B Gate

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D        10.1.2.0/24 [90/156160] via 10.1.23.2, 00:00:20, FastEthernet0/0
```

On R2:

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

      10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D        10.1.3.0/24 [90/156160] via 10.1.23.3, 00:07:07, FastEthernet0/0
```

NOTE: R2 and R3 were receiving routing information for networks 10.1.2.0 /24 and 10.1.3.0 /24 from two different sources, BGP and EIGRP.

R2 and R3 were using the routing information from BGP because it had a lower administrative distance (20 for EBGP versus 90 for Eigrp).

The Network command with the "backdoor" option is a BGP solution to this problem; the BGP "backdoor" option assigns an administrative distance of 200 to networks 10.1.2.0 /24 and 10.1.3.0 /24, therefore, making the Eigrp more attractive.

Task 6

Remove the IP address from the F0/0 interfaces of R2 and R3 and ensure that the F0/0 interfaces of both routers are in administratively down state. You should also remove the Loopback1 interface from these two routers.

On R2 and R3:

```
Rx(config)#Default interface F0/0

Rx(config)#Interface F0/0
Rx(config-if)#Shutdown

Rx(config)#No int lo1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 7

Configure R1 as follows:

- If network 2.0.0.0 /8 is up and it's advertised to R1, R1 should NOT advertise its network 1.0.0.0 /8 to R3.
- R1 should advertise network 1.0.0.0 /8 to R3 ONLY if network 2.0.0.0 /8 is down.

Before configuring this task you should verify the current BGP table of these routers:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
*> 3.0.0.0          13.1.1.3                 0             0 300 i

R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          12.1.1.1                               0 100 300 i

R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 2.0.0.0          13.1.1.1                               0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To implement conditional advertisement of selected prefixes, the following can be used:

- **Advertise-map**
- **Non-exist-map**
- **Exist-map**
- **Inject-map**

This situation calls for the use of the "advertise-map" and "non-exist-map", basically using these two commands we are saying advertise network 1.0.0.0 ONLY if network 2.0.0.0 is down, if network 2.0.0.0 is NOT down, then don't advertise network 1.0.0.0. To configure this task:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Step #1 – Identify the prefixes using two access-list/prefix-list:

```
R1(config)#Access-list 1 permit 1.0.0.0 0.255.255.255

R1(config)#Access-list 2 permit 2.0.0.0 0.255.255.255
```

Step #2 – Configure two route-maps, one to reference access-list 1 and the second one to reference access-list 2. To prevent confusion you should select meaningful names for the route-maps:

```
R1(config)#Route-map ADV permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#exit

R1(config)#Route-map NotThere permit 10
R1(config-route-map)#match ip addr 2
R1(config-route-map)#exit
```

Final Step – the route-maps are referenced by the "Advertise-map" and "non-exist-map" options:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 13.1.1.3 advertise-map ADV non-exist-map NotThere
```

The neighbor command has the following route-maps:

- The advertise-map – Specifies the name of the route-map that will be advertised if the condition of the non-exist-map is met.
- Non-exist-map – specifies the name of the route-map that will be compared to the advertise-map. If the condition is met and no match occurs, the route will be advertised. If a match occurs, then the condition is NOT met, and the route is withdrawn.

If network 2.0.0.0 is up, then network 1.0.0.0 should NOT be advertised to R3, since all the networks are up and advertised, R1 should withdraw its network (1.0.0.0 /8):

On R1:

NOTE: Network 2.0.0.0 is up so network 1.0.0.0 /8 should NOT be advertised to R3.

```
R1#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*> 3.0.0.0          13.1.1.3                 0             0 300 i
```

The following show command reveals that R1 does NOT advertise its network (1.0.0.0 /8) to R3:

```
R1#Show ip bgp neighbors 13.1.1.3 advertised-routes | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          12.1.1.2                 0             0 200 i
```

To verify this configuration

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          13.1.1.1                               0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the condition:

On R2:

```
R2(config)#int lo0
R2(config-if)#Shut
```

The output of the following show command reveals that network 2.0.0.0 is DOWN; and R1 is advertising its network (1.0.0.0 /8) to R3. It may take few seconds for this policy to be implemented:

On R1:

```
R1#Show ip bgp neighbors 13.1.1.3 advertised-routes | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```

To see the test on R3:

On R3:

```
R3#Show ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 8

Remove the configuration commands entered in the previous task before you proceed to the next task. Ensure that the routers have the advertised networks in their BGP table.

On R1:

```
R1(config)#No access-list 1
R1(config)#No access-list 2

R1(config)#No route-map ADV

R1(config)#No route-map NotThere

R1(config)#Router bgp 100
R1(config-router)#No Neighbor 13.1.1.3 advertise-map ADV non-exist-map NotThere

R1#Clear ip bgp *
```

On R2:

```
R2(config)#int lo0
R2(config-if)#No shut
```

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
*> 3.0.0.0          13.1.1.3                 0             0 300 i
```

On R2:

```
R2#Show ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          12.1.1.1                               0 100 300 i
```

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 2.0.0.0          13.1.1.1                               0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

Task 9

R1 should be configured based on the following policy:

1. If both networks (1.0.0.0 /8 and 2.0.0.0 /8) are up, then both networks should be advertised to R3.
2. If network 1.0.0.0 /8 is down, R1 should NOT advertise network 2.0.0.0 /8 to R3.
3. If network 2.0.0.0 /8 is down, then R1 should only advertise network 1.0.0.0 /8 to R3.

The logic in the following configuration says "ONLY advertise network 2.0.0.0/8 if network 1.0.0.0/8 is up, so if network 1.0.0.0/8 is NOT up, then DON'T advertise network 2.0.0.0/8.

On R1:

Step #1 - The following two access-lists identify the two networks (1.0.0.0 /8 and 2.0.0.0 /8):

```
R1(config)#Access-list 1 permit 1.0.0.0 0.255.255.255
R1(config)#Access-list 2 permit 2.0.0.0 0.255.255.255
```

Step #2 - The following route-maps are configured to reference the two access-lists from the previous step:

```
R1(config)#Route-map ADV permit 10
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-route-map)#match ip addr 2

R1(config)#Route-map EXIST permit 10
R1(config-route-map)#match ip addr 1
```

Final Step – With the following configuration, we are instructing BGP for the conditions of the task's requirements:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 13.1.1.3 advertise-map ADV exist-map EXIST
```

To test the first condition

If both networks (1.0.0.0 /8 and 2.0.0.0 /8) are up, then both networks should be advertised to R3.

NOTE: Both prefixes are up:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
*> 3.0.0.0          13.1.1.3                 0             0 300 i
```

Let's see the prefixes that R1 is advertising to R3:

On R1:

```
R1#Show ip bgp neighbor 13.1.1.3 advertised-routes | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
```

As you can see both prefixes are being advertised to R3, let's check R3's BGP table:

On R3:

```
R3#Show ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 2.0.0.0          13.1.1.1                 0             0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the second condition

If network 1.0.0.0 /8 is down, R1 should NOT advertise network 2.0.0.0 /8 to R3.

Let's shutdown R1's Lo0 interface:

On R1:

```
R1(config)#Int lo0
R1(config-if)#Shut
```

To force the change much faster:

On R1:

```
R1#Clear ip bgp *
```

Let's see the prefixes that R1 is advertising to R3:

```
R1#Show ip bgp neighbor 13.1.1.3 advertised-routes | B Network
R1#

R1#Sh ip bgp 2.0.0.0

BGP routing table entry for 2.0.0.0/8, version 4
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  200
    12.1.1.2 from 12.1.1.2 (10.1.2.2)
      Origin IGP, metric 0, localpref 100, valid, external, best
```

NOTE: The output of the above show command reveals that R1 is NOT advertising any prefixes to R3. Let's check R3's BGP table to verify:

On R3:

```
R3#Show ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 3.0.0.0          0.0.0.0                  0         32768 i
```

To test the third condition

If network 2.0.0.0 /8 is down, then R1 should only advertise network 1.0.0.0 /8 to R3.

Let's "NO shut" R1's Lo0 and shutdown R2's Lo0:

On R1:

```
R1(config)#Int Lo0
R1(config-if)#NO Shut
```

On R2:

```
R2(config)#Int Lo0
R2(config-if)#Shut
```

To force the change much faster:

On R1:

```
R1#Clear ip bgp *
```

Let's see which prefixes are advertised to R3 by R1:

```
R1#Show ip bgp neighbor 13.1.1.3 advertised-routes | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```

To verify the configuration

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 10

Configure R2 and R1 such that R1 advertises network 200.1.1.0/24. DO NOT configure another interface or logical interface with this IP address to accomplish this task. You may NOT have reachability to this network.

In order to accomplish this task, the "Exist-map" and the "Inject-map" is used. Basically the "exist-map" identifies the network/s that exist in the local router's (R1's) BGP table and the route-source (The router that has advertised it to R1); the "inject-map" identifies the network that should be injected into R1's BGP table if the condition of the "exist-map" is true.

Let's configure R2 to inject a default route, so R1 can match that default route in its "exist-map".

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 default-originate

R2#Clear ip bgp *
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Net
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	0.0.0.0	12.1.1.2			0	200 i
*>	1.0.0.0	0.0.0.0	0		32768	i
*>	3.0.0.0	13.1.1.3	0		0	300 i

Let's configure the task:

Step 1:

Configure three prefix-lists, one should reference the default route, the second one should reference R2 and the third one should refererence network 200.1.1.0/24:

```
R1(config)#Ip prefix-list R2 seq 5 permit 12.1.1.2/32
R1(config)#Ip prefix-list tst seq 5 permit 0.0.0.0/0
R1(config)#Ip prefix-list NET seq 5 permit 200.1.1.0/24
```

Step 2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Configure two route-maps, the first one will match on the route-source; this is the router that has advertised the default route to R1, in this case R2.
The second route-map will set the reuested network (200.1.1.0/24) into R1's BGP table:

```
R1(config)#Route-map EXIST permit 10
R1(config-route-map)#match ip addr prefix tst
R1(config-route-map)#match ip route-source prefix R2

R1(config)#Route-map INJECT permit 10
R1(config-route-map)#set ip addr prefix NET
```

The final step is to instruct BGP to inject network 200.1.1.0/24 into R1's BGP table ONLY if the route-source (R2) and the default route exist:

```
R1(config)#Router bgp 100
R1(config-router)#BGP inject-map INJECT exist-map EXIST

R1#Clear ip bgp *
```

To test and verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          12.1.1.2                             0 200 i
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
*> 3.0.0.0          13.1.1.3                 0             0 300 i
*> 200.1.1.0        12.1.1.2                               0 ?
```

On R3:

```
R3#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          13.1.1.1                             0 100 200 i
*> 1.0.0.0          13.1.1.1                 0           0 100 i
*> 2.0.0.0          13.1.1.1                             0 100 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 200.1.1.0        13.1.1.1                               0 100 ?
```

On R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Show ip bgp | B Net
   Network          Next Hop            Metric LocPrf Weight Path
   0.0.0.0          0.0.0.0                              0 i
*> 1.0.0.0          12.1.1.1                 0           0 100 i
*> 2.0.0.0          0.0.0.0                  0       32768 i
*> 3.0.0.0          12.1.1.1                             0 100 300 i
```

Since R2 is the router that is generating the default route, it will NOT get the route, meaning that R1 will NOT advertise the route to the route-source.

On R1:

```
R1#Show ip bgp neighbors 12.1.1.2 advertised-routes | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0       32768 i
*> 3.0.0.0          13.1.1.3                 0           0 300 i

Total number of prefixes 2
```

On R3:

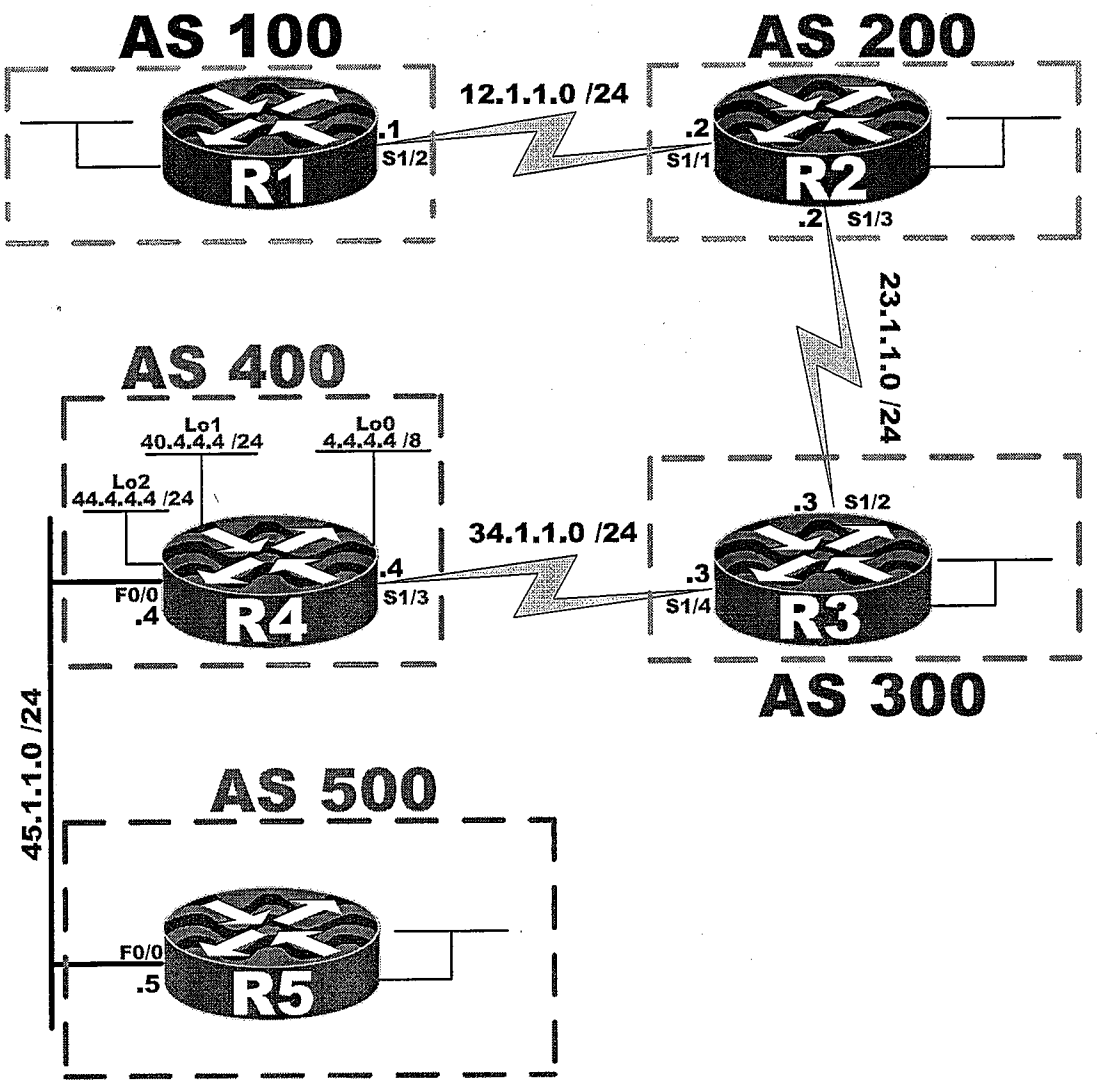
```
R3#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          13.1.1.1                             0 100 200 i
*> 1.0.0.0          13.1.1.1                 0           0 100 i
*> 2.0.0.0          13.1.1.1                             0 100 200 i
*> 3.0.0.0          0.0.0.0                  0       32768 i
*> 200.1.1.0        13.1.1.1                             0 100 ?
```

Task 11

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 4
Route Dampening



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-4".

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Configure an EBGP peer session between R1 and R2 and only advertise their loopback0 interfaces in BGP. Ensure that these routers have NLRI to each other's advertised prefix. R1 should be in AS 100 and R2 should be configured in AS 200.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Network 1.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Network 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          10.1.12.2                0             0 200 i
```

Task 2

Configure route dampening on R1 using the default parameters.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#bgp dampening
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The parameters of BGP dampening are as follows:

- Half time – Once a route has been assigned a penalty (1000 per flap), the accumulated penalty is decreased every 5 seconds such that when the half-time period expires. The accumulated penalties are reduced by half. The default value of half time is 15 minutes and the range is 1 to 45 minutes.
- Reuse – If the penalties for a flapping route is decreased enough to fall below this value, the route is reusable. The default is 750 and the range is 1 to 20000.
- Suppress – Once the accumulated penalties reaches this value, the route is suppressed. The default value is 2000 and the range is 1 – 20000.
- Max-Suppress-Time – The maximum time in minutes that a route can be suppressed. The default value is 4 times the half-time value (60 minutes) and the range is 1 to 255.

Therefore this configuration performs the following:

Half time = 15 minutes, reuse = 750, Suppress = 2000 and Max-Suppress-Time = 60.

To see the parameters for dampening:

On R1:

```
R1#Sh ip bgp dampening parameters

dampening 15 750 2000 60 (DEFAULT)
  Half-life time         : 15 mins       Decay Time            : 2320 secs
  Max suppress penalty: 12000            Max suppress time: 60 mins
  Suppress penalty     :  2000           Reuse penalty      : 750
```

If network 2.0.0.0 is shut down and then brought backup few times, the flap-statistics can be viewed in the output of the "Show ip bgp dampening flap-statistics" command.

Task 3

Configure an EBGP peer session between R2 and R3, and advertise their loopback0 interface in BGP. Ensure that these routers have NLRI to each other's loopback interface/s. R3 should be configured in AS 300.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 23.1.1.3 remote-as 300
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Network 3.0.0.0
R3(config-router)#Neighbor 23.1.1.2 remote-as 200
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          23.1.1.3                 0             0 300 i
```

Task 4

Configure route dampening on R3 such that the half-life parameter is set to 30 minutes. This router should use the default parameters for suppress-limit, reuse, and maximum suppress time.

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Bgp dampening 30 750 2000 60
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp dampening parameters
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
dampening 30 750 2000 60
  Half-life time         : 30 mins       Decay Time          : 1045 secs
  Max suppress penalty:  3000            Max suppress time: 60 mins
  Suppress penalty     :  2000           Reuse penalty     : 750
```

Note: You may receive a "% dampening reconfiguration in progress for IPv4 unicast" message; you should wait few seconds and try again.

Task 5

Configure an EBGP session between R3 and R4. Advertise loopback0, loopback1 and loopback2 interfaces of R4 in BGP. R4 should be configured in AS 400.

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Netw 40.4.4.0 mask 255.255.255.0
R4(config-router)#Netw 44.4.4.0 mask 255.255.255.0
R4(config-router)#Neighbor 34.1.1.3 remote-as 300
R4(config-router)#No au
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Neighbor 34.1.1.4 remote-as 400
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	23.1.1.2			0	200 100 i
*> 2.0.0.0	23.1.1.2	0		0	200 i
*> 3.0.0.0	0.0.0.0	0		32768	i
*> 4.0.0.0	34.1.1.4	0		0	400 i
*> 40.4.4.0/24	34.1.1.4	0		0	400 i
*> 44.4.4.0/24	34.1.1.4	0		0	400 i

© 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
*> 1.0.0.0          34.1.1.3                               0 300 200 100 i
*> 2.0.0.0          34.1.1.3                               0 300 200 i
*> 3.0.0.0          34.1.1.3                 0             0 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 40.4.4.0/24      0.0.0.0                  0         32768 i
*> 44.4.4.0/24      0.0.0.0                  0         32768 i
```

Task 6

Configure an EBGP peer session between R4 and R5. R5 should advertise its loopback 0 interface in BGP. Ensure that these routers have NLRI to each other's loopback interface. R5 should be configured in AS 500.

On R5:

```
R5(config)#Router bgp 500
R5(config-router)#No au
R5(config-router)#Neighbor 45.1.1.4 remote-as 400
R5(config-router)#Netw 5.0.0.0
```

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#Neighbor 45.1.1.5 remote-as 500
```

To verify the configuration:

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          45.1.1.4                               0 400 300 200 100 i
*> 2.0.0.0          45.1.1.4                               0 400 300 200 i
*> 3.0.0.0          45.1.1.4                               0 400 300 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 4.0.0.0          45.1.1.4                 0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 40.4.4.0/24      45.1.1.4                 0             0 400 i
*> 44.4.4.0/24      45.1.1.4                 0             0 400 i
```

Task 7

Configure route dampening On R4: as follows:

- Network 40.4.4.0 /24 should have the following dampening parameters applied: **max-suppress-time** of 90, **reuse** 800, **suppress** 2400 and a **half-time** of 20
- Network 44.4.4.0 /24 should have the following dampening parameters applied: **max-suppress-time** of 60, **reuse** 700, **suppress** 2000 and a **half-time** of 15

When route dampening is enabled, it affects all prefixes; but it can be configured to affect a specific prefix; this can be done through the configuration of route-map. Route-map gives us this flexibility.

In this case, we have applied two different route-dampening parameters to different routes. The two networks are identified with access-lists. The route-map references the access-lists and sets the dampening parameters based on the networks.

On R4:

```
R4(config)#Access-list 40 permit 40.4.4.0 0.0.0.255

R4(config)#Access-list 44 permit 44.4.4.0 0.0.0.255

R4(config)#Route-map TST permit 10
R4(config-route-map)#match ip addr 40
R4(config-route-map)#set dampening 20 800 2400 90

R4(config)#Route-map TST permit 20
R4(config-route-map)#match ip addr 44
R4(config-route-map)#set dampening 15 700 2000 60

R4(config)#Route-map TST permit 90

R4(config)#Router bgp 400
R4(config-router)#bgp dampening route-map TST
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R4:

```
R4#Sh ip bgp damp parameters

dampening 20 800 2400 90 (route-map TST 10)
  Half-life time      : 20 mins       Decay Time        : 3490 secs
  Max suppress penalty: 18075         Max suppress time: 90 mins
  Suppress penalty    :  2400         Reuse penalty     : 800

dampening 15 700 2000 60 (route-map TST 20)
  Half-life time      : 15 mins       Decay Time        : 2235 secs
  Max suppress penalty: 11200         Max suppress time: 60 mins
  Suppress penalty    :  2000         Reuse penalty     : 700
```

Task 8

Configure route dampening on R2 using the following policy:

- All the existing and future prefixes that originate in AS 300 should have the following parameters applied:

  Max-Suppress-Time of 80
  Reuse 750
  Suppress 2200
  Half Time 30

On R2:

```
R2(config)#Ip as-path access-list 1 permit ^300$

R2(config)#Route-map TST permit 10
R2(config-route-map)#match as-path 1
R2(config-route-map)#set dampening 30 750 2200 80
R2(config)#Route-map TST permit 90

R2(config)#Router bgp 200
R2(config-router)#bgp dampening route-map TST
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The combination of "route-map" and the "as-path access-list" command can apply bgp dampening to an AS based on the AS number, this will affect the existing and future prefixes that are advertised by this AS.

To verify the configuration:

On R2:

```
R2#Show ip bgp damp parameters

dampening 30 750 2200 80 (route-map TST 10)
  Half-life time       : 30 mins       Decay Time        : 1995 secs
  Max suppress penalty:  4755          Max suppress time: 80 mins
  Suppress penalty     :  2200         Reuse penalty     : 750
```

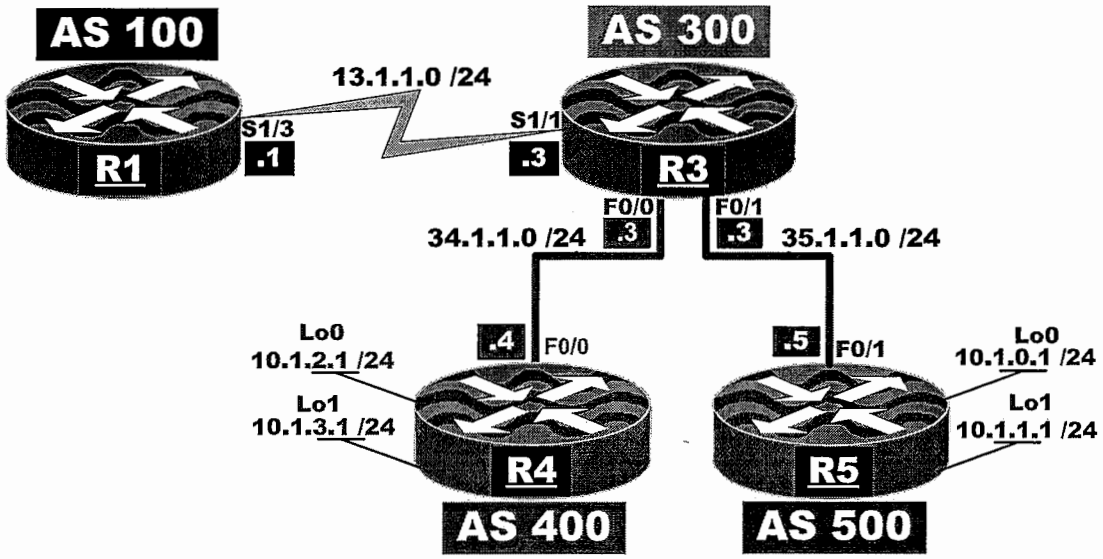
Task 9

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 5 – BGP Aggregation



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-5".

Task 1

R4 is configured in AS 400, and R5 is configured in AS 500. Configure R4 and R5 to advertise their Loopback interfaces in BGP, if this configuration is done successfully, R3 in AS 300 should see 4 prefixes in its BGP table. You should use "Network" command to accomplish this task.

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#Network 10.1.2.0 mask 255.255.255.0
R4(config-router)#Network 10.1.3.0 mask 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R5:

```
R5(config)#Router bgp 500
R5(config-router)#Network 10.1.0.0 mask 255.255.255.0
R5(config-router)#Network 10.1.1.0 mask 255.255.255.0
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	10.1.0.0/24	35.1.1.5	0		0	500 i
*>	10.1.1.0/24	35.1.1.5	0		0	500 i
*>	10.1.2.0/24	34.1.1.4	0		0	400 i
*>	10.1.3.0/24	34.1.1.4	0		0	400 i

Task 2

Configure R3 in AS 300 to aggregate all four prefixes received from AS 400 and AS 500. R3 should ONLY advertise the aggregate route to existing and future neighbor(s). R3 and R1 should establish a peer session after R3 aggregates the prefixes. If this configuration is performed successfully, R1 in AS 100 should ONLY receive a single prefix representing the four prefixes.

Note: In BGP, an aggregate is only created if at least one of the specific routes of the aggregate exists in the BGP table. It is recommended to configure most if not all the specific routes with a network statement, because if only a single network is configured to satisfy this requirement, and that particular network goes down, then, the aggregate will be removed.

There are many ways to advertise an aggregate, one way to advertise an aggregate is by creating a static route to "NULL0" that matches the aggregate route and then advertising the aggregate in BGP as follows:

```
(config)#Ip route 1.1.0.0 255.255.252.0 null 0

(config)#Router bgp 200
(config-router)#Network 1.1.0.0 mask 255.255.252.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

However, since that is not an option here, every specific prefix of the aggregate must be advertised individually.

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Aggregate-address 10.1.0.0 255.255.252.0
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/24      35.1.1.5                 0             0 500 i
*> 10.1.0.0/22      0.0.0.0                                32768 i
*> 10.1.1.0/24      35.1.1.5                 0             0 500 i
*> 10.1.2.0/24      34.1.1.4                 0             0 400 i
*> 10.1.3.0/24      34.1.1.4                 0             0 400 i
```

NOTE: R3 is advertising the aggregate plus all the specific prefixes. To suppress all the specific prefixes and ONLY advertise the aggregate, the "summary-only", or the "Suppress-map", or the "Unsuppress-map" can be used. In this example, the "Summary-only" keyword is used.

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Aggregate-address 10.1.0.0 255.255.252.0 Summary-only
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
s> 10.1.0.0/24      35.1.1.5                 0             0 500 i
*> 10.1.0.0/22      0.0.0.0                                32768 i
s> 10.1.1.0/24      35.1.1.5                 0             0 500 i
s> 10.1.2.0/24      34.1.1.4                 0             0 400 i
s> 10.1.3.0/24      34.1.1.4                 0             0 400 i
```

The letter "S" in the output of the above command reveals that the specific prefixes are suppressed and

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

the aggregate route is the ONLY prefix advertised to R1.

The second section of this task states to establish an EBGP peer session between R1 in AS 100 and R3 in AS 300. Let's configure this task:

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 13.1.1.3 remote-as 300
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Neighbor 13.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/22      13.1.1.3                 0             0 300 i
```

Task 3

Configure R3 such that R1 in AS 100 can see the AS number(s) where some or all the specific prefixes originated. DO NOT create a route-map to accomplish this task.

Before performing the task, let's see the BGP table for the aggregate route on R1:

On R1:

```
R1#Show ip bgp 10.1.0.0/22

BGP routing table entry for 10.1.0.0/22, version 2
Paths: (1 available, best #1, table default)
  Not advertised to any peer
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  300, (aggregated by 300 35.1.1.3)
    13.1.1.3 from 13.1.1.3 (35.1.1.3)
      Origin IGP, metric 0, localpref 100, valid, external, atomic-
aggregate, best
```

NOTE: When aggregation is configured, BGP will automatically attach two attributes to the aggregate route, the "Aggregator" and the "Atomic-aggregate". The aggregator attribute identifies the AS number and the RID of the router that performed the aggregation, and the Atomic-aggregate attribute tells the local router that some thing maybe hidden. In this case, we know that the specific prefixes are suppressed/hidden, but what else is hidden? The AS number(s) that some of the prefixes have originated in, let's check the BGP table of R1:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/22      13.1.1.3                 0             0 300 i
```

In the above output we can see that the specific prefixes are suppressed, we can also see that the AS 400 and 500 are also hidden. Let's reveal these AS numbers by configuring R3:

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Aggregate-address 10.1.0.0 255.255.252.0 Summary-only AS-Set
```

The "AS-Set" keyword reveals the AS number where some or all the specific routes have originated. To see and verify this point:

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
s> 10.1.0.0/24      35.1.1.5                 0             0 500 i
*> 10.1.0.0/22      0.0.0.0                                100  32768 {500,400} i
s> 10.1.1.0/24      35.1.1.5                 0             0 500 i
s> 10.1.2.0/24      34.1.1.4                 0             0 400 i
s> 10.1.3.0/24      34.1.1.4                 0             0 400 i
```

You can see that the AS numbers of the specific routes are now revealed. Let's check R1:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Show ip bgp | B Network
R1#
```

We don't see any prefixes. Maybe a "Clear ip bgp *" will do the trick.

On R1:

```
R1#Clear ip bgp *
```

Let's try the "Show ip bgp" command after 30 seconds:

```
R1#Show ip bgp | B Network
R1#
```

Is R3 advertising any prefixes to the local router?

```
R1#Show ip bgp neighbor 13.1.1.3 routes
```



```
Total number of prefixes 0
```

Let's check and see if R3 is advertising it:

On R3:

```
R3#Show ip bgp neighbor 13.1.1.1 advertised-routes
```

```
Total number of prefixes 0
```

Let's check the BGP table for the aggregate route:

On R3:

```
R3#Show ip bgp 10.1.0.0/22

BGP routing table entry for 10.1.0.0/22, version 11
Paths: (1 available, best #1, table default, not advertised to any peer)
  Not advertised to any peer
  {500,400}, (aggregated by 300 35.1.1.3)
    0.0.0.0 from 0.0.0.0 (35.1.1.3)
      Origin IGP, localpref 100, weight 32768, valid, aggregated, local, best
      Community: no-advertise
```

We can see the problem, a community of "No-Advertise" is attached to this aggregate route, how did this happen?

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

When the "AS-Set" keyword is used in aggregation, it cumulatively attaches all the attributes that were assigned to the specific prefixes to the aggregate route. To see the router that assigned a "No-advertise" community attribute to one of the specific prefixes:

On R3:

```
R3#Show ip bgp community no-advertise | B Network

   Network          Next Hop            Metric LocPrf Weight Path
s> 10.1.0.0/24      35.1.1.5                 0             0 500 i
*> 10.1.0.0/22      0.0.0.0                            100  32768 {500,400} i
```

We can see that it's 35.1.1.5 (R5), let's check R5's configuration:

On R5:

```
R5#Show run | s route-map

neighbor 35.1.1.3 route-map TST out
route-map TST permit 10
 match ip address 1
 set community no-advertise
route-map TST permit 90
```

The route-map called "TST" is assigning a community attribute of "No-Advertise" to any network(s) that is permitted in access-list 1, and the route-map is given to R3. Let's check access-list 1:

On R5:

```
R5#Show access-list 1

Standard IP access list 1
    10 permit 10.1.0.0, wildcard bits 0.0.0.255 (2 matches)
```

Since we can NOT configure any other router to accomplish this task. Let's configure R3.

On R3:

```
R3(config)#Route-map TST permit 10
R3(config-route-map)#Set community Internet

R3(config-route-map)#Router bgp 300
R3(config-router)#Aggregate-addr 10.1.0.0 255.255.252.0 summ as-set
attribute-map TST
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/22      13.1.1.3                 0             0 300 {500,400} i
```

The "Attribute-map" can be configured to override all the attributes attached to the aggregate route. The "attribute-map" references a route-map, in this case called "TST", and the route-map TST sets a community of "Internet" to the aggregate route, telling it to advertise it to all neighbors.

Task 4

Configure R3 such that ONLY R1 in AS 100 and R5 in AS 500 accept the aggregate, the other AS(es) in this topology (AS 400) should receive the aggregate route, but they should discard it.

On R3:

```
R3(config)#IP As-path access-list 1 permit ^400$

R3(config)#Route-map Task4 permit 10
R3(config-route-map)#Match as-path 1

R3(config)#Router bgp 300
R3(config-router)#Aggregate-addr 10.1.0.0 255.255.252.0 summ as-set
attribute-map TST advertise-map Task4
```

NOTE: The above configuration ONLY allows AS 400 to be advertised and it denies the other ASes to be advertised.

To verify the configuration:

On R1:

NOTE: R1 in AS 100 does NOT see its own AS number in the aggregate route advertised by R3, therefore, it accepts the update:

© 2014 Narbik Kocharians, All rights reserved

https://t.me/learningnets

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/22      13.1.1.3                 0             0 300 400 i
```

On R5:

NOTE: R5 in AS 500 does NOT see its own AS number in the aggregate route advertised by R3, therefore, it accepts the update:

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/24      0.0.0.0                  0         32768 i
*> 10.1.0.0/22      35.1.1.3                 0             0 300 400 i
*> 10.1.1.0/24      0.0.0.0                  0         32768 i
```

On R4:

NOTE: Since R4 sees its own AS number in the advertised aggregate route, it discards the route because of BGP's loop avoidance mechanism.

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.2.0/24      0.0.0.0                  0         32768 i
*> 10.1.3.0/24      0.0.0.0                  0         32768 i
```

Task 5

Configure R3 to advertise the aggregate plus 10.1.2.0 /24 prefix to R1 and R5.

On R3:

In the following configuration, we are denying Network 10.1.2.0 /24 from being suppressed, and allowing the rest of the prefixes that are part of the aggregate to be suppressed.

```
R3(config)#Access-list 1 deny 10.1.2.0 0.0.0.255
R3(config)#Access-list 1 permit any
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config)#Route-map Unsupp permit 10
R3(config-route-map)#match ip addr 1

R3(config)#Router bgp 300
R3(config-router)#Aggregate-addr 10.1.0.0 255.255.252.0 summ as-set
attribute-map TST advertise-map Task4 suppress-map Unsupp
```

To verify the configuration:

On R3:

NOTE: R3 is NO LONGER suppressing the 10.1.2.0 /24 prefix:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
s> 10.1.0.0/24      35.1.1.5                 0             0 500 i
*> 10.1.0.0/22      0.0.0.0                            100  32768 400 i
s> 10.1.1.0/24      35.1.1.5                 0             0 500 i
*> 10.1.2.0/24      34.1.1.4                 0             0 400 i
s> 10.1.3.0/24      34.1.1.4                 0             0 400 i
```

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/22      13.1.1.3                 0             0 300 400 i
*> 10.1.2.0/24      13.1.1.3                                 0 300 400 i
```

On R5:

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/24      0.0.0.0                  0         32768 i
*> 10.1.0.0/22      35.1.1.3                 0             0 300 400 i
*> 10.1.1.0/24      0.0.0.0                  0         32768 i
*> 10.1.2.0/24      35.1.1.3                                 0 300 400 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 7

Configure R3 to advertise 10.1.3.0 /24 to R1 ONLY. R5 should NOT receive this prefix.

The "Unsuppress-map" unsuppresses one or more of the specific prefixe(s) of an aggregate, the affect of this keyword is on a per neighbor basis, whereas, the "Suppress-map" affects all neighbors receiving the aggregate:

On R3:

Before configuring any access-list, you should check and see if one is configured already that is using the same number. Let's verify:

```
R3#Show access-list

Standard IP access list 1
    10 deny   10.1.2.0, wildcard bits 0.0.0.255 (3 matches)
    20 permit any (6 matches)

R3(config)#Access-list 3 permit 10.1.3.0 0.0.0.255

R3(config)#Route-map Task7 permit 10
R3(config-route-map)#match ip addr 3

R3(config)#Router bgp 300
R3(config-router)#Neighbor 13.1.1.1 unsuppress-map Task7
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/22      13.1.1.3                 0             0 300 400 i
*> 10.1.2.0/24      13.1.1.3                               0 300 400 i
*> 10.1.3.0/24      13.1.1.3                               0 300 400 i
```

On R5:

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/24      0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 10.1.0.0/22       35.1.1.3                 0             0 300 400 i
*> 10.1.1.0/24       0.0.0.0                  0         32768 i
*> 10.1.2.0/24       35.1.1.3                               0 300 400 i
```

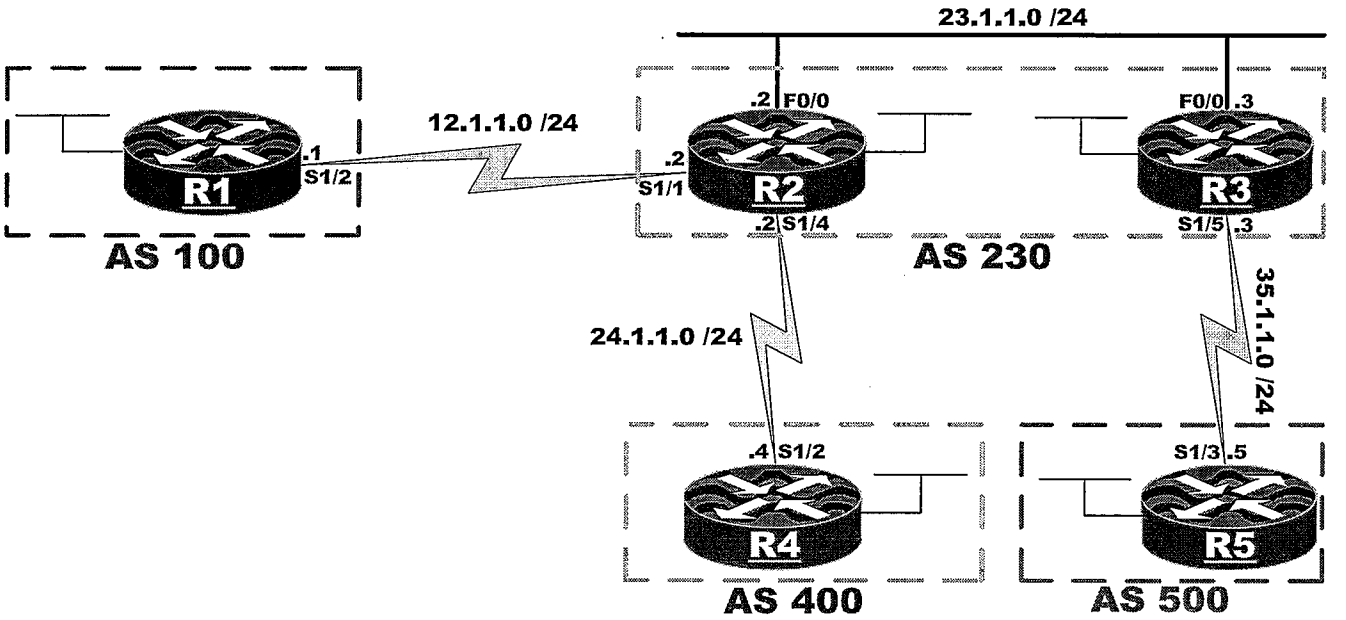
Task 8

Erase the startup configuration of the routers and reload them before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 6
The Community Attribute



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-6"→"Task-1-5"→"Lab-6-T1-5".

Task 1

Configure BGP peering based on the above topology, these routers should advertise their loopback0 interfaces in BGP. Ensure reachability between the loopback interfaces. DO NOT configure an IGP, or default routes to accomplish this task.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Netw 1.0.0.0
R1(config-router)#No au
R1(config-router)#Neighbor 12.1.1.2 remote-as 230
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R2:

```
R2(config)#Router bgp 230
R2(config-router)#netw 2.0.0.0
R2(config-router)#No au
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 23.1.1.3 remote-as 230
R2(config-router)#Neighbor 23.1.1.3 next-hop-self
R2(config-router)#Neighbor 24.1.1.4 remote-as 400
```

On R3:

```
R3(config)#Router bgp 230
R3(config-router)#Netw 3.0.0.0
R3(config-router)#No au
R3(config-router)#Neighbor 23.1.1.2 remote-as 230
R3(config-router)#Neighbor 23.1.1.2 next-hop-self
R3(config-router)#Neighbor 35.1.1.5 remote-as 500
```

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#Netw 4.0.0.0
R4(config-router)#No au
R4(config-router)#Neighbor 24.1.1.2 remote-as 230
```

On R5:

```
R5(config)#Router bgp 500
R5(config-router)#Netw 5.0.0.0
R5(config-router)#No au
R5(config-router)#Neighbor 35.1.1.3 remote-as 230
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 230 i
*> 3.0.0.0          12.1.1.2                               0 230 i
*> 4.0.0.0          12.1.1.2                               0 230 400 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*> 5.0.0.0          12.1.1.2                               0 230 500 i
```

On R2:

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          23.1.1.3                 0    100      0 i
*> 4.0.0.0          24.1.1.4                 0             0 400 i
*>i5.0.0.0          23.1.1.3                 0    100      0 500 i
```

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          23.1.1.2                 0    100      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          23.1.1.2                 0    100      0 400 i
*> 5.0.0.0          35.1.1.5                 0             0 500 i
```

On R4:

```
R4#Show ip bgp | B Netw

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          24.1.1.2                               0 230 100 i
*> 2.0.0.0          24.1.1.2                 0             0 230 i
*> 3.0.0.0          24.1.1.2                               0 230 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          24.1.1.2                               0 230 500 i
```

On R5:

```
R5#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          35.1.1.3                               0 230 100 i
*> 2.0.0.0          35.1.1.3                               0 230 i
*> 3.0.0.0          35.1.1.3                 0             0 230 i
*> 4.0.0.0          35.1.1.3                               0 230 400 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

To test the configuration:

On R1:

```
R1#Ping 2.2.2.2 source 1.1.1.1

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 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

R1#Ping 3.3.3.3 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms

R1#Ping 4.4.4.4 source 1.1.1.1

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 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60 ms

R1#Ping 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R3:

```
R3#Ping 4.4.4.4 source 3.3.3.3

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 3.3.3.3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
```

On R2:

```
R2#Ping 5.5.5.5 source 2.2.2.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

Task 2

Using the community attribute, configure R1 such that when it advertises network 1.0.0.0 /8 to R2 in AS 230, the network is not advertised to any of R2's IBGP or EBGP neighbor/s.

The community attribute is a numerical value that can be attached to a given prefix and advertised to a specific neighbor, once the neighbor receives the prefix, it will examine the community value and it will perform either filtering or use that value for route selection process.

By default, no community attribute/s is sent to any neighbor. To specify that a community attribute should be sent to a BGP neighbor, the "neighbor send-community" command is configured in the router config mode.

The well-known communities are as follows:

- **Internet – If assigned to a network/s, that network/s should be advertised to all neighbors.**
- **Local-as – If assigned to a network/s, that network/s should only be advertised within that AS.**
- **No-advertise – If assigned to a network/s. that network/s should not be advertised to any BGP neighbor.**
- **No-export – If assigned to a network/s, that network/s should not be advertised to an EBGP neighbor.**

On R1:

Note: Before configuring an access-list, always performs a "Show access-list" command to

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

ensure that one does not exist.

```
R1(config)#Access-list 1 permit 1.0.0.0 0.255.255.255
```

The access-list is used to identify the network; Prefix-lists can also be used for this purpose.

```
R1(config)#Route-map TEST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set community no-advertise

R1(config-route-map)#Route-map TEST permit 90
```

Note: The above route-map matches on the access-list and sets the community to "no-advertise", this well-known community attribute tells the receiving router not to advertise the prefix to any of its neighbor/s (IBGP or EBGP).

The "route-map TST permit 90" is the catchall route-map; it matches any network not matched with the match keyword in the "route-map TST permit 10".

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 12.1.1.2 send-community
R1(config-router)#Neighbor 12.1.1.2 route-map TEST out
```

In the above commands, we are sending the community and assigning the route-map to a given neighbor. The direction of the route-map specifies which router's decision should be influenced by this policy, if it should affect neighbor's decision, then, the direction of the route-map should be "out", but if the local router's decision should be influenced, then, the direction of the route-map should be "in".

To verify the configuration:

On R2:

```
R2#Sh ip bgp 1.0.0.0

BGP routing table entry for 1.0.0.0/8, version 18
Paths: (1 available, best #1, table default, not advertised to any peer)
  Not advertised to any peer
  100
    12.1.1.1 from 12.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-advertise
```

Since R2 does not advertise the network, R3 and the other EBGP and/or IBGP neighbor/s will not have any

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

knowledge of this prefix.

On R3:

```
R3#Show ip bgp 1.0.0.0

% Network not in table
```

On R4:

```
R4#Show ip bgp 1.0.0.0
% Network not in table
```

On R5:

```
R5#Show ip bgp 1.0.0.0
% Network not in table
```

Task 3

Configure R5 such that when it advertises its network 5.0.0.0 to R3 in AS 230, the routers in AS 230 **do not** advertise that network to any of their EBGP peer/s. **Do not** configure filtering on R3 to accomplish this task.

On R5:

```
R5(config)#IP Prefix-list NET-5 permit 5.0.0.0/8

R5(config)#Route-map TST permit 10
R5(config-route-map)#match ip addr prefix NET-5
R5(config-route-map)#set community no-export

R5(config-route-map)#Route-map TST permit 90

R5(config)#Router bgp 500
R5(config-router)#Neighbor 35.1.1.3 send-community
R5(config-router)#Neighbor 35.1.1.3 route-map TST out
```

This is another well-known community. In this case network 5.0.0.0 will only be advertised to the routers in AS 230. The routers in AS 230 will not advertise this network to any of their EBGP neighbor/s.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Remember....by default routers will strip the community attribute when they send the community in the route-map to a given neighbor (IBGP, or EBGP), therefore, in this case, R3 should be configured to send community to R2, or else, R2 will advertise that network to its EBGP peers.

On R3:

```
R3(config)#Router bgp 230
R3(config-router)#Neighbor 23.1.1.2 send-community
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 8
Paths: (1 available, best #1, table default, not advertised to EBGP peer)
  Advertised to update-groups:
     10
  500
    35.1.1.5 from 35.1.1.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-export
```

On R2:

```
R2#Show ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 19
Paths: (1 available, best #1, table default, not advertised to EBGP peer)
  Not advertised to any peer
  500
    23.1.1.3 from 23.1.1.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: no-export
```

On R4:

```
R4#Show ip bgp 5.0.0.0
% Network not in table
```

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip bgp 5.0.0.0
% Network not in table
```

Task 4

Configure R3 in AS 230 to advertise network 3.0.0.0 /8 to the routers in its own AS **only**. YOU MUST use a "Network" command to accomplish this task.

On R3:

```
R3(config)#Route-map TST permit 10
R3(config-route-map)#set community local-as

R3(config)#Router bgp 230
R3(config-router)#Network 3.0.0.0 route-map TST
R3(config-router)#Neighbor 23.1.1.2 send-community
```

To verify the configuration:

On R3:

```
R3#Sh ip bgp 3.0.0.0

BGP routing table entry for 3.0.0.0/8, version 9
Paths: (1 available, best #1, table default, not advertised outside local
AS)
  Advertised to update-groups:
     10
  Local
    0.0.0.0 from 0.0.0.0 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced,
local, best
      Community: local-AS
```

On R5:

```
R5#Show ip bgp 3.0.0.0

% Network not in table
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R2:

```
R2#Show ip bgp 3.0.0.0

BGP routing table entry for 3.0.0.0/8, version 20
Paths: (1 available, best #1, table default, not advertised outside local
AS)
  Not advertised to any peer
  Local
    23.1.1.3 from 23.1.1.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: local-AS
```

On R1:

```
R1#Show ip bgp 3.0.0.0

% Network not in table
```

In this topology the "No-export" and the "Local-as" have the same behavior; the "No-export" community does not look at the "BGP Confederation identifier" command, therefore, it will keep the prefix/es local in the AS number that the BGP router is configured in, whereas, the "Local-as" community looks at the "BGP Confederation identifier" command first, if it's configured, it will keep the prefix/es within that AS, but if the "BGP Confederation identifier" is NOT configured, it will keep the prefix/es local to the AS that the BGP is configured on.

Task 5

R1 is advertising network 1.0.0.0 which has an attached community attribute of "no-advertise" to R2 (task 2). Router R2 should be configured to advertise network 1.0.0.0 to all of its IBGP and EBGP peers. You should utilize a community attribute to accomplish this task.

On R2:

Before configuring R2, we should display the prefix in the BGP table:

```
R2#Show ip bgp 1.0.0.0

BGP routing table entry for 1.0.0.0/8, version 18
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Paths: (1 available, best #1, table default, not advertised to any peer)
  Not advertised to any peer
  100
    12.1.1.1 from 12.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: no-advertise
```

R2 can be configured to assign the "Internet" community to this network; when the "Internet" community is assigned to a network, that network will be advertised to all peers (IBGP or EBGP).

To configure:

On R2:

```
R2(config)#Access-list 1 permit 1.0.0.0

R2(config)#Route-map TST permit 10
R2(config-route-map)#match ip addr 1
R2(config-route-map)#set community internet

R2(config)#Router bgp 230
R2(config-router)#Neighbor 12.1.1.1 route-map TST in
```

To verify the configuration:

On R2:

```
R2#Show ip bgp 1.0.0.0

R2#sh ip bgp 1.0.0.0
BGP routing table entry for 1.0.0.0/8, version 21
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     6           7
  100
    12.1.1.1 from 12.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: internet
```

On R4:

```
R4#Show ip bgp | b Network
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          24.1.1.2                               0 230 100 i
*> 2.0.0.0          24.1.1.2                 0             0 230 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

On R5:

```
R5#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          35.1.1.3                               0 230 100 i
*> 2.0.0.0          35.1.1.3                               0 230 i
*> 4.0.0.0          35.1.1.3                               0 230 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 6

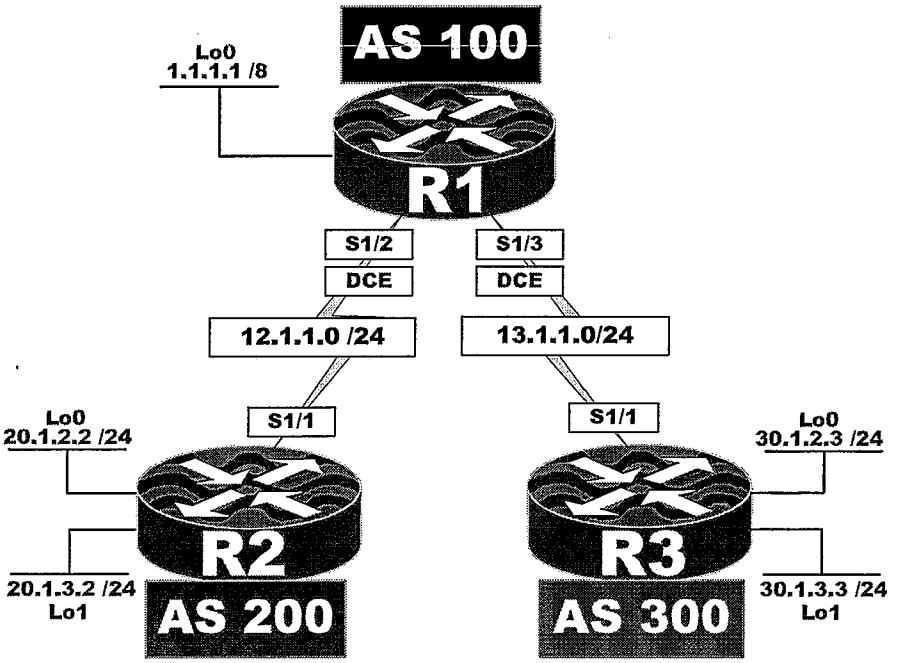
Erase the startup config and reload the routers before proceeding to the next task.

Task 7

Configure the routers according to the diagram below and use the IP addressing and AS numbers identified in the chart. These routers should advertise their Loopback interfaces in BGP.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-6"→"Task-7-8"→"Lab-6-T7-8".

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Netw 1.0.0.0
R1(config-router)#No au
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Neighbor 13.1.1.3 remote-as 300
```

On R2:

```
R2(config-if)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Network 20.1.2.0 mask 255.255.255.0
R2(config-router)#Network 20.1.3.0 mask 255.255.255.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
```

On R3:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config-if)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Netw 30.1.2.0 mask 255.255.255.0
R3(config-router)#Netw 30.1.3.0 mask 255.255.255.0
R3(config-router)#Neighbor 13.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 20.1.2.0/24      12.1.1.2                 0             0 200 i
*> 20.1.3.0/24      12.1.1.2                 0             0 200 i
*> 30.1.2.0/24      13.1.1.3                 0             0 300 i
*> 30.1.3.0/24      13.1.1.3                 0             0 300 i
```

On R2:

```
R2#Show ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 20.1.2.0/24      0.0.0.0                  0         32768 i
*> 20.1.3.0/24      0.0.0.0                  0         32768 i
*> 30.1.2.0/24      12.1.1.1                                 0 100 300 i
*> 30.1.3.0/24      12.1.1.1                                 0 100 300 i
```

On R3:

```
R3#Show ip bgp b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          13.1.1.1                 0             0 100 i
*> 20.1.2.0/24      13.1.1.1                                 0 100 200 i
*> 20.1.3.0/24      13.1.1.1                                 0 100 200 i
*> 30.1.2.0/24      0.0.0.0                  0         32768 i
*> 30.1.3.0/24      0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 8

Ensure that R1 uses AS 200 to connect to networks in subnet 2 (20.1.2.0 /24 and 30.1.2.0 /24) and AS 300 to connect to networks in subnet 3 (20.1.3.0 /24 and 30.1.3.0 /24). You must use community tags in AS 200 and 300 and neighbor commands on R1 to accomplish this task.

On R2:

```
R2(config)#Access-list 2 permit 20.1.2.0 0.0.0.255

R2(config)#Access-list 3 permit 20.1.3.0 0.0.0.255

R2(config)#Route-map TST permit 10
R2(config-route-map)#match ip addr 2
R2(config-route-map)#set community 2

R2(config)#Route-map TST permit 20
R2(config-route-map)#match ip addr 3
R2(config-route-map)#set community 3
```

Note: The above command "set community" tags the route/s identified in the access-list.

```
R2(config)#Route-map TST permit 90

R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 send-community
R2(config-router)#Neighbor 12.1.1.1 route-map TST out
```

On R3:

```
R3(config)#Access-list 2 permit 30.1.2.0 0.0.0.255

R3(config)#Access-list 3 permit 30.1.3.0 0.0.0.255

R3(config)#Route-map TST permit 10
R3(config-route-map)#match ip addr 2
R3(config-route-map)#set community 2

R3(config)#Route-map TST permit 20
R3(config-route-map)#match ip addr 3
R3(config-route-map)#set community 3

R3(config)#Route-map TST permit 90
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3(config)#Router bgp 300
R3(config-router)#Neighbor 13.1.1.1 send-community
R3(config-router)#Neighbor 13.1.1.1 route-map TST out
```

Note: R1 can display the routes via their assigned community tags:

On R1:

```
R1#Show ip bgp 20.1.2.0/24

BGP routing table entry for 20.1.2.0/24, version 8
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  200
    12.1.1.2 from 12.1.1.2 (20.1.3.2)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 2
```

```
R1#Show ip bgp community 2 | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.1.2.0/24      12.1.1.2                 0             0 200 i
*> 30.1.2.0/24      13.1.1.3                 0             0 300 i
```

Note: These are the routes that R2 and R3 tagged using community 2.

```
R1#Show ip bgp community 3 | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 20.1.3.0/24      12.1.1.2                 0             0 200 i
*> 30.1.3.0/24      13.1.1.3                 0             0 300 i
```

Note: These are the routes that R2 and R3 tagged using community 3.

On R1:

R1 is identifying the community tags using a community-list. This is like writing an access-list to identify a given route/s.

```
R1(config)#Ip community-list standard TST2 permit 2

R1(config)#Ip community-list standard TST3 permit 3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The communities are matched and the policy is assigned.

```
R1(config)#Route-map TEST permit 10
R1(config-route-map)#match community TST2
R1(config-route-map)#set ip next-hop 12.1.1.2

R1(config)#Route-map TEST permit 20
R1(config-route-map)#match community TST3
R1(config-route-map)#set ip next-hop 13.1.1.3

R1(config)#Route-map TEST permit 90
```

The policy is applied inbound, since R1's decision must be affected:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 12.1.1.2 route-map TEST in
R1(config-router)#Neighbor 13.1.1.3 route-map TEST in
```

The policies are applied to the neighbors using the "neighbor route-map" commands.

Do not try to test reachability to the network through the newly assigned next hop IP addresses; the purpose of this lab is to understand the scope of the community attributes and its uses.

To verify this lab enter "Show ip bgp" to see the next hop attribute.

```
R1#Show ip bgp | b Network
```

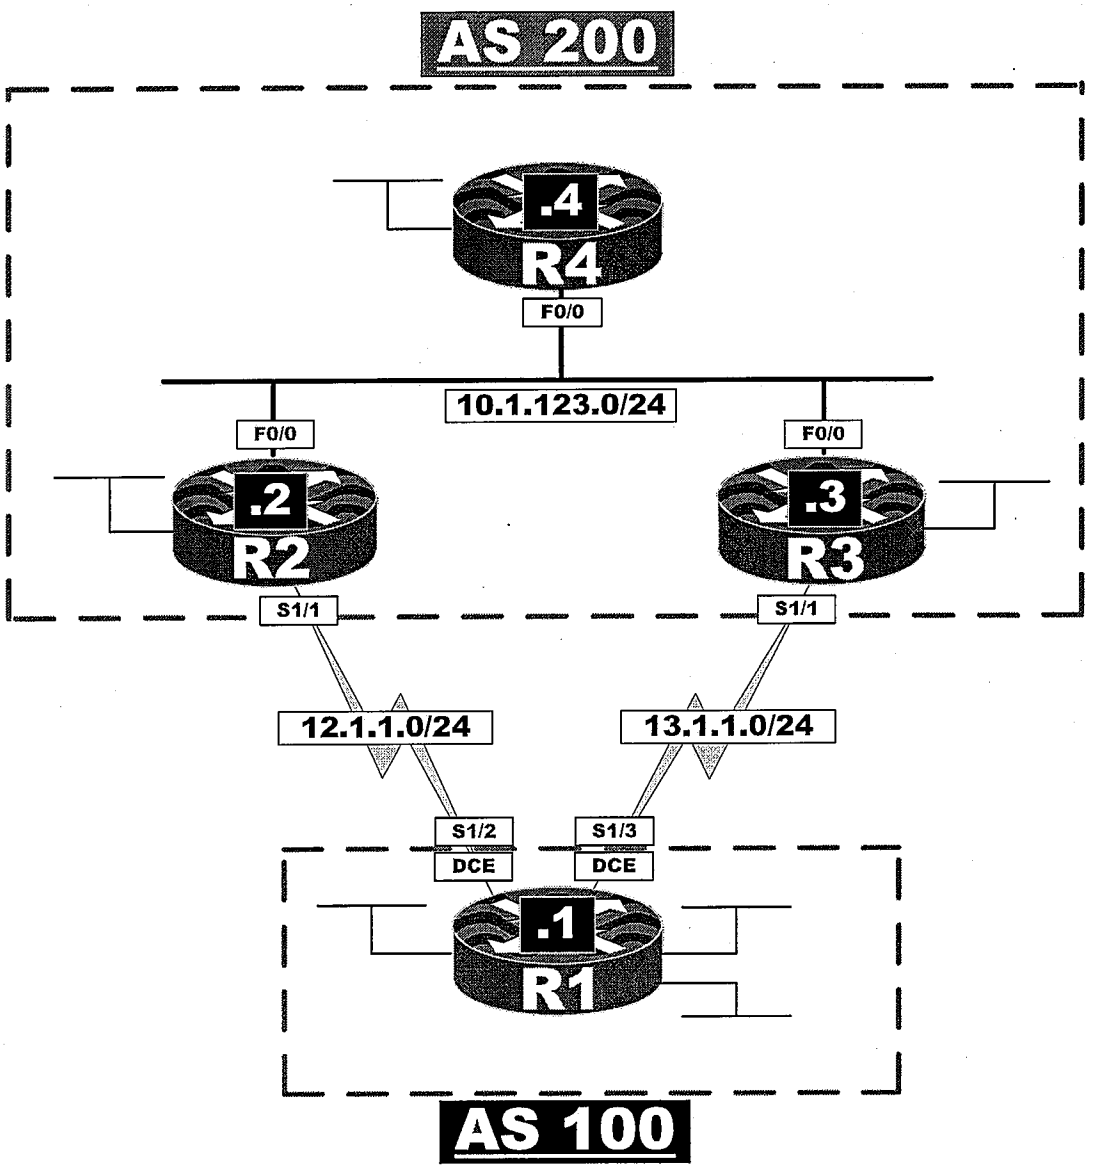
```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 20.1.2.0/24      12.1.1.2                 0             0 200 i
*> 20.1.3.0/24      13.1.1.3                 0             0 200 i
*> 30.1.2.0/24      12.1.1.2                 0             0 300 i
*> 30.1.3.0/24      13.1.1.3                 0             0 300 i
```

Task 9

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 7 – BGP Cost Community



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-7".

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Configure R1 in AS 100; this router should establish an EBGP peer session with R2 and R3 in AS 200. R1 should advertise its Lo1 and Lo2 interfaces in BGP. All BGP routers should use their loopback 0 interface as their router-id.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#bgp router-id 1.1.1.1
R1(config-router)#No auto-summary
R1(config-router)#Neigh 12.1.1.2 remote-as 200
R1(config-router)#Neigh 13.1.1.3 remote-as 200
R1(config-router)#Network 100.1.1.0 mask 255.255.255.0
R1(config-router)#Network 200.1.1.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#bgp router-id 2.2.2.2
R2(config-router)#No auto-summary
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#bgp router-id 3.3.3.3
R3(config-router)#No auto-summary
R3(config-router)#Neighbor 13.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp summary | b Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.1.1.2        4   200       4       5        3    0    0 00:01:07        0
13.1.1.3        4   200       4       5        3    0    0 00:00:20        0

R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 100.1.1.0/24     0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 200.1.1.0        0.0.0.0                  0         32768 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 100.1.1.0/24     12.1.1.1                 0             0 100 i
*> 200.1.1.0        12.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 100.1.1.0/24     13.1.1.1                 0             0 100 i
*> 200.1.1.0        13.1.1.1                 0             0 100 i
```

Task 2

Configure an IBGP peer session between R2, R3 and R4; these routers should establish their peer session based on their loopback 0 interface. You should provide NLRI using EIGRP AS 100 with auto summarization disabled.

On R2:

```
R2(config)#Router eigrp 100
R2(config-router)#No au
R2(config-router)#Network 2.2.2.2 0.0.0.0
R2(config-router)#Network 10.1.234.2 0.0.0.0

R2(config)#Router bgp 200
R2(config-router)#Neighbor 3.3.3.3 remote-as 200
R2(config-router)#Neighbor 3.3.3.3 Update-source Lo0
R2(config-router)#Neighbor 4.4.4.4 remote-as 200
R2(config-router)#Neighbor 4.4.4.4 update-source Lo0
```

On R3:

```
R3(config)#Router eigrp 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-router)#No au
R3(config-router)#Network 10.1.234.3 0.0.0.0
R3(config-router)#Network 3.3.3.3 0.0.0.0

R3(config)#Router bgp 200
R3(config-router)#Neighbor 4.4.4.4 remote-as 200
R3(config-router)#Neighbor 4.4.4.4 update-source Lo0
R3(config-router)#Neighbor 2.2.2.2 remote-as 200
R3(config-router)#Neighbor 2.2.2.2 update-source Lo0
```

On R4:

```
R4(config)#Router eigrp 100
R4(config-router)#No au
R4(config-router)#Network 10.1.234.4 0.0.0.0
R4(config-router)#Network 4.4.4.4 0.0.0.0

R4(config)#Router bgp 200
R4(config-router)#No au
R4(config-router)#Neighbor 2.2.2.2 remote-as 200
R4(config-router)#Neighbor 2.2.2.2 update-source lo0
R4(config-router)#Neighbor 3.3.3.3 remote-as 200
R4(config-router)#Neighbor 3.3.3.3 update-source Lo0
```

To verify the configuration:

On R2:

```
R2#Show ip bgp sum | b Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4   200       7       7        3    0    0 00:03:25        2
4.4.4.4         4   200       5       7        3    0    0 00:03:16        0
12.1.1.1        4   100       7       6        3    0    0 00:03:25        2

R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     13.1.1.1                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
* i200.1.1.0        13.1.1.1                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

```
R3#Show ip bgp sum | b Neigh
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
2.2.2.2	4	200	5	5	3	0	0	00:01:50	2
4.4.4.4	4	200	3	5	3	0	0	00:01:38	0
13.1.1.1	4	100	5	4	3	0	0	00:01:38	2

```
R3#Show ip bgp | B Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	i100.1.1.0/24	12.1.1.1	0	100	0	100 i
*>		13.1.1.1	0		0	100 i
*	i200.1.1.0	12.1.1.1	0	100	0	100 i
*>		13.1.1.1	0		0	100 i

On R4:

```
R4#Show ip bgp sum | b Neigh
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
2.2.2.2	4	200	8	7	3	0	0	00:03:23	2
3.3.3.3	4	200	8	7	3	0	0	00:03:07	2

```
R4#Show ip bgp | B Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	i100.1.1.0/24	12.1.1.1	0	100	0	100 i
*	i	13.1.1.1	0	100	0	100 i
*	i200.1.1.0	12.1.1.1	0	100	0	100 i
*	i	13.1.1.1	0	100	0	100 i

Task 3

Configure R2 and R3 such that the output of their "**Show ip bgp**" command resembles the following:

On R2:

```
R2#Show ip bgp | b Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>                  13.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     2.2.2.2                  0    100      0 100 i
*>                  13.1.1.1                 0             0 100 i
* i200.1.1.0        2.2.2.2                  0    100      0 100 i
*>                  13.1.1.1                 0             0 100 i
```

On R4:

```
R4#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
```

Let's check the BGP table of R2, R3 and R4:

On R2:

```
R2#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     13.1.1.1                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
* i200.1.1.0        13.1.1.1                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     12.1.1.1                 0    100      0 100 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*>                  13.1.1.1                 0             0 100 i
* i200.1.1.0        12.1.1.1                 0    100      0 100 i
*>                  13.1.1.1                 0             0 100 i

R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     12.1.1.1                 0    100      0 100 i
* i                 13.1.1.1                 0    100      0 100 i
* i200.1.1.0        12.1.1.1                 0    100      0 100 i
* i                 13.1.1.1                 0    100      0 100 i
```

Let's configure R2 first:

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 3.3.3.3 next-hop-self
R2(config-router)#Neighbor 4.4.4.4 next-hop-self
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#Neighbor 2.2.2.2 next-hop-self
R3(config-router)#Neighbor 4.4.4.4 next-hop-self
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     2.2.2.2                  0    100      0 100 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*>                    13.1.1.1                 0             0 100 i
* i200.1.1.0          2.2.2.2                  0    100      0 100 i
*>                    13.1.1.1                 0             0 100 i
```

On R4:

```
R4#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
* i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
```

Task 4

Using an extended community attribute, configure R4 such that if a "**Show ip bgp | b Network**" command is entered, the output of that show command matches the output of the following show command:

On R4:

```
R4#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
* i                 2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                 2.2.2.2                  0    100      0 100 i
```

In this scenario, the cost extended community attribute is used.

- **Cost is an extended community attribute**
- **It is a non-transitive extended community attribute that allows you to customize the local route's preference, which can influence the best path selection process.**
- **This attribute is applied by configuring the "set ext community cost" command, using a route-map. This command is configured with a cost community id (0-255) and a cost value (0-4.29 Billion) with a default cost value of 2.145 billion. The lower value has more preference, and the lower cost**

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

community id value is used as the tiebreaker.

On R4:

The following identifies the prefix:

```
R4(config)#IP Prefix-list NET permit 100.1.1.0/24
```

A route-map matches the prefix-list and applies the extcommunity cost attribute with two numbers, the first number is the community id and the second number is the community value.

```
R4(config)#Route-map TST permit 10
R4(config-route-map)#match ip addr prefix NET
R4(config-route-map)#Set extcommunity cost 1 1
R4(config)#Route-map TST per 90
```

Lastly, it's applied on a per neighbor basis using the "neighbor" command:

```
R4(config-route-map)#Router bgp 200
R4(config-router)#Neighbor 3.3.3.3 route-map TST in

R4#Clear ip bgp * in
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i100.1.1.0/24     3.3.3.3                  0    100      0 100 i
* i                  2.2.2.2                  0    100      0 100 i
* i200.1.1.0        3.3.3.3                  0    100      0 100 i
*>i                  2.2.2.2                  0    100      0 100 i
```

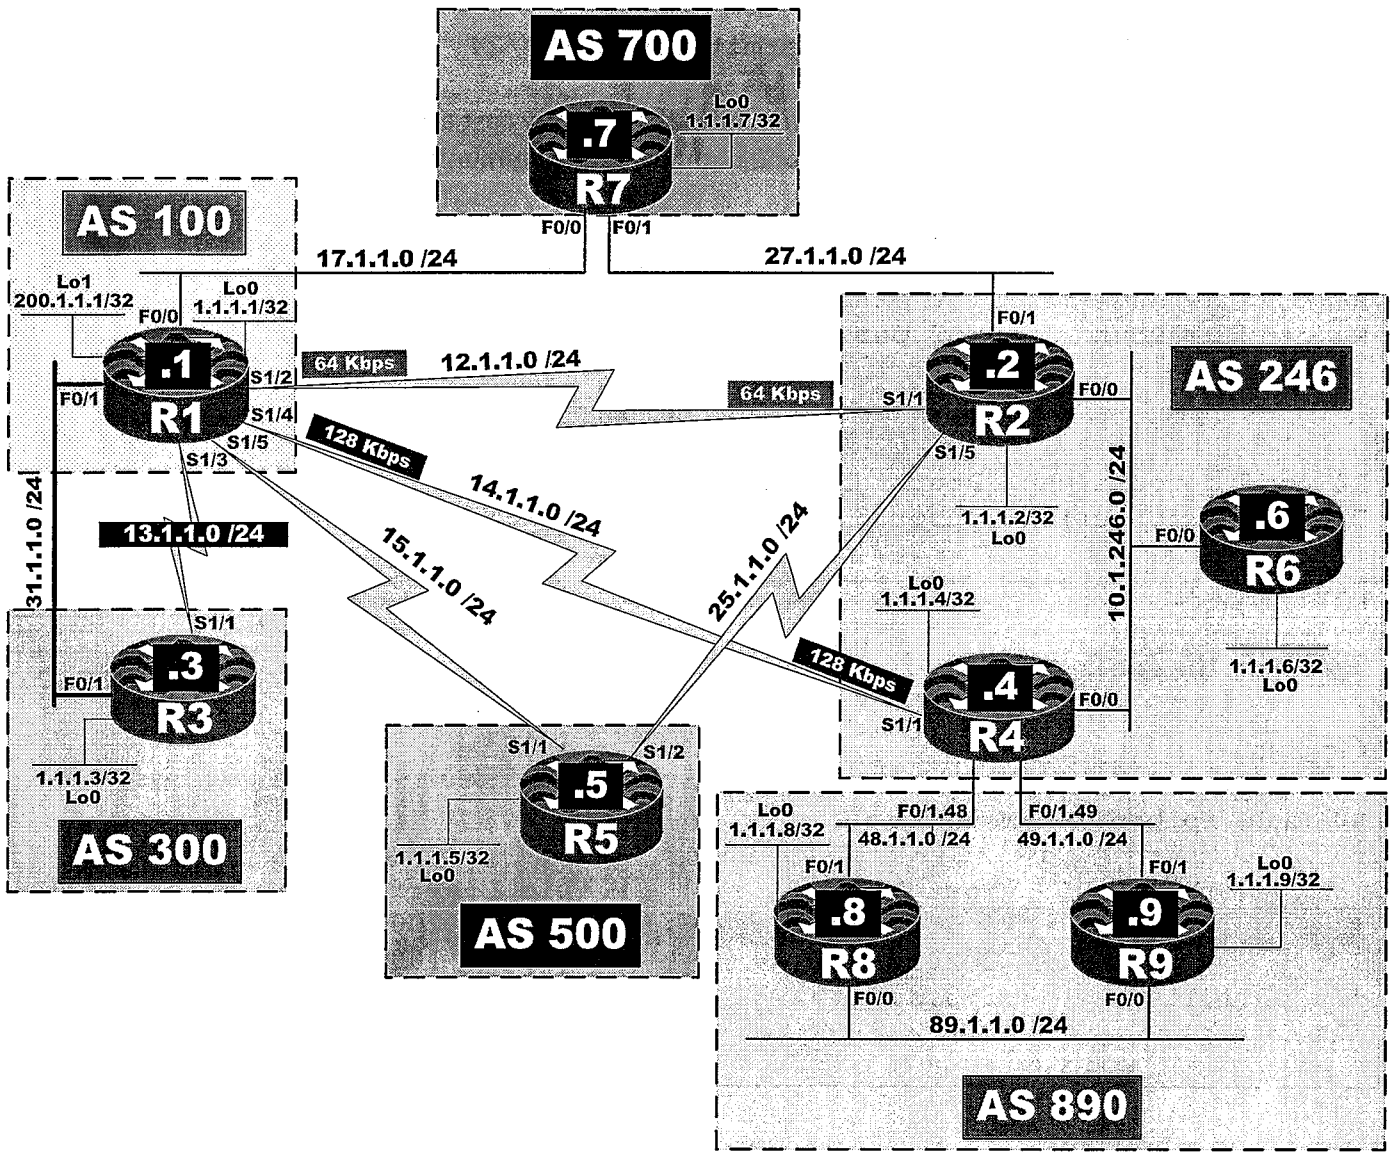
Task 5

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 8 - BGP Load Balancing



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-8".

Task 1

Configure BGP AS 100 on R1 and BGP AS 300 on R3. These routers should be configured such that they load share using their F0/1 and their Serial interfaces (R1's S1/3, and R3's S1/1) connections. You should use RIPv2 to provide NLRI.

In this topology, since the routers are directly connected the load balancing can be performed if the EBGP peer session is established based on the loopback interfaces of these two routers. In this scenario you must use an IGP to accomplish this task, the idea is that if both links have an equal cost to the loopback0's interface and the BGP's update source is set based on the loopback0 interface, we will do load sharing when the two BGP peers communicate. The reason RIPv2 was chosen is because OSPF and EIGRP use bandwith, and since the links have different bandwidths (100M versus 1M) it is much easier to load share equally between the two links using RIPv2.

On Both Routers:

```
Rx(config)#Router rip
Rx(config-router)#No au
Rx(config-router)#Ver 2
Rx(config-router)#Network 13.0.0.0
Rx(config-router)#Network 31.0.0.0
Rx(config-router)#Network 1.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gateway
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
R        1.1.1.3 [120/1] via 31.1.1.3, 00:00:24, FastEthernet0/1
                 [120/1] via 13.1.1.3, 00:00:17, Serial1/3
```

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

      1.0.0.0/32 is subnetted, 2 subnets
R        1.1.1.1 [120/1] via 31.1.1.1, 00:00:23, FastEthernet0/1
                 [120/1] via 13.1.1.1, 00:00:05, Serial1/1
```

We can see that there are multiple paths to the loopback 0 interface of the other router. In the following BGP configuration the update-source is based on the loopback0 interface of the peer.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 1.1.1.3 remote-as 300
R1(config-router)#Neighbor 1.1.1.3 update-source Lo0
R1(config-router)#Neighbor 1.1.1.3 ebgp-multihop 2
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Neighbor 1.1.1.1 remote-as 100
R3(config-router)#Neighbor 1.1.1.1 update-source lo0
R3(config-router)#Neighbor 1.1.1.1 ebgp-multihop 2
```

To test the configuration:

On Both Routers:

To test this scenario let's set the load sharing based on "Per-Packet" and test:

```
Rx(config)#Int F0/1
Rx(config-if)#IP load-sharing per-packet

Rx(config)#Int S1/3
Rx(config-if)#IP load-sharing per-packet
```

On R1:

```
R1#Show ip cef exact-route 1.1.1.1 1.1.1.3
1.1.1.1 -> 1.1.1.3 => IP adj out of Serial1/3
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Show ip cef exact-route 1.1.1.1 1.1.1.3
1.1.1.1 -> 1.1.1.3 => IP adj out of FastEthernet0/1, addr 31.1.1.3

R1#Show ip cef exact-route 1.1.1.1 1.1.1.3
1.1.1.1 -> 1.1.1.3 => IP adj out of Serial1/3

R1#Show ip cef exact-route 1.1.1.1 1.1.1.3
1.1.1.1  -> 1.1.1.3 => IP adj out of FastEthernet0/1, addr 31.1.1.3
```

Task 2

R2, R4 and R6 should be configured with IBGP peering in a full mesh manner in AS 246, and R8 and R9 in AS 890; the routers in AS 890 should use their Loopback 0 interface for this peering. R4 in AS 246, should be configured such that it performs load sharing between the two EBGP neighbors (R8 and R9) to reach 89.1.1.0 /24 network.

Let's configure the routers in AS 246 in a full mesh manner.

On R2:

```
R2(config)#Router bgp 246
R2(config-router)#No au
R2(config-router)#Neighbor 10.1.246.4 remote-as 246
R2(config-router)#Neighbor 10.1.246.6 remote-as 246
```

On R4:

```
R4(config)#Router bgp 246
R4(config-router)#No au
R4(config-router)#Neighbor 10.1.246.2 remote-as 246
R4(config-router)#Neighbor 10.1.246.6 remote-as 246
```

On R6:

```
R6(config)#Router bgp 246
R6(config-router)#No au
R6(config-router)#Neighbor 10.1.246.2 remote-as 246
R6(config-router)#Neighbor 10.1.246.4 remote-as 246
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Show ip bgp summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.246.4      4   246       6       7        1    0    0 00:03:29        0
10.1.246.6      4   246       4       4        1    0    0 00:02:37        0
```

Let's configure the routers in AS 890, but before the BGP peering can successfully be configured an IGP must be configured to advertise their loopback0 interfaces so R8 and R9 can have NLRI to each other's loopback0 interfaces. In this case OSPF was chosen.

On R8 and R9:

```
Rx(config)#Router ospf 1
Rx(config-router)#Network 0.0.0.0 0.0.0.0 area 0
```

On R8:

```
R8(config)#Router bgp 890
R8(config-router)#No au
R8(config-router)#Neighbor 1.1.1.9 remote-as 890
R8(config-router)#Neighbor 1.1.1.9 update-source Lo0

R8(config)#Router bgp 890
R8(config-router)#No au
R8(config-router)#Neighbor 1.1.1.9 remote-as 890
R8(config-router)#Neighbor 1.1.1.9 update-source Lo0
```

On R9:

```
R9(config)#Router bgp 890
R9(config-router)#No au
R9(config-router)#Neighbor 1.1.1.8 remote-as 890
R9(config-router)#Neighbor 1.1.1.8 update-source lo0

R9(config)#Router bgp 890
R9(config-router)#No au
R9(config-router)#Neighbor 1.1.1.8 remote-as 890
R9(config-router)#Neighbor 1.1.1.8 update-source Lo0
```

To verify the configuration:

On R9:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R9#Show ip bgp summary | b Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
1:1.1.8         4   890       5       5        1    0    0 00:01:12        0
```

Let's configure the peering between R4 in AS 246 and R8 and R9 in AS 890:

On R4:

```
R4(config)#Router bgp 246
R4(config-router)#Neighbor 48.1.1.8 remote-as 890
R4(config-router)#Neighbor 49.1.1.9 remote-as 890
```

Now we have to configure the peering from AS 890 (R8 and R9) with R4:

On R8:

```
R8(config)#Router bgp 890
R8(config-router)#Neighbor 48.1.1.4 remote-as 246
```

On R9:

```
R9(config)#Router bgp 890
R9(config-router)#Neighbor 49.1.1.4 remote-as 246
```

To verify the configuration:

On R4:

```
R4#Show ip bgp summary | B Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.246.2      4   246      14      12        1    0    0 00:10:00        0
10.1.246.6      4   246      13      11        1    0    0 00:09:36        0
48.1.1.8        4   890       5       3        1    0    0 00:01:18        0
49.1.1.9        4   890       4       3        1    0    0 00:00:49        0
```

Now we have to configure R8 and R9 to advertise network 89.1.1.0/24:

On R8 and R9:

```
Rx(config)#Router bgp 890
Rx(config-router)#Network 89.1.1.0 mask 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  89.1.1.0/24      49.1.1.9                 0             0 890 i
*>                  48.1.1.8                 0             0 890 i

R4#Show ip route bgp

     89.0.0.0/24 is subnetted, 1 subnets
B       89.1.1.0 [20/0] via 48.1.1.8, 00:02:52
```

NOTE: By default, BGP will NOT load balance amongst multiple equal cost paths. The "maximum-paths" command can be configured to change this behavior.

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#maximum-paths 2
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  89.1.1.0/24      49.1.1.9                 0             0 890 i
*>                  48.1.1.8                 0             0 890 i
```

The output of the above show command will NOT show the multiple active paths, but a "Show ip route" command will verify that R4 is performing load sharing:

```
R4#Show ip route bgp

     89.0.0.0/24 is subnetted, 1 subnets
B       89.1.1.0 [20/0] via 49.1.1.9, 00:01:28
                 [20/0] via 48.1.1.8, 00:05:20
```

To test the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R4:

```
R4(config)#Int F0/1.48
R4(config-subif)#IP load-sharing per-packet

R4(config)#Int F0/1.49
R4(config-subif)#IP load-sharing per-packet

R4#Show ip cef exact-route 48.1.1.4 89.1.1.8
48.1.1.4 -> 89.1.1.8 => IP adj out of FastEthernet0/1.48, addr 48.1.1.8

R4#Show ip cef exact-route 48.1.1.4 89.1.1.8
48.1.1.4 -> 89.1.1.8 => IP adj out of FastEthernet0/1.49, addr 49.1.1.9

R4#Show ip cef exact-route 49.1.1.4 89.1.1.8
49.1.1.4 -> 89.1.1.8 => IP adj out of FastEthernet0/1.48, addr 48.1.1.8

R4#Show ip cef exact-route 49.1.1.4 89.1.1.8
49.1.1.4 -> 89.1.1.8 => IP adj out of FastEthernet0/1.49, addr 49.1.1.9

R4#Show ip cef exact-route 49.1.1.4 89.1.1.9
49.1.1.4 -> 89.1.1.9 => IP adj out of FastEthernet0/1.48, addr 48.1.1.8

R4#Show ip cef exact-route 49.1.1.4 89.1.1.9
49.1.1.4 -> 89.1.1.9 => IP adj out of FastEthernet0/1.49, addr 49.1.1.9
```

Task 3

Configure R1 to advertise its Lo1 interface in BGP. Configure the appropriate router/routers such that R6 performs unequal cost load sharing for R1's Lo1 based on the bandwidth of R1→ R2 and R1→R4's connection. Configure the appropriate BGP peering to accomplish this task.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 12.1.1.2 remote-as 246
R1(config-router)#Neighbor 14.1.1.4 remote-as 246
R1(config-router)#Network 200.1.1.1 mask 255.255.255.255
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#Router bgp 246
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 10.1.246.4 next-hop-self
R2(config-router)#Neighbor 10.1.246.6 next-hop-self
```

On R4:

```
R4(config)#Router bgp 246
R4(config-router)#Neighbor 14.1.1.1 remote-as 100
R4(config-router)#Neighbor 10.1.246.6 next-hop-self
R4(config-router)#Neighbor 10.1.246.2 next-hop-self
```

To verify the configuration:

On R6:

```
R6#Show ip bgp | B Net
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i89.1.1.0/24      10.1.246.4               0    100      0 890 i
* i200.1.1.1/32     10.1.246.4               0    100      0 100 i
*>i                 10.1.246.2               0    100      0 100 i
```

On R2:

```
R2#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i89.1.1.0/24      10.1.246.4               0    100      0 890 i
* i200.1.1.1/32     10.1.246.4               0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```

On R4:

```
R4#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  89.1.1.0/24      49.1.1.9                 0             0 890 i
*>                  48.1.1.8                 0             0 890 i
*> 200.1.1.1/32     14.1.1.1                 0             0 100 i
* i                 10.1.246.2               0    100      0 100 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Let's check the routing table of R6:

On R6:

```
R6#Show ip route bgp | B Gate
Gateway of last resort is not set

      89.0.0.0/24 is subnetted, 1 subnets
B        89.1.1.0 [200/0] via 10.1.246.4, 00:29:44
      200.1.1.0/32 is subnetted, 1 subnets
B        200.1.1.1 [200/0] via 10.1.246.2, 00:04:58
```

To verify the configuration:

On R6:

NOTE: R6 only takes one of the paths to reach prefix 200.1.1.1 /32.

```
R6#Show ip bgp 200.1.1.1

BGP routing table entry for 200.1.1.1/32, version 6
Paths: (2 available, best #2, table default)
  Not advertised to any peer
  100
    10.1.246.4 from 10.1.246.4 (1.1.1.4)
      Origin IGP, metric 0, localpref 100, valid, internal
  100
    10.1.246.2 from 10.1.246.2 (1.1.1.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

The following steps outlines the configuration of unequal cost load sharing in BGP:

1. Enable the BGP dmzlink-bw feature:
   This is accomplished by configuring the "BGP dmzlink-bw" router configuration command, this must be configured on all routers (R2, R4 and R6).

2. Configure R2 and R4 to include their external link's bandwidth value in extended community; this is done so the bandwidth value can be propagated to the IBGP peer/s (R6). This is accomplished through the "Neighbor dmzlink-bw" router configuration command.

3. R6 must have the "maximum-path ibgp" configured so it can load sharre across two paths.

Remember, for this feature to work, R6 must have an equal IGP cost to R2 and R4, and all BGP attributes

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

must be identical or else the feature will NOT work.

Note the bandwidth of the links connecting the border routers R2, R4 to AS 100 is set to 64K and 128K respectively.

To configure Step 1:

On R2, R4 and R6:

```
RX(config)#Router bgp 246
RX(config-router)#BGP dmzlink-bw
```

To configure Step 2:

The following commands MUST be configured on the border routers of AS 246 so these routers can include the software bandwidth value (This is accomplished by configuring the bandwidth command) of their link to R1 and send it as an extended community to R6:

On R2:

```
R2(config)#Router bgp 246
R2(config-router)#Neighbor 12.1.1.1 dmzlink-bw
R2(config-router)#Neighbor 10.1.246.6 send-community extended
```

On R4:

```
R4(config)#Router bgp 246
R4(config-router)#Neighbor 14.1.1.1 dmzlink-bw
R4(config-router)#Neighbor 10.1.246.6 send-community extended
```

To configure Step 3:

Configure R6 to perform load sharing across two links, on this version and this platform the BGP routers can use this command to load share across 32 links:

On R6:

```
R6(config)#Router bgp 246
R6(config-router)#maximum-paths ibgp 2
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R6:

```
R6#Clear ip bgp *

R6#Show ip bgp 200.1.1.1

BGP routing table entry for 200.1.1.1/32, version 3
Paths: (2 available, best #2, table default)
Multipath: iBGP
  Not advertised to any peer
  100
    10.1.246.4 from 10.1.246.4 (1.1.1.4)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      DMZ-Link Bw 16 kbytes
  100
    10.1.246.2 from 10.1.246.2 (1.1.1.2)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      DMZ-Link Bw 8 kbytes
```

If the output of the following show command does NOT resemble the output on your router, and the traffic share count is set to one on both links, and you have saved the configuration on R6 and reloaded the router, then the problem is NOT you or your configuration, it is the IOS, I tried this on "c3725-adventerprisek9-mz.124-15.T11.bin" and it worked without any problems. This bug only affects this task.

```
R6#Show ip route 200.1.1.1

Routing entry for 200.1.1.1/32
  Known via "bgp 400", distance 200, metric 0
  Tag 100, type internal
  Last update from 10.1.246.4 00:00:54 ago
  Routing Descriptor Blocks:
    10.1.246.4, from 10.1.246.4, 00:00:54 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
  * 10.1.246.2, from 10.1.246.2, 00:00:54 ago
      Route metric is 0, traffic share count is 16
      AS Hops 1
      Route tag 100
```

To test the configuration:

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6(config)#int f0/0
R6(config-if)#ip load-sharing per-packet

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.2)

R6#sh ip cef exact-route 10.1.246.6 200.1.1.1
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
10.1.246.6      -> 200.1.1.1: FastEthernet0/0 (next hop 10.1.246.4)
```

Task 4

Configure R7 in AS 700, this router should establish an EBGP peer session with R2 in AS 246 and R1 in AS 100. R1 in AS 100 and R2 in AS 246 should each establish an EBGP session with R5 in AS 500.

R7 should advertise its Lo0 interface in AS 700.

Configure the appropriate router such that R5 in AS 500, load shares between R1 in AS 100 and R2 in AS 246 to reach the loopback 0 interface of R7 in AS 700.

On R7:

```
R7(config)#Router bgp 700
R7(config-router)#No au
R7(config-router)#Neighbor 17.1.1.1 remote-as 100
R7(config-router)#Neighbor 27.1.1.2 remote-as 246
R7(config-router)#Network 1.1.1.7 Mask 255.255.255.255
```

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 17.1.1.7 remote-as 700
R1(config-router)#Neighbor 15.1.1.5 remote-as 500
```

On R2:

```
R2(config)#Router bgp 246
R2(config-router)#Neighbor 27.1.1.7 remote-as 700
R2(config-router)#Neighbor 25.1.1.5 remote-as 500
```

On R5:

```
R5(config)#Router bgp 500
R5(config-router)#Neighbor 15.1.1.1 remote-as 100
R5(config-router)#Neighbor 25.1.1.2 remote-as 246
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R7:

```
R7#Show ip bgp summary | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
17.1.1.1        4   100       7       7        6    0    0 00:00:49        2
27.1.1.2        4   246      28      28        6    0    0 00:20:41        2
```

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.7/32       25.1.1.2                               0 246 700 i
*                   15.1.1.1                               0 100 700 i
*> 89.1.1.0/24      25.1.1.2                               0 246 890 i
*                   15.1.1.1                               0 100 246 890 i
*  200.1.1.1/32     25.1.1.2                               0 246 100 i
*>                  15.1.1.1                 0             0 100 i

R5#Show ip bgp 1.1.1.7

BGP routing table entry for 1.1.1.7/32, version 6
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     1
  246 700
    25.1.1.2 from 25.1.1.2 (1.1.1.2)
      Origin IGP, localpref 100, valid, external, best
  100 700
    15.1.1.1 from 15.1.1.1 (200.1.1.1)
      Origin IGP, localpref 100, valid, external

R5#Show ip route 1.1.1.7

Routing entry for 1.1.1.7/32
  Known via "bgp 500", distance 20, metric 0
  Tag 246, type external
  Last update from 25.1.1.2 00:14:57 ago
  Routing Descriptor Blocks:
  * 25.1.1.2, from 25.1.1.2, 00:14:57 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 246
      MPLS label: none
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The output of the above "Show ip bgp 1.1.1.7" command reveals that the prefix originated in AS 700, and it was advertised to R5 in AS 500 by AS 100 and 246. Note the Local-preference, AS-Path, Origin types are all the same on both paths.

In order to resolve this task, R5 in AS 500 MUST be configured with the "maximum-paths 2" router configuration mode, and of course a task can NOT be sweet without the use of a hidden command. In this case we have to use the "BGP Bestpath as-path multipath-relax" BGP router configuration command. This is a hidden command, meaning that a "?" will NOT reveal the command, and the "tab" key will NOT complete the command.

On R5:

```
R5(config)#Router bgp 500
R5(config-router)#maximum-paths 2
R5(config-router)#BGP bestpath as-path multipath-relax

R5#Clear ip bgp *
```

To verify the configuration:

On R5:

```
R5#Show ip route 1.1.1.7

Routing entry for 1.1.1.7/32
  Known via "bgp 500", distance 20, metric 0
  Tag 246, type external
  Last update from 15.1.1.1 00:00:35 ago
  Routing Descriptor Blocks:
  * 25.1.1.2, from 25.1.1.2, 00:00:35 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 246
      MPLS label: none
    15.1.1.1, from 15.1.1.1, 00:00:35 ago
      Route metric is 0, traffic share count is 1
      AS Hops 2
      Route tag 246
      MPLS label: none
```

NOTE: One reason that this command works is because the route tag which is based on the AS number matches on both routes. If this does not match, R5 will NOT be able to load share going across two different ASes to connect to AS 700.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show ip route bgp | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
B        1.1.1.7 [20/0] via 25.1.1.2, 00:02:29
                 [20/0] via 15.1.1.1, 00:02:29
      89.0.0.0/24 is subnetted, 1 subnets
B        89.1.1.0 [20/0] via 25.1.1.2, 00:02:29
      200.1.1.0/32 is subnetted, 1 subnets
B        200.1.1.1 [20/0] via 15.1.1.1, 00:02:29
```

To test the configuration:

On R5:

```
R5(config)#Int S1/1
R5(config-if)#ip load-sharing per-packet

R5(config)#Int S1/2
R5(config-if)#ip load-sharing per-packet

R5#Show ip cef exact-route 15.1.1.5 1.1.1.7
15.1.1.5 -> 1.1.1.7 => IP adj out of Serial1/1

R5#Show ip cef exact-route 15.1.1.5 1.1.1.7
15.1.1.5 -> 1.1.1.7 => IP adj out of Serial1/2

R5#Show ip cef exact-route 25.1.1.2 1.1.1.7
25.1.1.2 -> 1.1.1.7 => IP adj out of Serial1/1

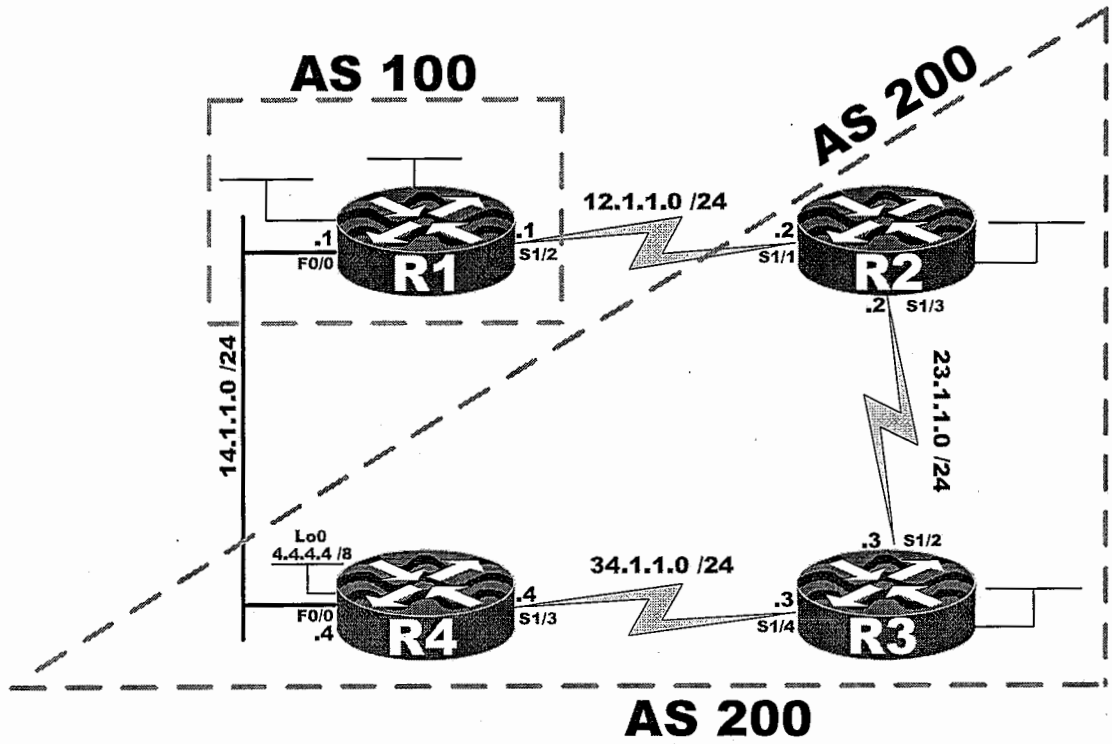
R5#Show ip cef exact-route 25.1.1.2 1.1.1.7
25.1.1.2 -> 1.1.1.7 => IP adj out of Serial1/2
```

Task 5

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

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 9 - Local-Preference -I



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-9".

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. R2 and R4 should have an EBGP peer session to R1 in AS 100. BGP routers should **only** advertise their loopback interface/s in BGP. Provide NLRI for the links in AS 200 using OSPF.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 14.1.1.4 remote-as 200
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Netw 11.0.0.0
R1(config-router)#Netw 1.0.0.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 23.1.1.3 remote-as 200
R2(config-router)#Neighbor 34.1.1.4 remote-as 200
R2(config-router)#Neighbor 23.1.1.3 next-hop-self
R2(config-router)#Neighbor 34.1.1.4 next-hop-self

R2(config)#Router ospf 1
R2(config-router)#Network 23.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Neighbor 34.1.1.4 remote-as 200
R3(config-router)#Neighbor 23.1.1.2 remote-as 200

R3(config)#Router ospf 1
R3(config-router)#Network 23.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 34.1.1.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#Router bgp 200
R4(config-router)#No au
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Neighbor 34.1.1.3 remote-as 200
R4(config-router)#Neighbor 23.1.1.2 remote-as 200
R4(config-router)#Neighbor 14.1.1.1 remote-as 100
R4(config-router)#Neighbor 34.1.1.3 next-hop-self
R4(config-router)#Neighbor 23.1.1.2 next-hop-self
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config)#Router ospf 1
R4(config-router)#Network 34.1.1.4 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                               0 200 i
*>                  12.1.1.2                 0             0 200 i
*  3.0.0.0          14.1.1.4                               0 200 i
*>                  12.1.1.2                               0 200 i
*  4.0.0.0          12.1.1.2                               0 200 i
*>                  14.1.1.4                 0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

Task 2

Ensure that the routers in AS 200 use R4 to reach network 1.0.0.0 /8 in AS 100. Local-Pref attribute must be used to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          34.1.1.4                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          23.1.1.3                 0    100      0 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
* i11.0.0.0         34.1.1.4                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         34.1.1.4                 0    100      0 100 i
*>i                23.1.1.2                 0    100      0 100 i
*>i2.0.0.0         23.1.1.2                 0    100      0 i
*> 3.0.0.0         0.0.0.0                  0         32768 i
*>i4.0.0.0         34.1.1.4                 0    100      0 i
* i11.0.0.0        34.1.1.4                 0    100      0 100 i
*>i                23.1.1.2                 0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         23.1.1.2                 0    100      0 100 i
*>                 14.1.1.1                 0             0 100 i
*>i2.0.0.0         23.1.1.2                 0    100      0 i
*>i3.0.0.0         34.1.1.3                 0    100      0 i
*> 4.0.0.0         0.0.0.0                  0         32768 i
* i11.0.0.0        23.1.1.2                 0    100      0 100 i
*>                 14.1.1.1                 0             0 100 i
```

NOTE: R2 and R3 are taking the R2-R1 link to connect to network 1.0.0.0 /8, whereas, R4 is taking the R4-R1 link. To configure the task:

On R4:

```
R4(config)#Access-list 1 permit 1.0.0.0 0.255.255.255

R4(config)#Route-map TST permit 10
R4(config-route-map)#match ip addr 1
R4(config-route-map)#set Local-preference 101
R4(config)#Route-map TST permit 90

R4(config-route-map)#Router bgp 200
R4(config-router)#Neighbor 14.1.1.1 route-map TST in

R4#Clear ip bgp * in
```

The local preference attribute is used to prefer an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS, hence an AS-Wide

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

attribute. If there are multiple exit points from the local AS, the local preference attribute is used to select the exit point for a specific or all prefix/es. Since the local preference attribute affects the routers within the AS, the route-map should be configured in the "in" direction.

Remember..... with local preference, the higher value has better preference. When R4 and R2 in AS 200 receive an update for network 1.0.0.0 /8 from R1 in AS 100, they set the local-preference to 100 and share this information with all of their IBGP peer/s. In the above configuration, R4 has changed the local preference for network 1.0.0.0/8 to 101, which is higher than the default value of 100, and since R4 will share this local preference value with all of its IBGP peers, they will all take the path through R4 to reach network 1.0.0.0/8 because it has a higher local preference.

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i1.0.0.0	34.1.1.4	0	101	0	100 i
*	12.1.1.1	0		0	100 i
*> 2.0.0.0	0.0.0.0	0		32768	i
*>i3.0.0.0	23.1.1.3	0	100	0	i
*>i4.0.0.0	34.1.1.4	0	100	0	i
* i11.0.0.0	34.1.1.4	0	100	0	100 i
*>	12.1.1.1	0		0	100 i

On R3:

```
R3#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i1.0.0.0	34.1.1.4	0	101	0	100 i
*>i2.0.0.0	23.1.1.2	0	100	0	i
*> 3.0.0.0	0.0.0.0	0		32768	i
*>i4.0.0.0	34.1.1.4	0	100	0	i
* i11.0.0.0	34.1.1.4	0	100	0	100 i
*>i	23.1.1.2	0	100	0	100 i

On R4:

```
R4#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*> 1.0.0.0              14.1.1.1                 0    101          0 100 i
*>i2.0.0.0              23.1.1.2                 0    100          0 i
*>i3.0.0.0              34.1.1.3                 0    100          0 i
*> 4.0.0.0              0.0.0.0                  0              32768 i
* i11.0.0.0             23.1.1.2                 0    100          0 100 i
*>                       14.1.1.1                 0                 0 100 i
```

Note: The routers in AS 200 take the R4-R1 link to connect to network 1.0.0.0/8.

To test the configuration

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 default-originate

R2#Clear ip bgp *
```

On R4:

```
R4(config)#Router bgp 200
R4(config-router)#Neighbor 14.1.1.1 default-originate

R4#Clear ip bgp *
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  0.0.0.0          14.1.1.4                 0             0 200 i
*>                  12.1.1.2                                0 200 i
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                                0 200 i
*>                  12.1.1.2                 0              0 200 i
*  3.0.0.0          14.1.1.4                                0 200 i
*>                  12.1.1.2                                0 200 i
*  4.0.0.0          12.1.1.2                                0 200 i
*>                  14.1.1.4                 0              0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R2:

```
R2#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 23.1.1.3 16 msec 16 msec 16 msec
  2 34.1.1.4 28 msec 28 msec 28 msec
  3 14.1.1.1 20 msec *  20 msec
```

On R3:

```
R3#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 34.1.1.4 12 msec 16 msec 16 msec
  2 14.1.1.1 24 msec *  24 msec
```

Task 3

Ensure that the routers in AS 200 use R2 to reach network 11.0.0.0 /8. Local-Pref attribute must be used to accomplish this task.

Before this attribute is configured, the existing BGP table of the routers in AS 200 should be examined:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
   0.0.0.0          0.0.0.0                                0 i
*>i1.0.0.0          34.1.1.4                 0    101      0 100 i
*                   12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          23.1.1.3                 0    100      0 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
* i11.0.0.0         34.1.1.4                 0    100      0 100 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*>                    12.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.4                 0    101      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
* i11.0.0.0         34.1.1.4                 0    100      0 100 i
*>i                 23.1.1.2                 0    100      0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          14.1.1.1                 0    101      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*>i3.0.0.0          34.1.1.3                 0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         23.1.1.2                 0    100      0 100 i
*>                  14.1.1.1                 0             0 100 i
```

Note: R4 is taking the R4-R1 link to connect to network 11.0.0.0, whereas, the other routers in AS 200 are taking the R1-R2 link. Let's configure R2 for this task, in the previous task an access-list was used, in this task a prefix-list is used so we can see how they are both configured:

On R2:

```
R2(config)#Ip prefix-list NET-11 permit 11.0.0.0/8

R2(config)#Route-map TST permit 10
R2(config-route-map)#match ip addr prefix NET-11
R2(config-route-map)#set Local-preference 101
R2(config)#Route-map TST permit 90

R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 route-map TST in

R2#Clear ip bgp * in
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
   0.0.0.0          0.0.0.0                                  0 i
*>i1.0.0.0          34.1.1.4                 0    101       0 100 i
*                   12.1.1.1                 0              0 100 i
*> 2.0.0.0          0.0.0.0                  0          32768 i
*>i3.0.0.0          23.1.1.3                 0    100       0 i
*>i4.0.0.0          34.1.1.4                 0    100       0 i
*> 11.0.0.0         12.1.1.1                 0    101       0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.4                 0    101       0 100 i
*>i2.0.0.0          23.1.1.2                 0    100       0 i
*> 3.0.0.0          0.0.0.0                  0          32768 i
*>i4.0.0.0          34.1.1.4                 0    100       0 i
*>i11.0.0.0         23.1.1.2                 0    101       0 100 i

R3#Traceroute 11.1.1.1 numeric

Type escape sequence to abort.
Tracing the route to 11.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 23.1.1.2 16 msec 16 msec 16 msec
  2 12.1.1.1 20 msec *  20 msec
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          14.1.1.1                 0    101       0 100 i
*>i2.0.0.0          23.1.1.2                 0    100       0 i
*>i3.0.0.0          34.1.1.3                 0    100       0 i
*> 4.0.0.0          0.0.0.0                  0          32768 i
*>i11.0.0.0         23.1.1.2                 0    101       0 100 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*                    14.1.1.1                 0             0 100 i

R4#Traceroute 11.1.1.1 numeric

Type escape sequence to abort.
Tracing the route to 11.1.1.1

  1 34.1.1.3 12 msec 16 msec 16 msec
  2 23.1.1.2 28 msec 28 msec 32 msec
  3 12.1.1.1 16 msec *  16 msec
```

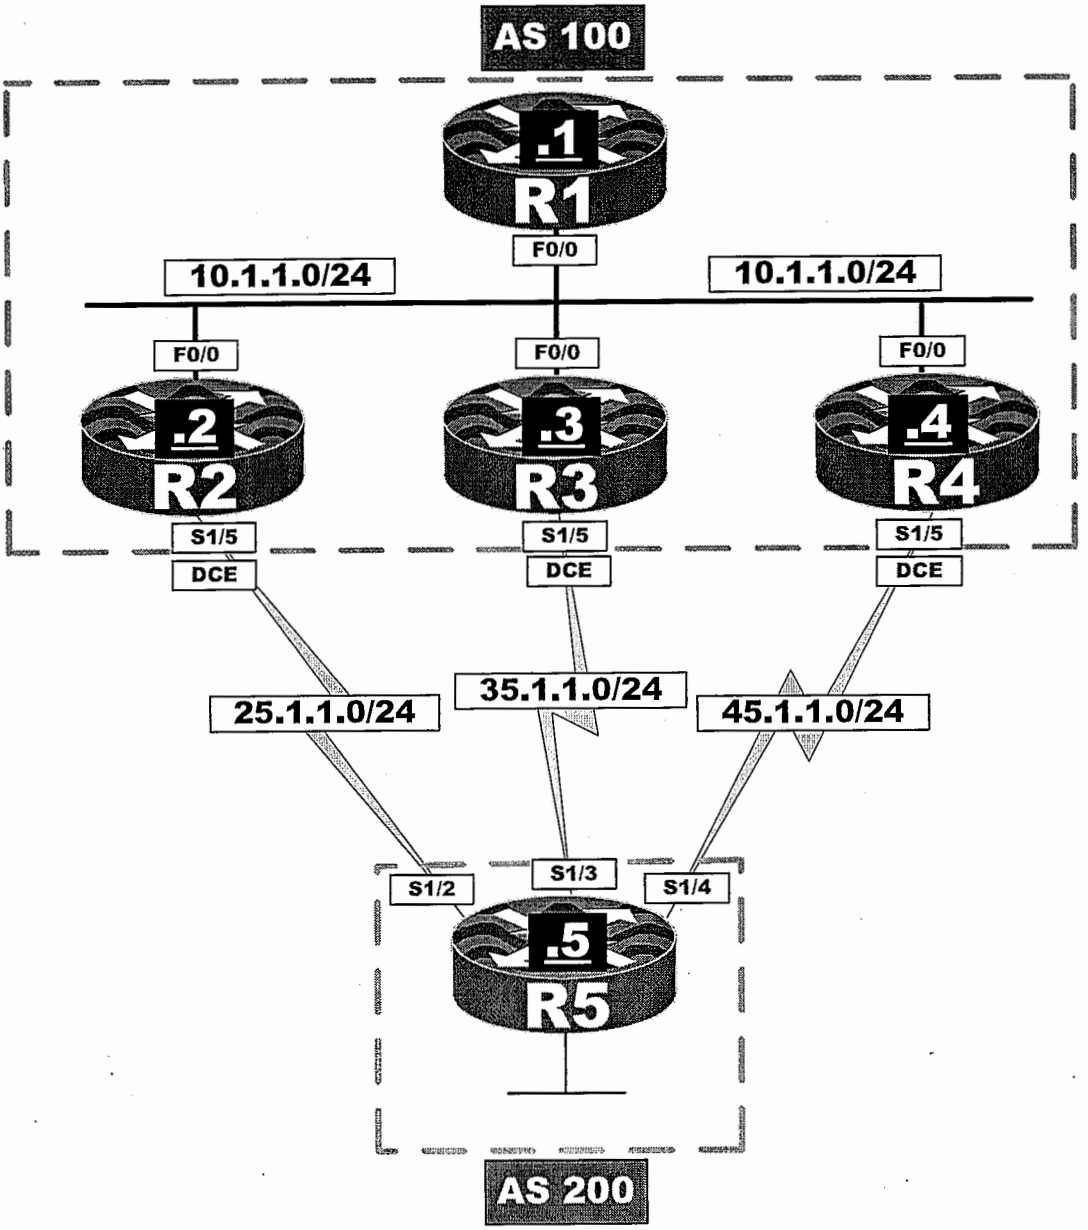
Note: The routers in AS 200 connect to network 11.0.0.0 /8 by going through R2-R1 link.

Task 4

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 10 – BGP Local-Preference - II



© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-10".

Task 1

Configure peering according to the diagram. R5 should advertise its Loopback0 interface in AS 200.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 10.1.1.2 remote 100
R1(config-router)#Neighbor 10.1.1.3 remote 100
R1(config-router)#Neighbor 10.1.1.4 remote 100
```

On R2:

```
R2(config)#Router bgp 100
R2(config-router)#No au
R2(config-router)#Neighbor 10.1.1.1 remote 100
R2(config-router)#Neighbor 10.1.1.3 remote 100
R2(config-router)#Neighbor 10.1.1.4 remote 100
R2(config-router)#Neighbor 25.1.1.5 remote 200
```

On R3:

```
R3(config)#Router bgp 100
R3(config-router)#No au
R3(config-router)#Neighbor 10.1.1.1 remote 100
R3(config-router)#Neighbor 10.1.1.2 remote 100
R3(config-router)#Neighbor 10.1.1.4 remote 100
R3(config-router)#Neighbor 35.1.1.5 remote 200
```

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#No au
R4(config-router)#Neighbor 10.1.1.1 remote 100
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config-router)#Neighbor 10.1.1.2 remote 100

R4(config-router)#Neighbor 10.1.1.3 remote 100
R4(config-router)#Neighbor 45.1.1.5 remote 200
```

On R5:

```
R5(config)#Router bgp 200
R5(config-router)#No au
R5(config-router)#Neighbor 25.1.1.2 remote 100
R5(config-router)#Neighbor 35.1.1.3 remote 100
R5(config-router)#Neighbor 45.1.1.4 remote 100
R5(config-router)#Network 5.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         35.1.1.5                 0    100      0 200 i
* i                45.1.1.5                 0    100      0 200 i
* i                25.1.1.5                 0    100      0 200 i
```

We can see that R1 does NOT have NLRI to the next-hop IP address to reach network 5.0.0.0/8. To resolve this problem we can configure R2 – R4 to change the next-hop IP address using the "Next-hop-self" bgp neighbor command. Let's configure this on R2, R3 and R4:

On R2, R3 and R4:

```
Rx(config)#Router bgp 100
Rx(config-router)#Neighbor 10.1.1.1 next-hop-self
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.1.3                 0    100      0 200 i
* i                10.1.1.4                 0    100      0 200 i
*>i               10.1.1.2                 0    100      0 200 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 2

Configure R2, R3 and R4 such that R1 takes R4 as the primary and R3 as the backup, if R4 and R3 are both down, then, R1 should take R2 to reach network 5.0.0.0 /8. You must use local preference and minimum number of commands to accomplish this task.

On R3:

```
R3(config)#Router bgp 100
R3(config-router)#bgp default local-preference 300

R3#Clear ip bgp * in
```

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#bgp default local-preference 400

R4#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.1.3                 0    300      0 200 i
*>i               10.1.1.4                 0    400      0 200 i
* i               10.1.1.2                 0    100      0 200 i

R1#Show ip bgp 5.0.0.0

BGP routing table entry for 5.0.0.0/8, version 3
Paths: (3 available, best #2, table default)
  Not advertised to any peer
  200
    10.1.1.3 from 10.1.1.3 (35.1.1.3)
      Origin IGP, metric 0, localpref 300, valid, internal
  200
    10.1.1.4 from 10.1.1.4 (45.1.1.4)
      Origin IGP, metric 0, localpref 400, valid, internal, best
  200
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
    10.1.1.2 from 10.1.1.2 (25.1.1.2)
      Origin IGP, metric 0, localpref 100, valid, internal
```

To test the configuration:

On R2:

```
R2(config)#Router bgp 100
R2(config-router)#Neighbor 25.1.1.5 default-originate

R2#Clear ip bgp *
```

On R3:

```
R3(config)#Router bgp 100
R3(config-router)#Neighbor 35.1.1.5 default-originate

R3#Clear ip bgp *
```

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#Neighbor 45.1.1.5 default-originate

R4#Clear ip bgp *
```

On R1:

```
R1#Traceroute 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.4 0 msec 4 msec 0 msec
  2 45.1.1.5 16 msec *  12 msec
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 3

Remove the configuration from the previous task and re-configure the same task using another method. **Do not** use neighbor 25.1.1.5, 35.1.1.5 or 45.1.1.5. You should use local preference to accomplish this task.

Let's remove the commands from the previous task first:

On R3:

```
R3(config)#Router bgp 100
R3(config-router)#No bgp default local-preference 300

R3#Clear ip bgp * in
```

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#No bgp default local-preference 400

R4#Clear ip bgp * in
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
* i5.0.0.0         10.1.1.4                 0    100      0 200 i
* i                10.1.1.3                 0    100      0 200 i
*>i               10.1.1.2                 0    100      0 200 i
```

To configure the task:

On R3:

```
R3(config)#Route-map TST permit 10
R3(config-route-map)#set local-preference 300
R3(config-route-map)#route-map TST permit 90

R3(config)#Router bgp 100
R3(config-router)#Neighbor 10.1.1.1 route-map TST OUT
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R3#Clear ip bgp * in
```

On R4:

```
R4(config)#Route-map TST permit 10
R4(config-route-map)#set local-preference 400
R4(config-route-map)#route-map TST permit 90

R4(config)#Router bgp 100
R4(config-router)#Neighbor 10.1.1.1 route-map TST OUT

R4#Clear ip bgp * in
```

To test the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0          10.1.1.4                 0    400      0 200 i
* i                  10.1.1.3                 0    300      0 200 i
* i                  10.1.1.2                 0    100      0 200 i

R1#Traceroute 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.4 0 msec 4 msec 0 msec
  2 45.1.1.5 16 msec *  12 msec
```

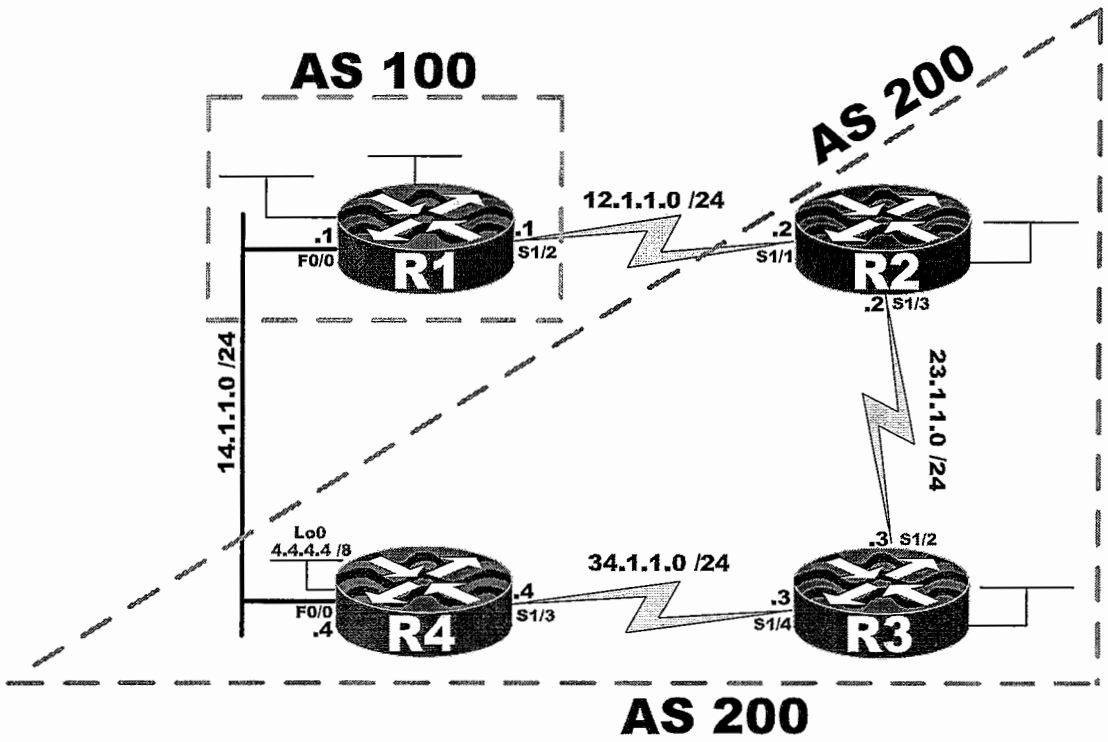
Task 4

Erase the startup configuration on all routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 11 - The AS-Path Attribute



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-11".

Task 1

Configure R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. R2 and R4 should have an EBGP peer session with R1 in AS 100. These routers should ONLY advertise their loopback interface/s in BGP. Provide NLRI for the links using OSPF.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-router)#Neighbor 14.1.1.4 remote-as 200
R1(config-router)#Network 1.0.0.0
R1(config-router)#Network 11.0.0.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 23.1.1.3 remote-as 200
R2(config-router)#Neighbor 23.1.1.3 next-hop-self
R2(config-router)#Neighbor 34.1.1.4 next-hop-self
R2(config-router)#Neighbor 34.1.1.4 remote-as 200

R2(config)#Router ospf 1
R2(config-router)#Network 23.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Neighbor 34.1.1.4 remote-as 200
R3(config-router)#Neighbor 23.1.1.2 remote-as 200

R3(config)#Router ospf 1
R3(config-router)#Network 23.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 34.1.1.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#Router bgp 200
R4(config-router)#No au
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Neighbor 34.1.1.3 remote-as 200
R4(config-router)#Neighbor 34.1.1.3 next-hop-self
R4(config-router)#Neighbor 23.1.1.2 remote-as 200
R4(config-router)#Neighbor 23.1.1.2 next-hop-self
R4(config-router)#Neighbor 14.1.1.1 remote-as 100

R4(config)#Router ospf 1
R4(config-router)#Network 34.1.1.4 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                               0 200 i
*>                  12.1.1.2                 0             0 200 i
*  3.0.0.0          12.1.1.2                               0 200 i
*>                  14.1.1.4                               0 200 i
*  4.0.0.0          14.1.1.4                 0             0 200 i
*>                  12.1.1.2                               0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

On R2:

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          14.1.1.1                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          23.1.1.3                 0    100      0 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
* i11.0.0.0         14.1.1.1                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          34.1.1.4                 0    100      0 100 i
*>i                 23.1.1.2                 0    100      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
* i11.0.0.0         34.1.1.4                 0    100      0 100 i
*>i                 23.1.1.2                 0    100      0 100 i
```

On R4:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0         23.1.1.2                 0    100      0 100 i
*>                 14.1.1.1                 0             0 100 i
*>i2.0.0.0        23.1.1.2                 0    100      0 i
*>i3.0.0.0        34.1.1.3                 0    100      0 i
*> 4.0.0.0        0.0.0.0                  0         32768 i
* i11.0.0.0       23.1.1.2                 0    100      0 100 i
*>                 14.1.1.1                 0             0 100 i
```

Task 2

Configure R1 in AS 100 such that routers in AS 200 use the link through R4-R1 to reach its network 1.0.0.0 /8. Use the AS-Path attribute to accomplish this task.

The AS-path attribute can be manipulated to accomplish this task. This attribute can be manipulated by prepending the AS numbers to existing AS paths. Typically, you perform AS-Path prepending on the outgoing EBGP updates over the non-desired return path, making the AS-Path over the non-desired link become longer than the desired path, and as a result making the desired path the primary.

On R1:

```
R1(config)#Access-list 1 permit 1.0.0.0 0.255.255.255

R1(config)#Route-map TST1 permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set as-path prepend 100
R1(config)#route-map TST1 permit 90

R1(config)#Router bgp 100
R1(config-router)#Neighbor 12.1.1.2 route-map TST1 out

R1#Clear ip bgp * Out
```

To verify the configuration:

On R2:

```
R2#Sh ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.4                 0    100      0 100 i
*                   12.1.1.1                 0             0 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          23.1.1.3                 0    100      0 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
* i11.0.0.0         34.1.1.4                 0    100      0 100 i
*>                  12.1.1.1                 0             0 100 i
```

To test the configuration:

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 default-originate

R2#Clear ip bgp * Out
```

On R4:

```
R4(config)#Router bgp 200
R4(config-router)#Neighbor 14.1.1.1 default-originate

R4#Clear ip bgp * out
```

On R2:

```
R2#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 23.1.1.3 12 msec 12 msec 16 msec
  2 34.1.1.4 28 msec 32 msec 28 msec
  3 14.1.1.1 20 msec *  20 msec
```

On R3:

```
R3#Sh ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.4                 0    100      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*>i4.0.0.0              34.1.1.4                 0    100      0 i
* i11.0.0.0             34.1.1.4                 0    100      0 100 i
*>i                     23.1.1.2                 0    100      0 100 i

R3#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 34.1.1.4 16 msec 16 msec 12 msec
  2 14.1.1.1 20 msec *  24 msec
```

On R4:

```
R4#Sh ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          14.1.1.1                 0             0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*>i3.0.0.0          34.1.1.3                 0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         23.1.1.2                 0    100      0 100 i
*>                  14.1.1.1                 0             0 100 i

R4#Traceroute 1.1.1.1 numeric

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 14.1.1.1 0 msec *  0 msec
```

Note the AS-path attribute is used to influence the degree of preference in another AS. R2, R3 and R4 will go through R4 to reach network 1.0.0.0 /8 because going through R4 they have to traverse through less AS Hops.

Task 3

Configure R1 in AS 100 such that the routers in AS 200 use the link through R2-R1 to reach network 11.0.0.0 /8. Use the AS-Path attribute and an IP Prefix-list to accomplish this task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#IP Prefix-list NET11 permit 11.0.0.0/8

R1(config)#Route-map TST11 permit 10
R1(config-route-map)#match ip addr prefix NET11
R1(config-route-map)#set as-path prepend 100 100
R1(config)#route-map TST11 permit 90

R1(config-route-map)#Router bgp 100
R1(config-router)#Neighbor 14.1.1.4 route-map TST11 out

R1#Clear ip bgp * Out
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
   0.0.0.0          0.0.0.0                                0 i
*>i1.0.0.0          34.1.1.4                 0    100      0 100 i
*                   12.1.1.1                 0             0 100 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*>i3.0.0.0          23.1.1.3                 0    100      0 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
*> 11.0.0.0         12.1.1.1                 0             0 100 i

R2#Traceroute 11.1.1.1

Type escape sequence to abort.
Tracing the route to 11.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.1 16 msec *  12 msec
```

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.4                 0    100      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 3.0.0.0           0.0.0.0                  0         32768 i
*>i4.0.0.0           34.1.1.4                 0    100      0 i
*>i11.0.0.0          23.1.1.2                 0    100      0 100 i

R3#Traceroute 11.1.1.1

Type escape sequence to abort.
Tracing the route to 11.1.1.1

  1 23.1.1.2 12 msec 12 msec 16 msec
  2 12.1.1.1 28 msec *  28 msec
```

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          14.1.1.1                 0             0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*>i3.0.0.0          34.1.1.3                 0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*>i11.0.0.0         23.1.1.2                 0    100      0 100 i
*                   14.1.1.1                 0             0 100 100 100 i

R4#Traceroute 11.1.1.1

Type escape sequence to abort.
Tracing the route to 11.1.1.1

  1 34.1.1.3 12 msec 12 msec 16 msec
  2 23.1.1.2 28 msec 28 msec 32 msec
  3 12.1.1.1 40 msec *  40 msec
```

Note all the routers will take the R2-R1 link to connect to network 11.0.0.0/8.

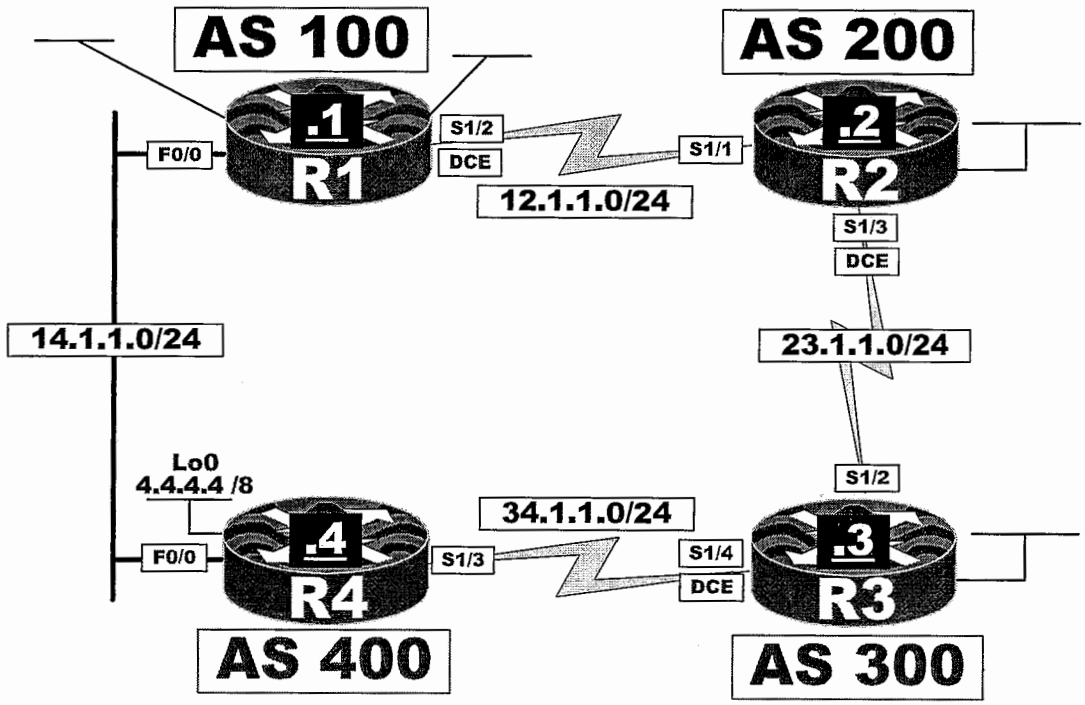
Task 4

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 12 - The Weight Attribute



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-12".

The initial configuration file configures RIPv2 on all routers advertising the links to provide NLRI.

Task 1

Configure R1 in AS 100 to establish an EBGP peer session with R2 in AS 200 and R4 in AS 400.
Router R2 should establish EBGP peer sessions with R1 in AS 100 and R3 in AS 300.
Router R3 should establish EBGP peer sessions with R2 in AS 200 and R4 in AS 400.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Router R4 should establish EBGP peer sessions with R1 in AS 100 and R3 in AS 300.
The BGP routers should ONLY advertise their loopback/s in BGP.

On R1:

```
R1(config-if)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Netw 11.0.0.0
R1(config-router)#Neighbor 14.1.1.4 remote-as 400
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 23.1.1.3 remote-as 300
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Neighbor 34.1.1.4 remote-as 400
R3(config-router)#Neighbor 23.1.1.2 remote-as 200
```

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#No au
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Neighbor 34.1.1.3 remote-as 300
R4(config-router)#Neighbor 14.1.1.1 remote-as 100
```

To verify the configuration:

On R4:

```
R4#Sh ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*  1.0.0.0          34.1.1.3                               0 300 200 100 i
*>                  14.1.1.1                 0             0 100 i
*> 2.0.0.0          34.1.1.3                               0 300 200 i
*                   14.1.1.1                               0 100 200 i
*> 3.0.0.0          34.1.1.3                 0             0 300 i
*                   14.1.1.1                               0 100 200 300 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*  11.0.0.0         34.1.1.3                               0 300 200 100 i
*>                  14.1.1.1                 0             0 100 i
```

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                               0 400 300 200 i
*>                  12.1.1.2                 0             0 200 i
*  3.0.0.0          14.1.1.4                               0 400 300 i
*>                  12.1.1.2                               0 200 300 i
*> 4.0.0.0          14.1.1.4                 0             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

Task 3

Configure R1 in AS 100 to use AS 200 to reach all the prefixes within this topology; you must use the "Weight" attribute to accomplish this task.

The weight attribute is a Cisco-defined attribute that is local to the router. This attribute is NOT advertised to any BGP neighbor/s. If there is more than one route to a given destination, the weight attribute can decide which path is the best path. The higher the value the better the preference.

The following assigns a weight attribute of 40000 to all existing and future prefixes:

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Neighbor 12.1.1.2 weight 40000

R1#Clear ip bgp *
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                               0 400 300 200 i
*>                  12.1.1.2                 0             40000 200 i
*  3.0.0.0          14.1.1.4                               0 400 300 i
*>                  12.1.1.2                               40000 200 300 i
*  4.0.0.0          14.1.1.4                 0             0 400 i
*>                  12.1.1.2                               40000 200 300 400 i
*> 11.0.0.0          0.0.0.0                  0         32768 i
```

We can see that R1 has to traverse through R2 to reach all the networks within this topology.

```
R1#Traceroute 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 52 msec 16 msec 16 msec
  2 23.1.1.3 28 msec 28 msec 32 msec
  3 34.1.1.4 16 msec *  16 msec

R1#Traceroute 3.3.3.3

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 16 msec 12 msec 16 msec
  2 23.1.1.3 32 msec *  28 msec

R1#Traceroute 2.2.2.2

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 16 msec *  12 msec
```

Task 4

Configure and advertise the following Loopback interface on R3.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Loopback 1 – 33.3.3.3 /24

On R3:

```
R3(config)#Int Lo1
R3(config-if)#IP addr 33.3.3.3 255.255.255.0

R3(config)#Router bgp 300
R3(config-router)#Network 33.3.3.0 mask 255.255.255.0
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          34.1.1.4                               0 400 100 i
*>                  23.1.1.2                               0 200 100 i
*> 2.0.0.0          23.1.1.2                 0             0 200 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          34.1.1.4                 0             0 400 i
*  11.0.0.0         34.1.1.4                               0 400 100 i
*>                  23.1.1.2                               0 200 100 i
*> 33.3.3.0/24      0.0.0.0                  0         32768 i
```

On R1:

```
R1#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                               0 400 300 200 i
*>                  12.1.1.2                 0         40000 200 i
*  3.0.0.0          14.1.1.4                               0 400 300 i
*>                  12.1.1.2                           40000 200 300 i
*  4.0.0.0          14.1.1.4                 0             0 400 i
*>                  12.1.1.2                           40000 200 300 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*> 33.3.3.0/24      12.1.1.2                           40000 200 300 i
*                   14.1.1.4                               0 400 300 i
```

NOTE: The "Weight" attribute configured in Task 3 affected the existing and the future prefix/es within the topology. To verify this information:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Traceroute 33.3.3.3

Type escape sequence to abort.
Tracing the route to 33.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 12 msec 16 msec 16 msec
  2 23.1.1.3 28 msec *  28 msec
```

Task 5

Re-configure R1 based on the following policy, this policy may override the previously implemented policy:

- R1 in AS 100 should use R4 in AS 400 to reach networks 33.0.0.0 /8 and network 4.0.0.0 /8.
- R1 in AS 100 should use R2 in AS 200 to reach network 3.0.0.0 /8, and existing and future prefixes advertised in AS 200.

Let's remove the "Weight" configuration that was implemented in task 3:

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No Neigh 12.1.1.2 weight 40000

R1#Clear ip bgp *
```

To configure this task, two access-lists are configured, the first access-list identifies networks 4.0.0.0 and 33.3.3.0/8, and the second one identifies network 4.0.0.0/8, but how do we identify the existing and the future prefixes advertised in AS 200? The easiest way to accomplish this is to configure an AS-Path access-list identifying the AS number, because the AS number remains the same for existing and the future routes in AS 200:

The following access-list identifies networks 33.3.3.0 /24 and 4.0.0.0 /8:

```
R1(config)#Access-list 1 permit 33.3.3.0 0.0.0.255
R1(config)#Access-list 1 permit 4.0.0.0 0.255.255.255
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The following access-list identifies networks 3.0.0.0 /8:

```
R1(config)#Access-list 2 permit 3.0.0.0 0.255.255.255
```

The following as-path access-list identifies AS 200:

```
R1(config)#Ip as-path access-list 1 permit ^200$
```

Once the prefixes are all identified, they are referenced in a route-map and the "Weight" attribute is assigned to them based on the requirements of this task.

The following route-map (TST) assigns a weight attribute of 45000 to the networks identified in the access-list 1.

```
R1(config)#Route-map TST permit 10
R1(config-route-map)#match ip addr 1
R1(config-route-map)#set weight 45000
R1(config)#Route-map TST permit 90
```

The following route-map (TEST) assigns a weight attribute of 50000 to the networks identified in the access-list 2 and AS 200.

```
R1(config)#Route-map TEST per 10
R1(config-route-map)#match as-path 1
R1(config-route-map)#set weight 50000

R1(config)#Route-map TEST per 20
R1(config-route-map)#match ip addr 2
R1(config-route-map)#set weight 50000
R1(config)#Route-map TEST per 90

R1(config)#Router bgp 100
R1(config-router)#Neighbor 14.1.1.4 route-map TST in
R1(config-router)#Neighbor 12.1.1.2 route-map TEST in
```

The above commands assign the attributes to neighbors R2 and R4.

```
R1#Clear ip bgp *
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Sh ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0         0.0.0.0                  0         32768 i
*  2.0.0.0         14.1.1.4                               0 400 300 200 i
*>                 12.1.1.2                 0             50000 200 i
*  3.0.0.0         14.1.1.4                               0 400 300 i
*>                 12.1.1.2                               50000 200 300 i
*> 4.0.0.0         14.1.1.4                 0             45000 400 i
*                  12.1.1.2                               0 200 300 400 i
*> 11.0.0.0        0.0.0.0                  0         32768 i
*> 33.3.3.0/24     14.1.1.4                               45000 400 300 i
*                  12.1.1.2                               0 200 300 i
```

NOTE: Network 3.0.0.0/8 and Networks that originated in AS 200 are assigned a weight attribute of 50000, and networks 4.0.0.0 and 33.3.3.0 /24 are assigned a weight attribute of 45000, therefore, R1 goes through R4 to reach networks 33.3.3.0/24 and 4.0.0.0/8, and it goes through R2 for all existing and future prefixes that originated in AS 200, and network 3.0.0.0/8. Let's verify:

```
R1#Traceroute 4.4.4.4 numeric

Type escape sequence to abort.
Tracing the route to 4.4.4.4
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.4 0 msec *  0 msec

R1#Traceroute 33.3.3.3 numeric

Type escape sequence to abort.
Tracing the route to 33.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.4 4 msec 0 msec 0 msec
  2 34.1.1.3 12 msec *  12 msec

R1#Traceroute 2.2.2.2 numeric

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 32 msec *  12 msec

R1#Traceroute 3.3.3.3 numeric

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
1 12.1.1.2 16 msec 16 msec 16 msec
2 23.1.1.3 28 msec *  28 msec
```

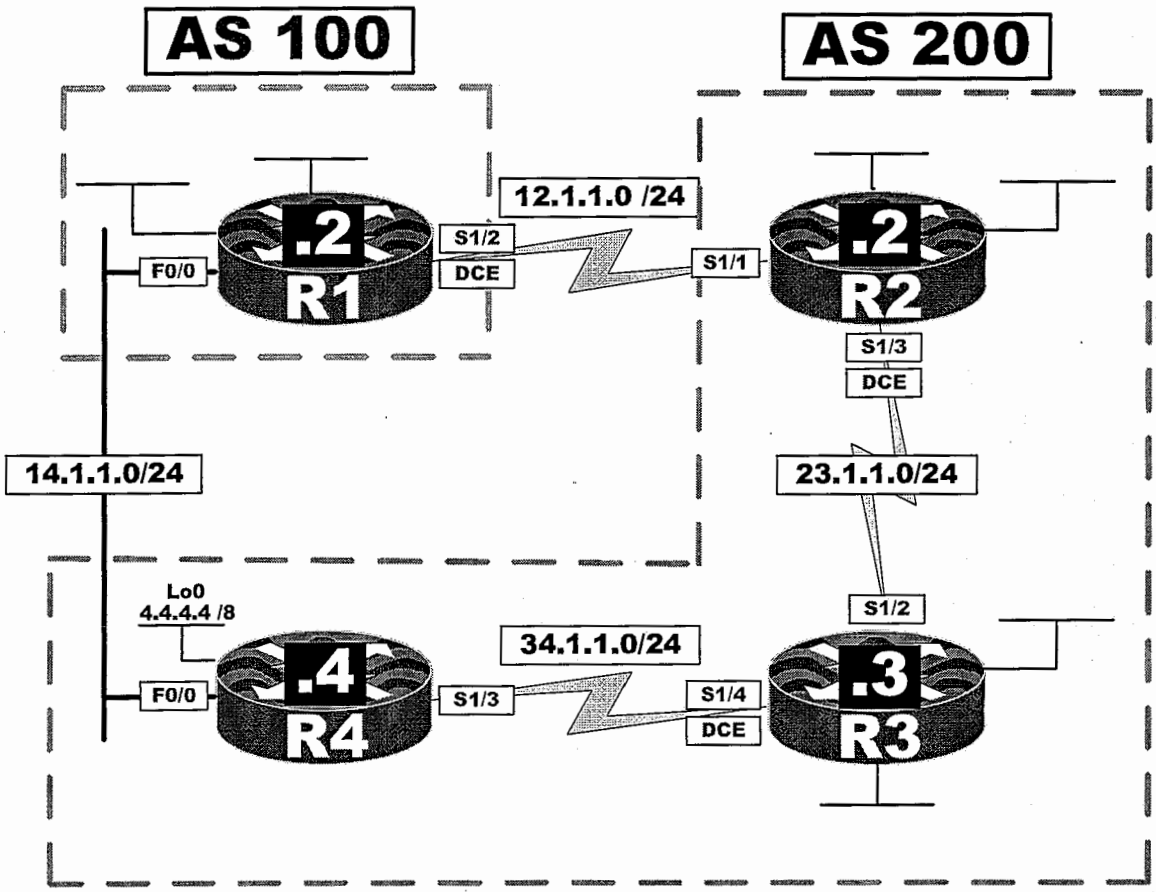
Task 6

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 13
Multi Exit Discriminator Attribute



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-13".

Task 1

Configure routers R2, R3 and R4 in AS 200, these routers should have full mesh peer session between them. Routers R2 and router R4 should have EBGP peer session to R1

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

in AS 100. BGP routers should **only** advertise their loopback interface/s in BGP. Provide NLRI for the links in As 200 using OSPF.

On R1:

```
R1(config-if)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 14.1.1.4 remote-as 200
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Netw 11.0.0.0
R1(config-router)#Netw 1.0.0.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Netw 22.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 23.1.1.3 remote-as 200
R2(config-router)#Neighbor 23.1.1.3 next-hop-self
R2(config-router)#Neighbor 34.1.1.4 remote-as 200
R2(config-router)#Neighbor 34.1.1.4 next-hop-self

R2(config)#Router ospf 1
R2(config-router)#Router-id 0.0.0.2
R2(config-router)#Netw 23.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Netw 33.0.0.0
R3(config-router)#Neighbor 34.1.1.4 remote-as 200
R3(config-router)#Neighbor 23.1.1.2 remote-as 200

R3(config)#Router ospf 1
R3(config-router)#Router-id 0.0.0.3
R3(config-router)#Netw 23.1.1.3 0.0.0.0 area 0
R3(config-router)#Netw 34.1.1.3 0.0.0.0 area 0
```

On R4:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config)#Router bgp 200
R4(config-router)#No au
R4(config-router)#netw 4.0.0.0
R4(config-router)#Neighbor 34.1.1.3 remote-as 200
R4(config-router)#Neighbor 34.1.1.3 next-hop-self
R4(config-router)#Neighbor 23.1.1.2 remote-as 200
R4(config-router)#Neighbor 23.1.1.2 next-hop-self
R4(config-router)#Neighbor 14.1.1.1 remote-as 100

R4(config)#Router ospf 1
R4(config-router)#Router-id 0.0.0.4
R4(config-router)#Netw 34.1.1.4 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1.0.0.0	0.0.0.0	0		32768	i
*	2.0.0.0	14.1.1.4			0	200 i
*>		12.1.1.2	0		0	200 i
*	3.0.0.0	14.1.1.4			0	200 i
*>		12.1.1.2			0	200 i
*	4.0.0.0	12.1.1.2			0	200 i
*>		14.1.1.4	0		0	200 i
*>	11.0.0.0	0.0.0.0	0		32768	i
*	22.0.0.0	14.1.1.4			0	200 i
*>		12.1.1.2	0		0	200 i
*	33.0.0.0	14.1.1.4			0	200 i
*>		12.1.1.2			0	200 i

On R2:

```
R2#Sh ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
* i	1.0.0.0	34.1.1.4	0	100	0	100 i
*>		12.1.1.1	0		0	100 i
*>	2.0.0.0	0.0.0.0	0		32768	i
*>i	3.0.0.0	23.1.1.3	0	100	0	i
*>i	4.0.0.0	34.1.1.4	0	100	0	i
* i	11.0.0.0	34.1.1.4	0	100	0	100 i
*>		12.1.1.1	0		0	100 i

© 2014 Narbik Kocharians, All rights reserved

https://t.me/learningnets

```
*> 22.0.0.0          0.0.0.0                  0         32768 i
*>i33.0.0.0          23.1.1.3                 0    100      0 i
```

On R3:

```
R3#Sh ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.4                 0    100      0 100 i
* i                 23.1.1.2                 0    100      0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*>i4.0.0.0          34.1.1.4                 0    100      0 i
*>i11.0.0.0         34.1.1.4                 0    100      0 100 i
* i                 23.1.1.2                 0    100      0 100 i
*>i22.0.0.0         23.1.1.2                 0    100      0 i
*> 33.0.0.0         0.0.0.0                  0         32768 i
```

On R4:

```
R4#Sh ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
* i1.0.0.0          23.1.1.2                 0    100      0 100 i
*>                  14.1.1.1                 0             0 100 i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*>i3.0.0.0          34.1.1.3                 0    100      0 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
* i11.0.0.0         23.1.1.2                 0    100      0 100 i
*>                  14.1.1.1                 0             0 100 i
*>i22.0.0.0         23.1.1.2                 0    100      0 i
*>i33.0.0.0         34.1.1.3                 0    100      0 i
```

Task 2

Configure AS 200 such that R1 in AS 100 takes R4 to reach any prefix advertised in AS 200. Manipulate MED to accomplish this task.

MED is used as a suggestion to an external AS regarding the preferred path into the AS that is advertising the metric. The reason suggestion is used here is because the AS that is receiving MED attribute could use another attribute such as "weight" that will override the MED attribute. As far as MED is concerned, the

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

lower value has better preference.

On R2:

```
R2(config)#Route-map TST permit 10
R2(config-route-map)#set metric 100

R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 route-map TST out

R2#Clear ip bgp * Out
```

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1.0.0.0	0.0.0.0	0		32768	i
*>	2.0.0.0	14.1.1.4			0	200 i
*		12.1.1.2	100		0	200 i
*>	3.0.0.0	14.1.1.4			0	200 i
*		12.1.1.2	100		0	200 i
*	4.0.0.0	12.1.1.2	100		0	200 i
*>		14.1.1.4	0		0	200 i
*>	11.0.0.0	0.0.0.0	0		32768	i
*>	22.0.0.0	14.1.1.4			0	200 i
*		12.1.1.2	100		0	200 i
*>	33.0.0.0	14.1.1.4			0	200 i
*		12.1.1.2	100		0	200 i

```
R1#Traceroute 2.2.2.2 source 1.1.1.1

Type escape sequence to abort.
Tracing the route to 2.2.2.2
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.4 0 msec 4 msec 0 msec
  2 34.1.1.3 16 msec 16 msec 16 msec
  3 23.1.1.2 20 msec *  20 msec
```

Note: R1 receives a MED value of 100 for all prefixes from R2, since R1 does not receive a MED value from R4, R1 considers the MED value coming from R4 as zero. When R1 compares the two MED values, it takes the path with a lower MED value, which is R4.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 3

Remove the configuration command from the previous task before proceeding to the next task.

On R2:

```
R2(config)#No route-map TST

R2(config)#Router bgp 200
R2(config-router)#No Neighbor 12.1.1.1 route-map TST out

R2#Clear ip bgp *
```

Task 4

Configure the appropriate router/s in AS 200 such that AS 100 goes through R4 to reach network 33.0.0.0 /8 and R2 to reach network 3.0.0.0 /8. Utilize MED to accomplish this task.

The following output shows the existing BGP table of R1 before manipulating the MED attribute:

On R1:

```
R1#Sh ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1.0.0.0	0.0.0.0	0		32768	i
*	2.0.0.0	14.1.1.4			0	200 i
*>		12.1.1.2	0		0	200 i
*	3.0.0.0	12.1.1.2			0	200 i
*>		14.1.1.4			0	200 i
*	4.0.0.0	12.1.1.2			0	200 i
*>		14.1.1.4	0		0	200 i
*>	11.0.0.0	0.0.0.0	0		32768	i
*	22.0.0.0	14.1.1.4			0	200 i
*>		12.1.1.2	0		0	200 i
*	33.0.0.0	12.1.1.2			0	200 i
*>		14.1.1.4			0	200 i

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Note: Both networks are reachable via R4

On R2:

```
R2(config)#IP Prefix-list NET3 permit 3.0.0.0/8
R2(config)#IP Prefix-list NET33 permit 33.0.0.0/8
```

Note: It is very easy to remember that Prefix-list "NET3" is referencing network 3.0.0.0/8 and Prefix-list "NET33" is referencing network 33.0.0.0/8. If possible, you should choose a meaningful name.

On R2, the following configuration sets the "metric" which is the MED value to 50 for network 3.0.0.0/8 and 100 for network 33.0.0.0/8, whereas, on R4, the MED values are reversed, meaning that the MED value for network 3.0.0.0/8 is set to 100 and for network 33.0.0.0/8 it is set to 50.

```
R2(config)#Route-map TST permit 10
R2(config-route-map)#match ip addr prefix NET3
R2(config-route-map)#set metric 50

R2(config)#Route-map TST permit 20
R2(config-route-map)#match ip addr prefix NET33
R2(config-route-map)#set metric 100

R2(config)#Route-map TST permit 90

R2(config)#Router bgp 200
R2(config-router)#Neighbor 10.1.12.1 route-map TST out

R2#Clear ip bgp *
```

On R4:

```
R4(config)#IP Prefix-list NET3 permit 3.0.0.0/8
R4(config)#IP Prefix-list NET33 permit 33.0.0.0/8

R4(config)#Route-map TST permit 10
R4(config-route-map)#match ip addr prefix NET3
R4(config-route-map)#set metric 100

R4(config)#Route-map TST permit 20
R4(config-route-map)#match ip addr prefix NET33
R4(config-route-map)#set metric 50
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#Route-map TST permit 90

R4(config)#Router bgp 200
R4(config-router)#Neighbor 10.1.14.1 route-map TST out

R4#Clear ip bgp *
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          14.1.1.4                               0 200 i
*>                  12.1.1.2                 0             0 200 i
*  3.0.0.0          14.1.1.4               100             0 200 i
*>                  12.1.1.2                50             0 200 i
*  4.0.0.0          12.1.1.2                               0 200 i
*>                  14.1.1.4                 0             0 200 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         14.1.1.4                               0 200 i
*>                  12.1.1.2                 0             0 200 i
*> 33.0.0.0         14.1.1.4                50             0 200 i
*                   12.1.1.2               100             0 200 i
```

Note: Network 3.0.0.0/8 is reachable via R2 and network 33.0.0.0/8 is reachable via R4

To test the configuration:

On R1:

```
R1#Traceroute 3.3.3.3 source 1.1.1.1

Type escape sequence to abort.
Tracing the route to 3.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 16 msec 16 msec 12 msec
  2 23.1.1.3 24 msec *  20 msec

R1#Traceroute 33.3.3.3 source 1.1.1.1

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
Tracing the route to 33.3.3.3
VRF info: (vrf in name/id, vrf out name/id)
  1 14.1.1.4 4 msec 0 msec 0 msec
  2 34.1.1.3 12 msec *  12 msec
```

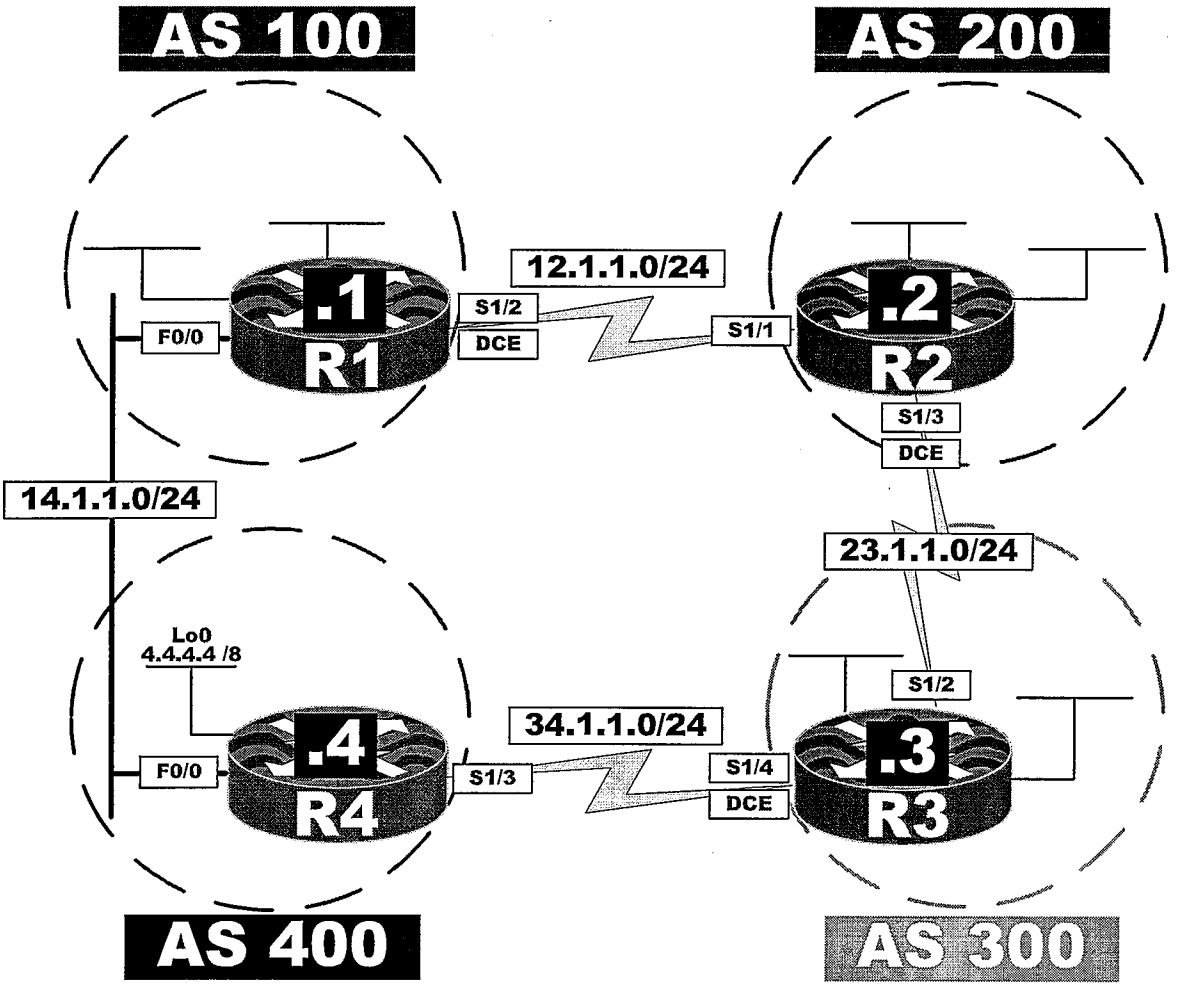
Task 5

Remove BGP routing protocol and any previous BGP related command/s from the previous tasks and reconfigure them based on the chart below. These routers should advertise their loopback interface/s in their own AS. The BGP peering should be established as follows:

R1 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R2 should establish EBGP peer sessions with R1 and R3 in AS 100 and 300 respectively.
R3 should establish EBGP peer sessions with R2 and R4 in AS 200 and 400 respectively.
R4 should establish EBGP peer sessions with R3 and R1 in AS 300 and 100 respectively.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets



On R1:

```
R1(config)#No router bgp 100

R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Netw 11.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Neighbor 14.1.1.4 remote-as 400
```

On R2:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config)#No Router bgp 200
R2(config)#No IP Prefix-list NET3
R2(config)#No IP Prefix-list NET33
R2(config)#No Route-map TST

R2(config)#Router bgp 200
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Netw 22.0.0.0
R2(config-router)#No au
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
R2(config-router)#Neighbor 23.1.1.3 remote-as 300
```

On R3:

```
R3(config)#No Router bgp 200

R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Netw 33.0.0.0
R3(config-router)#Neighbor 23.1.1.2 remote-as 200
R3(config-router)#Neighbor 34.1.1.4 remote-as 400
```

On R4:

```
R4(config)#No Router bgp 200
R4(config)#No IP Prefix-list NET3
R4(config)#No IP Prefix-list NET33
R4(config)#No Route-map TST

R4(config)#Router bgp 400
R4(config-router)#No au
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Neighbor 14.1.1.1 remote-as 100
R4(config-router)#Neighbor 34.1.1.3 remote-as 300
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*  2.0.0.0          14.1.1.4                               0 400 300 200 i
*>                  12.1.1.2                 0             0 200 i
*> 3.0.0.0          14.1.1.4                               0 400 300 i
*                   12.1.1.2                               0 200 300 i
*  4.0.0.0          12.1.1.2                               0 200 300 400 i
*>                  14.1.1.4                 0             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         14.1.1.4                               0 400 300 200 i
*>                  12.1.1.2                 0             0 200 i
*> 33.0.0.0         14.1.1.4                               0 400 300 i
*                   12.1.1.2                               0 200 300 i
```

Task 6

Configure R4 in AS 400 to pass a MED value of 100 to R1 in AS 100 and configure R2 in AS 200 to pass a MED of 120 to R1 in AS 100.

On R2:

```
R2(config)#Route-map TST permit 10
R2(config-route-map)#set metric 120

R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 route-map TST out

R2#Clear ip bgp *
```

On R4:

```
R4(config)#Route-map TST permit 10
R4(config-route-map)#set metric 100

R4(config)#Router bgp 400
R4(config-router)#Neighbor 14.1.1.1 route-map TST out

R4#Clear ip bgp * out
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1#Clear ip bgp *

R1#Show ip bgp | b Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1.0.0.0	0.0.0.0	0		32768	i
*>	2.0.0.0	12.1.1.2	120		0	200 i
*	3.0.0.0	12.1.1.2	120		0	200 300 i
*>		14.1.1.4	100		0	400 300 i
*	4.0.0.0	12.1.1.2	120		0	200 300 400 i
*>		14.1.1.4	100		0	400 i
*>	11.0.0.0	0.0.0.0	0		32768	i
*>	22.0.0.0	12.1.1.2	120		0	200 i
*	33.0.0.0	12.1.1.2	120		0	200 300 i
*>		14.1.1.4	100		0	400 300 i

Task 7

Ensure that R1 in AS 100 always takes R4 to reach all prefixes advertised in this topology. You should configure R1 to accomplish this task.

Note: R1 goes through R2 to reach network 22.0.0.0 /8 even though the MED value is lower through R4. The reason this happens is because the AS-PATH is longer through R4, and we know that the AS-Path is considered before MED.

There are two challenges in resolving this task:

The networks are being advertised by both ASes and by default the "MED" attribute is only considered among paths from the same AS.

AS-Path attribute is considered before MED.

To resolve the above challenges, we can change the default behavior of MED such that the local router looks at the MED attribute coming from two ASes, in this case AS 200 and AS 400; then, we MUST instruct the local router (R1) to ignore the AS-Path attribute when performing best path determination.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#Bgp always-compare-med
R1(config-router)#Bgp bestpath as-path ignore
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Clear ip bgp *
```

Note: The tab key will not work when entering the "bgp bestpath as-path ignore" command. This command is a hidden one.

The MED as stated in RFC 1771, is an optional non-transitive attribute. The value of this attribute may be used by BGP speakers for best path selection process to discriminate among multiple exit points to a neighboring AS. The lower the value the better the path. The MED value comparison is done only among paths from the same AS. The "bgp always-compare-med" command is used to change this behavior by enforcing MED comparison between all paths, regardless of the AS from which the paths are received.

Note: The second command "bgp bestpath as-path ignore" is also needed as part of the solution in accomplishing this task, this command is needed because as-path and the origin attributes are looked at before the MED attribute. Since the origin attribute is identical for all routes and BGP routing protocol is told to ignore the as-path attribute in best path calculation, BGP will look at the MED value next.

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          12.1.1.2               120             0 200 i
*>                  14.1.1.4               100             0 400 300 200 i
*  3.0.0.0          12.1.1.2               120             0 200 300 i
*>                  14.1.1.4               100             0 400 300 i
*  4.0.0.0          12.1.1.2               120             0 200 300 400 i
*>                  14.1.1.4               100             0 400 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
*  22.0.0.0         12.1.1.2               120             0 200 i
*>                  14.1.1.4               100             0 400 300 200 i
*  33.0.0.0         12.1.1.2               120             0 200 300 i
*>                  14.1.1.4               100             0 400 300 i
```

Note: R1 takes R4 to reach all the networks in this topology

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 8

Remove the BGP configuration performed in Task 6. Configure R4 to pass a MED value of 100 to R1; R2 in AS 200 should **NOT** pass any MED values to R1.

Ensure that R1 in AS 100 takes R4 to reach any prefix advertised in this topology except the ones originated by R2. **Do not** use any global config command/s as part of the solution to this task.

To remove the configuration performed in Task 6 and verify the configuration:

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No Neighbor 12.1.1.1 route-map TST out

R2(config)#No Route-map TST

R2#Clear ip bgp *
```

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#No Neighbor 14.1.1.1 route-map TST out

R4(config)#No Route-map TST

R4#Clear ip bgp *
```

Let's configure R4 to pass a MED value of 100 to R1:

On R4:

```
R4(config)#Route-map TST per 10
R4(config-route-map)#set metric 100

R4(config)#Router bgp 400
R4(config-router)#Neigh 14.1.1.1 route-map TST out

R4#Clear ip bgp *
```

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0         0.0.0.0                  0         32768 i
*> 2.0.0.0         12.1.1.2                 0             0 200 i
*  3.0.0.0         14.1.1.4               100             0 400 300 i
*>                 12.1.1.2                               0 200 300 i
*> 4.0.0.0         12.1.1.2                               0 200 300 400 i
*                  14.1.1.4               100             0 400 i
*> 11.0.0.0        0.0.0.0                  0         32768 i
*> 22.0.0.0        12.1.1.2                 0             0 200 i
*  33.0.0.0        14.1.1.4               100             0 400 300 i
*>                 12.1.1.2                               0 200 300 i
```

Since R2 is not giving a MED value to R1; R1 consider the MED value coiming from R2 as zero, therefore, R1 takes the path through R2 (The path with the lowest MED value). This behavior can be changed by instructing R1 to set the path with the missing MED value to the worst possible value.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#bgp bestpath med missing-as-worst
```

To verify the configuration:

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0         0.0.0.0                  0         32768 i
*> 2.0.0.0         12.1.1.2                 0             0 200 i
*> 3.0.0.0         14.1.1.4               100             0 400 300 i
*                  12.1.1.2            4294967295           0 200 300 i
*  4.0.0.0         12.1.1.2            4294967295           0 200 300 400 i
*>                 14.1.1.4               100             0 400 i
*> 11.0.0.0        0.0.0.0                  0         32768 i
*> 22.0.0.0        12.1.1.2                 0             0 200 i
*> 33.0.0.0        14.1.1.4               100             0 400 300 i
*                  12.1.1.2            4294967295           0 200 300 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

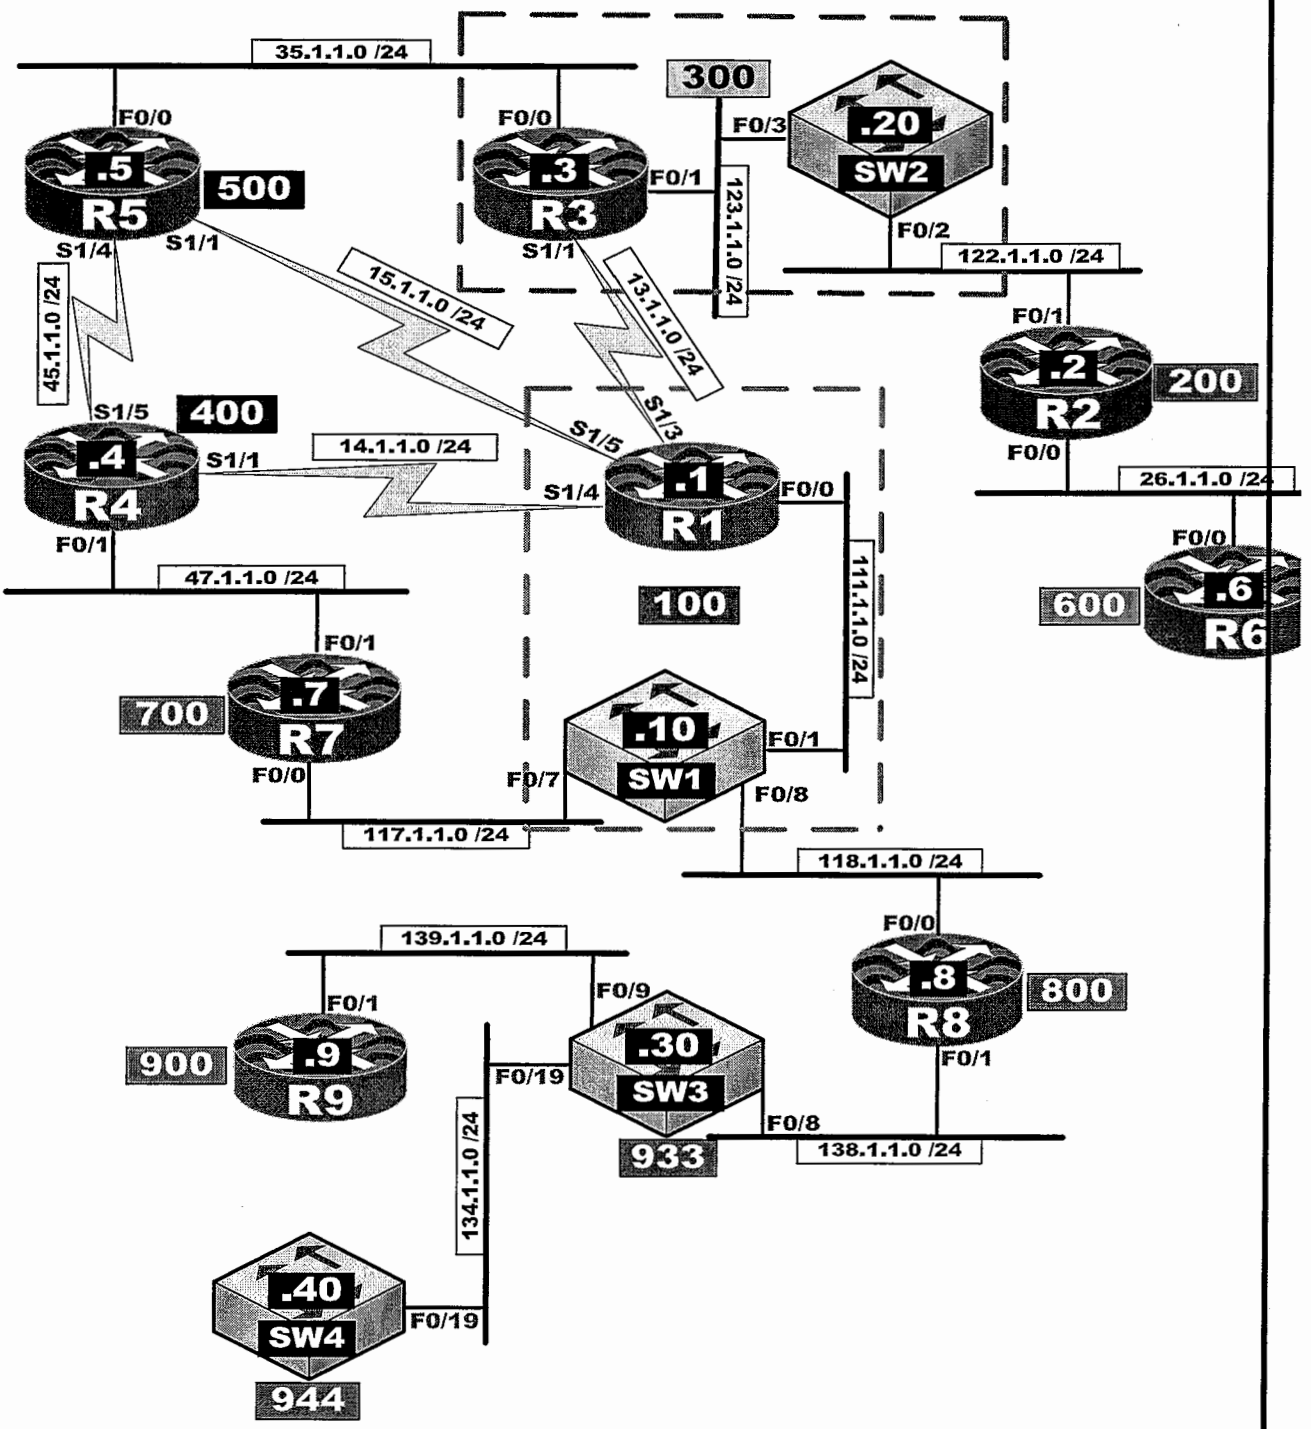
Task 9

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 14 – BGP Filtering



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-14".

The following loopback interfaces are pre-configured on the router:

Router	Loopback 0	Loopback 200
R1	1.1.1.1 /8	200.1.1.1 /32
R2	2.2.2.2 /8	200.1.1.2 /32
R3	3.3.3.3 /8	200.1.1.3 /32
R4	4.4.4.4 /8	200.1.1.4 /32
R5	5.5.5.5 /8	200.1.1.5 /32
R6	6.6.6.6 /8	200.1.1.6 /32
R7	7.7.7.7 /8	200.1.1.7 /32
R8	8.8.8.8 /8	200.1.1.8 /32
R9	9.9.9.9 /8	200.1.1.9 /32
SW1	10.1.1.10 /8	200.1.1.10 /32
SW2	20.1.1.20 /8	200.1.1.20 /32
SW3	30.1.1.30 /8	200.1.1.30 /32
SW4	40.1.1.40 /8	200.1.1.40 /32

Task 1

Configure RIPv2 on all Routers and Switches in the above topology, these routers should advertise the links and their Loopback 0 interface in RIPv2.

On R1:

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 13.0.0.0
R1(config-router)#Network 14.0.0.0
R1(config-router)#Network 15.0.0.0
R1(config-router)#Network 111.0.0.0
R1(config-router)#Network 1.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 122.0.0.0
R2(config-router)#Network 26.0.0.0
R2(config-router)#Network 2.0.0.0
```

On R3:

```
R3(config)#Router rip
R3(config-router)#No au
R3(config-router)#Ver 2
R3(config-router)#Network 35.0.0.0
R3(config-router)#Network 123.0.0.0
R3(config-router)#Network 13.0.0.0
R3(config-router)#Network 3.0.0.0
```

On R4:

```
R4(config)#Router rip
R4(config-router)#No au
R4(config-router)#Ver 2
R4(config-router)#Network 45.0.0.0
R4(config-router)#Network 14.0.0.0
R4(config-router)#Network 47.0.0.0
R4(config-router)#Network 4.0.0.0
```

On R5:

```
R5(config)#Router rip
R5(config-router)#No au
R5(config-router)#Ver 2
R5(config-router)#Network 35.0.0.0
R5(config-router)#Network 15.0.0.0
R5(config-router)#Network 45.0.0.0
R5(config-router)#Network 5.0.0.0
```

On R6:

```
R6(config)#Router rip
R6(config-router)#No au
R6(config-router)#Ver 2
R6(config-router)#Network 26.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R6(config-router)#Network 6.0.0.0
```

On R7:

```
R7(config)#Router rip
R7(config-router)#No au
R7(config-router)#Ver 2
R7(config-router)#Network 47.0.0.0
R7(config-router)#Network 117.0.0.0
R7(config-router)#Network 7.0.0.0
```

On R8:

```
R8(config)#Router rip
R8(config-router)#No au
R8(config-router)#Ver 2
R8(config-router)#Network 118.0.0.0
R8(config-router)#Network 138.1.0.0
R8(config-router)#Network 8.0.0.0
```

On R9:

```
R9(config)#Router rip
R9(config-router)#No au
R9(config-router)#Ver 2
R9(config-router)#Network 139.1.0.0
R9(config-router)#Network 9.0.0.0
```

On SW1:

```
SW1(config)#Router rip
SW1(config-router)#No au
SW1(config-router)#Ver 2
SW1(config-router)#Network 117.0.0.0
SW1(config-router)#Network 118.0.0.0
SW1(config-router)#Network 111.0.0.0
SW1(config-router)#Network 10.0.0.0
```

On SW2:

```
SW2(config)#Router rip
SW2(config-router)#No au
SW2(config-router)#Ver 2
SW2(config-router)#Network 123.0.0.0
SW2(config-router)#Network 122.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW2(config-router)#Network 20.0.0.0
```

On SW3:

```
SW3(config)#Router rip
SW3(config-router)#No au
SW3(config-router)#Ver 2
SW3(config-router)#Network 139.1.0.0
SW3(config-router)#Network 134.1.0.0
SW3(config-router)#Network 138.1.0.0
SW3(config-router)#Network 30.0.0.0
```

On SW4:

```
SW4(config)#Router rip
SW4(config-router)#No au
SW4(config-router)#Ver 2
SW4(config-router)#Network 134.1.0.0
SW4(config-router)#Network 40.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

R     2.0.0.0/8 [120/3] via 13.1.1.3, 00:00:03, Serial1/3
R     3.0.0.0/8 [120/1] via 13.1.1.3, 00:00:03, Serial1/3
R     4.0.0.0/8 [120/1] via 14.1.1.4, 00:00:27, Serial1/4
R     5.0.0.0/8 [120/1] via 15.1.1.5, 00:00:25, Serial1/5
R     6.0.0.0/8 [120/4] via 13.1.1.3, 00:00:03, Serial1/3
R     7.0.0.0/8 [120/2] via 111.1.1.10, 00:00:25, FastEthernet0/0
                [120/2] via 14.1.1.4, 00:00:27, Serial1/4
R     8.0.0.0/8 [120/2] via 111.1.1.10, 00:00:25, FastEthernet0/0
R     9.0.0.0/8 [120/4] via 111.1.1.10, 00:00:25, FastEthernet0/0
R     10.0.0.0/8 [120/1] via 111.1.1.10, 00:00:25, FastEthernet0/0
R     20.0.0.0/8 [120/2] via 13.1.1.3, 00:00:03, Serial1/3
      26.0.0.0/24 is subnetted, 1 subnets
R        26.1.1.0 [120/3] via 13.1.1.3, 00:00:03, Serial1/3
R     30.0.0.0/8 [120/3] via 111.1.1.10, 00:00:25, FastEthernet0/0
      35.0.0.0/24 is subnetted, 1 subnets
R        35.1.1.0 [120/1] via 15.1.1.5, 00:00:25, Serial1/5
                  [120/1] via 13.1.1.3, 00:00:03, Serial1/3
R     40.0.0.0/8 [120/4] via 111.1.1.10, 00:00:25, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
      45.0.0.0/24 is subnetted, 1 subnets

R        45.1.1.0 [120/1] via 15.1.1.5, 00:00:25, Serial1/5
                  [120/1] via 14.1.1.4, 00:00:27, Serial1/4
      47.0.0.0/24 is subnetted, 1 subnets
R        47.1.1.0 [120/1] via 14.1.1.4, 00:00:27, Serial1/4
      117.0.0.0/24 is subnetted, 1 subnets
R        117.1.1.0 [120/1] via 111.1.1.10, 00:00:25, FastEthernet0/0
      118.0.0.0/24 is subnetted, 1 subnets
R        118.1.1.0 [120/1] via 111.1.1.10, 00:00:25, FastEthernet0/0
      122.0.0.0/24 is subnetted, 1 subnets
R        122.1.1.0 [120/2] via 13.1.1.3, 00:00:03, Serial1/3
      123.0.0.0/24 is subnetted, 1 subnets
R        123.1.1.0 [120/1] via 13.1.1.3, 00:00:03, Serial1/3
      134.1.0.0/24 is subnetted, 1 subnets
R        134.1.1.0 [120/3] via 111.1.1.10, 00:00:25, FastEthernet0/0
      138.1.0.0/24 is subnetted, 1 subnets
R        138.1.1.0 [120/2] via 111.1.1.10, 00:00:25, FastEthernet0/0
      139.1.0.0/24 is subnetted, 1 subnets
R        139.1.1.0 [120/3] via 111.1.1.10, 00:00:25, FastEthernet0/0
```

On R6:

```
R6#Show ip route rip | B Gate
Gateway of last resort is not set

R     1.0.0.0/8 [120/4] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     2.0.0.0/8 [120/1] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     3.0.0.0/8 [120/3] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     4.0.0.0/8 [120/5] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     5.0.0.0/8 [120/4] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     7.0.0.0/8 [120/6] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     8.0.0.0/8 [120/6] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     9.0.0.0/8 [120/8] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     10.0.0.0/8 [120/5] via 26.1.1.2, 00:00:13, FastEthernet0/0
      13.0.0.0/24 is subnetted, 1 subnets
R        13.1.1.0 [120/3] via 26.1.1.2, 00:00:13, FastEthernet0/0
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/4] via 26.1.1.2, 00:00:13, FastEthernet0/0
      15.0.0.0/24 is subnetted, 1 subnets
R        15.1.1.0 [120/4] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     20.0.0.0/8 [120/2] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     30.0.0.0/8 [120/7] via 26.1.1.2, 00:00:13, FastEthernet0/0
      35.0.0.0/24 is subnetted, 1 subnets
R        35.1.1.0 [120/3] via 26.1.1.2, 00:00:13, FastEthernet0/0
R     40.0.0.0/8 [120/8] via 26.1.1.2, 00:00:13, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      45.0.0.0/24 is subnetted, 1 subnets

R        45.1.1.0 [120/4] via 26.1.1.2, 00:00:13, FastEthernet0/0
      47.0.0.0/24 is subnetted, 1 subnets
R        47.1.1.0 [120/5] via 26.1.1.2, 00:00:13, FastEthernet0/0
      111.0.0.0/24 is subnetted, 1 subnets
R        111.1.1.0 [120/4] via 26.1.1.2, 00:00:13, FastEthernet0/0
      117.0.0.0/24 is subnetted, 1 subnets
R        117.1.1.0 [120/5] via 26.1.1.2, 00:00:13, FastEthernet0/0
      118.0.0.0/24 is subnetted, 1 subnets
R        118.1.1.0 [120/5] via 26.1.1.2, 00:00:13, FastEthernet0/0
      122.0.0.0/24 is subnetted, 1 subnets
R        122.1.1.0 [120/1] via 26.1.1.2, 00:00:13, FastEthernet0/0
      123.0.0.0/24 is subnetted, 1 subnets
R        123.1.1.0 [120/2] via 26.1.1.2, 00:00:13, FastEthernet0/0
      134.1.0.0/24 is subnetted, 1 subnets
R        134.1.1.0 [120/7] via 26.1.1.2, 00:00:13, FastEthernet0/0
      138.1.0.0/24 is subnetted, 1 subnets
R        138.1.1.0 [120/6] via 26.1.1.2, 00:00:13, FastEthernet0/0
      139.1.0.0/24 is subnetted, 1 subnets
R        139.1.1.0 [120/7] via 26.1.1.2, 00:00:13, FastEthernet0/0
```

On SW4:

```
SW4#Show ip route rip | Inc R

R    1.0.0.0/8 [120/4] via 134.1.1.30, 00:00:27, FastEthernet0/19
R       35.1.1.0 [120/5] via 134.1.1.30, 00:00:27, FastEthernet0/19
R       118.1.1.0 [120/2] via 134.1.1.30, 00:00:27, FastEthernet0/19
R    2.0.0.0/8 [120/7] via 134.1.1.30, 00:00:27, FastEthernet0/19
R       117.1.1.0 [120/3] via 134.1.1.30, 00:00:27, FastEthernet0/19
R       138.1.1.0 [120/1] via 134.1.1.30, 00:00:27, FastEthernet0/19
R    3.0.0.0/8 [120/5] via 134.1.1.30, 00:00:27, FastEthernet0/19
R       139.1.1.0 [120/1] via 134.1.1.30, 00:00:27, FastEthernet0/19
R    4.0.0.0/8 [120/5] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    20.0.0.0/8 [120/6] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    5.0.0.0/8 [120/5] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    6.0.0.0/8 [120/8] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    7.0.0.0/8 [120/4] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    8.0.0.0/8 [120/2] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    9.0.0.0/8 [120/2] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       111.1.1.0 [120/3] via 134.1.1.30, 00:00:28, FastEthernet0/19
R    10.0.0.0/8 [120/3] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       26.1.1.0 [120/7] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       123.1.1.0 [120/5] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       122.1.1.0 [120/6] via 134.1.1.30, 00:00:28, FastEthernet0/19
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R       13.1.1.0 [120/4] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       47.1.1.0/24 [120/4] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       14.1.1.0 [120/4] via 134.1.1.30, 00:00:28, FastEthernet0/19
R       45.1.1.0/24 [120/5] via 134.1.1.30, 00:00:01, FastEthernet0/19
R    30.0.0.0/8 [120/1] via 134.1.1.30, 00:00:01, FastEthernet0/19
R       15.1.1.0 [120/4] via 134.1.1.30, 00:00:02, FastEthernet0/19
```

On R9:

```
R9#Show ip route rip | B Gate
Gateway of last resort is not set

R     1.0.0.0/8 [120/4] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     2.0.0.0/8 [120/7] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     3.0.0.0/8 [120/5] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     4.0.0.0/8 [120/5] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     5.0.0.0/8 [120/5] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     6.0.0.0/8 [120/8] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     7.0.0.0/8 [120/4] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     8.0.0.0/8 [120/2] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     10.0.0.0/8 [120/3] via 139.1.1.30, 00:00:21, FastEthernet0/1
      13.0.0.0/24 is subnetted, 1 subnets
R        13.1.1.0 [120/4] via 139.1.1.30, 00:00:21, FastEthernet0/1
      14.0.0.0/24 is subnetted, 1 subnets
R        14.1.1.0 [120/4] via 139.1.1.30, 00:00:21, FastEthernet0/1
      15.0.0.0/24 is subnetted, 1 subnets
R        15.1.1.0 [120/4] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     20.0.0.0/8 [120/6] via 139.1.1.30, 00:00:21, FastEthernet0/1
      26.0.0.0/24 is subnetted, 1 subnets
R        26.1.1.0 [120/7] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     30.0.0.0/8 [120/1] via 139.1.1.30, 00:00:21, FastEthernet0/1
      35.0.0.0/24 is subnetted, 1 subnets
R        35.1.1.0 [120/5] via 139.1.1.30, 00:00:21, FastEthernet0/1
R     40.0.0.0/8 [120/2] via 139.1.1.30, 00:00:21, FastEthernet0/1
      45.0.0.0/24 is subnetted, 1 subnets
R        45.1.1.0 [120/5] via 139.1.1.30, 00:00:21, FastEthernet0/1
      47.0.0.0/24 is subnetted, 1 subnets
R        47.1.1.0 [120/4] via 139.1.1.30, 00:00:21, FastEthernet0/1
      111.0.0.0/24 is subnetted, 1 subnets
R        111.1.1.0 [120/3] via 139.1.1.30, 00:00:21, FastEthernet0/1
      117.0.0.0/24 is subnetted, 1 subnets
R        117.1.1.0 [120/3] via 139.1.1.30, 00:00:21, FastEthernet0/1
      118.0.0.0/24 is subnetted, 1 subnets
R        118.1.1.0 [120/2] via 139.1.1.30, 00:00:21, FastEthernet0/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      122.0.0.0/24 is subnetted, 1 subnets

R        122.1.1.0 [120/6] via 139.1.1.30, 00:00:21, FastEthernet0/1
      123.0.0.0/24 is subnetted, 1 subnets
R        123.1.1.0 [120/5] via 139.1.1.30, 00:00:21, FastEthernet0/1
      134.1.0.0/24 is subnetted, 1 subnets
R        134.1.1.0 [120/1] via 139.1.1.30, 00:00:21, FastEthernet0/1
      138.1.0.0/24 is subnetted, 1 subnets
R        138.1.1.0 [120/1] via 139.1.1.30, 00:00:21, FastEthernet0/1
```

Task 2

Establish BGP session on the routers based on the diagram. The BGP sessions MUST be established based on the Loopback0 interface of the routers. These routers should advertise their Lo200 in BGP.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No Au
R1(config-router)#Neighbor 10.1.1.10 remote-as 100
R1(config-router)#Neighbor 10.1.1.10 update-source Lo0
R1(config-router)#Neighbor 10.1.1.10 next-hop-self

R1(config-router)#Neighbor 3.3.3.3 remote-as 300
R1(config-router)#Neighbor 3.3.3.3 update-source Lo0
R1(config-router)#Neighbor 3.3.3.3 ebgp-multihop 2

R1(config-router)#Neighbor 5.5.5.5 remote-as 500
R1(config-router)#Neighbor 5.5.5.5 update-source Lo0
R1(config-router)#Neighbor 5.5.5.5 ebgp-multihop 2

R1(config-router)#Neighbor 4.4.4.4 remote-as 400
R1(config-router)#Neighbor 4.4.4.4 update-source Lo0
R1(config-router)#Neighbor 4.4.4.4 ebgp-multihop 2
R1(config-router)#Network 200.1.1.1 mask 255.255.255.255
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No Au
R2(config-router)#Neighbor 6.6.6.6 remote-as 600
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Neighbor 6.6.6.6 update-source Lo0

R2(config-router)#Neighbor 6.6.6.6 ebgp-multihop 2

R2(config-router)#Neighbor 20.1.1.20 remote-as 300
R2(config-router)#Neighbor 20.1.1.20 update-source Lo0
R2(config-router)#Neighbor 20.1.1.20 ebgp-multihop 2
R2(config-router)#Network 200.1.1.2 mask 255.255.255.255
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No Au
R3(config-router)#Neighbor 20.1.1.20 remote-as 300
R3(config-router)#Neighbor 20.1.1.20 update-source Lo0
R3(config-router)#Neighbor 20.1.1.20 Next-hop-self

R3(config-router)#Neighbor 1.1.1.1 remote-as 100
R3(config-router)#Neighbor 1.1.1.1 update-source Lo0
R3(config-router)#Neighbor 1.1.1.1 ebgp-multihop 2

R3(config-router)#Neighbor 5.5.5.5 remote-as 500
R3(config-router)#Neighbor 5.5.5.5 update-source Lo0
R3(config-router)#Neighbor 5.5.5.5 ebgp-multihop 2
R3(config-router)#Network 200.1.1.3 mask 255.255.255.255
```

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#No Au
R4(config-router)#Neighbor 1.1.1.1 remote-as 100
R4(config-router)#Neighbor 1.1.1.1 update-source Lo0
R4(config-router)#Neighbor 1.1.1.1 ebgp-multihop 2

R4(config-router)#Neighbor 5.5.5.5 remote-as 500
R4(config-router)#Neighbor 5.5.5.5 update-source Lo0
R4(config-router)#Neighbor 5.5.5.5 ebgp-multihop 2

R4(config-router)#Neighbor 7.7.7.7 remote-as 700
R4(config-router)#Neighbor 7.7.7.7 update-source Lo0
R4(config-router)#Neighbor 7.7.7.7 ebgp-multihop 2
R4(config-router)#Network 200.1.1.4 mask 255.255.255.255
```

On R5:

```
R5(config)#Router bgp 500
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config-router)#No Au

R5(config-router)#Neighbor 4.4.4.4 remote-as 400
R5(config-router)#Neighbor 4.4.4.4 update-source Lo0
R5(config-router)#Neighbor 4.4.4.4 ebgp-multihop 2

R5(config-router)#Neighbor 3.3.3.3 remote-as 300
R5(config-router)#Neighbor 3.3.3.3 update-source Lo0
R5(config-router)#Neighbor 3.3.3.3 ebgp-multihop 2

R5(config-router)#Neighbor 1.1.1.1 remote-as 100
R5(config-router)#Neighbor 1.1.1.1 update-source Lo0
R5(config-router)#Neighbor 1.1.1.1 ebgp-multihop 2
R5(config-router)#Network 200.1.1.5 mask 255.255.255.255
```

On R6:

```
R6(config)#Router bgp 600
R6(config-router)#No Au
R6(config-router)#Neighbor 2.2.2.2 remote-as 200
R6(config-router)#Neighbor 2.2.2.2 update-source Lo0
R6(config-router)#Neighbor 2.2.2.2 ebgp-multihop 2
R6(config-router)#Network 200.1.1.6 mask 255.255.255.255
```

On R7:

```
R7(config)#Router bgp 700
R7(config-router)#No Au
R7(config-router)#Neighbor 10.1.1.10 remote-as 100
R7(config-router)#Neighbor 10.1.1.10 update-source Lo0
R7(config-router)#Neighbor 10.1.1.10 ebgp-multihop 2
R7(config-router)#Neighbor 4.4.4.4 remote-as 400
R7(config-router)#Neighbor 4.4.4.4 update-source Lo0
R7(config-router)#Neighbor 4.4.4.4 ebgp-multihop 2
R7(config-router)#Network 200.1.1.7 mask 255.255.255.255
```

On R8:

```
R8(config)#Router bgp 800
R8(config-router)#No Au
R8(config-router)#Neighbor 10.1.1.10 remote-as 100
R8(config-router)#Neighbor 10.1.1.10 update-source Lo0
R8(config-router)#Neighbor 10.1.1.10 ebgp-multihop 2
R8(config-router)#Neighbor 30.1.1.30 remote-as 933
R8(config-router)#Neighbor 30.1.1.30 update-source Lo0
R8(config-router)#Neighbor 30.1.1.30 ebgp-multihop 2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R8(config-router)#Network 200.1.1.8 mask 255.255.255.255
```

On R9:

```
R9(config)#Router bgp 900
R9(config-router)#No Au
R9(config-router)#Neighbor 30.1.1.30 remote-as 933
R9(config-router)#Neighbor 30.1.1.30 update-source Lo0
R9(config-router)#Neighbor 30.1.1.30 ebgp-multihop 2
R9(config-router)#Network 200.1.1.9 mask 255.255.255.255
```

On SW1:

```
SW1(config)#Router bgp 100
SW1(config-router)#No Au
SW1(config-router)#Neighbor 1.1.1.1 remote-as 100
SW1(config-router)#Neighbor 1.1.1.1 update-source Lo0
SW1(config-router)#Neighbor 1.1.1.1 Next-hop-self

SW1(config-router)#Neighbor 7.7.7.7 remote-as 700
SW1(config-router)#Neighbor 7.7.7.7 update-source Lo0
SW1(config-router)#Neighbor 7.7.7.7 ebgp-multihop 2

SW1(config-router)#Neighbor 8.8.8.8 remote-as 800
SW1(config-router)#Neighbor 8.8.8.8 Update-source Lo0
SW1(config-router)#Neighbor 8.8.8.8 ebgp-multihop 2
SW1(config-router)#Network 200.1.1.10 mask 255.255.255.255
```

On SW2:

```
SW2(config)#Router bgp 300
SW2(config-router)#No Au
SW2(config-router)#Neighbor 2.2.2.2 remote-as 200
SW2(config-router)#Neighbor 2.2.2.2 Update-source Lo0
SW2(config-router)#Neighbor 2.2.2.2 ebgp-multihop 2

SW2(config-router)#Neighbor 3.3.3.3 remote-as 300
SW2(config-router)#Neighbor 3.3.3.3 Update-source Lo0
SW2(config-router)#Neighbor 3.3.3.3 Next-hop-self
SW2(config-router)#Network 200.1.1.20 mask 255.255.255.255
```

On SW3:

```
SW3(config)#Router bgp 933
SW3(config-router)#No Au
SW3(config-router)#Neighbor 40.1.1.40 remote-as 944
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW3(config-router)#Neighbor 40.1.1.40 Update-source Lo0

SW3(config-router)#Neighbor 40.1.1.40 ebgp-multihop 2

SW3(config-router)#Neighbor 8.8.8.8 remote-as 800
SW3(config-router)#Neighbor 8.8.8.8 Update-source Lo0
SW3(config-router)#Neighbor 8.8.8.8 ebgp-multihop 2

SW3(config-router)#Neighbor 9.9.9.9 remote-as 900
SW3(config-router)#Neighbor 9.9.9.9 Update-source Lo0
SW3(config-router)#Neighbor 9.9.9.9 ebgp-multihop 2
SW3(config-router)#Network 200.1.1.30 mask 255.255.255.255
```

On SW4:

```
SW4(config)#Router bgp 944
SW4(config-router)#No Au
SW4(config-router)#Neighbor 30.1.1.30 remote-as 933
SW4(config-router)#Neighbor 30.1.1.30 update-source Lo0
SW4(config-router)#Neighbor 30.1.1.30 ebgp-multihop 2
SW4(config-router)#Network 200.1.1.40 mask 255.255.255.255
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     0.0.0.0                  0         32768 i
*  200.1.1.2/32     4.4.4.4                                0 400 500 300 200 i
*                   5.5.5.5                                0 500 300 200 i
*>                  3.3.3.3                                0 300 200 i
*  200.1.1.3/32     5.5.5.5                                0 500 300 i
*>                  3.3.3.3                  0             0 300 i
*  200.1.1.4/32     5.5.5.5                                0 500 400 i
*>                  4.4.4.4                  0             0 400 i
*  200.1.1.5/32     3.3.3.3                                0 300 500 i
*                   4.4.4.4                                0 400 500 i
*>                  5.5.5.5                  0             0 500 i
*  200.1.1.6/32     4.4.4.4                                0 400 500 300 200 600 i
*                   5.5.5.5                                0 500 300 200 600 i
*>                  3.3.3.3                                0 300 200 600 i
*>i200.1.1.7/32     10.1.1.10                0    100       0 700 i
*                   5.5.5.5                                0 500 400 700 i
*                   4.4.4.4                                0 400 700 i
*>i200.1.1.8/32     10.1.1.10                0    100       0 800 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*>i200.1.1.9/32     10.1.1.10                0    100      0 800 933 900 i
*>i200.1.1.10/32    10.1.1.10                0    100      0 i
*  200.1.1.20/32    4.4.4.4                                0 400 500 300 i
*                   5.5.5.5                                0 500 300 i
*>                  3.3.3.3                                0 300 i
*>i200.1.1.30/32    10.1.1.10                0    100      0 800 933 i
*>i200.1.1.40/32    10.1.1.10                0    100      0 800 933 944 i
```

On R6:

```
R6#Show ip bgp | B Network
```

```
   Network          Next Hop      Metric LocPrf Weight Path
*> 200.1.1.1/32     2.2.2.2                          0 200 300 100 i
*> 200.1.1.2/32     2.2.2.2                0         0 200 i
*> 200.1.1.3/32     2.2.2.2                          0 200 300 i
*> 200.1.1.4/32     2.2.2.2                          0 200 300 100 400 i
*> 200.1.1.5/32     2.2.2.2                          0 200 300 500 i
*> 200.1.1.6/32     0.0.0.0                0     32768 i
*> 200.1.1.7/32     2.2.2.2                          0 200 300 100 700 i
*> 200.1.1.8/32     2.2.2.2                          0 200 300 100 800 i
*> 200.1.1.9/32     2.2.2.2                          0 200 300 100 800 933 900 i
*> 200.1.1.10/32    2.2.2.2                          0 200 300 100 i
*> 200.1.1.20/32    2.2.2.2                          0 200 300 i
*> 200.1.1.30/32    2.2.2.2                          0 200 300 100 800 933 i
*> 200.1.1.40/32    2.2.2.2                          0 200 300 100 800 933 944 i
```

On SW4:

```
SW4#Show ip bgp | B Network
```

```
   Network          Next Hop      Metric LocPrf Weight Path
*> 200.1.1.1/32     30.1.1.30                        0 933 800 100 i
*> 200.1.1.2/32     30.1.1.30                        0 933 800 100 300 200 i
*> 200.1.1.3/32     30.1.1.30                        0 933 800 100 300 i
*> 200.1.1.4/32     30.1.1.30                        0 933 800 100 400 i
*> 200.1.1.5/32     30.1.1.30                        0 933 800 100 500 i
*> 200.1.1.6/32     30.1.1.30                        0 933 800 100 300 200 600 i
*> 200.1.1.7/32     30.1.1.30                        0 933 800 100 700 i
*> 200.1.1.8/32     30.1.1.30                        0 933 800 i
*> 200.1.1.9/32     30.1.1.30                        0 933 900 i
*> 200.1.1.10/32    30.1.1.30                        0 933 800 100 i
*> 200.1.1.20/32    30.1.1.30                        0 933 800 100 300 i
*> 200.1.1.30/32    30.1.1.30              0         0 933 i
*> 200.1.1.40/32    0.0.0.0                0     32768 i
```

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R9#Show ip bgp | B Network

   Network          Next Hop      Metric LocPrf Weight Path
*> 200.1.1.1/32     30.1.1.30                      0 933 800 100 i
*> 200.1.1.2/32     30.1.1.30                      0 933 800 100 300 200 i
*> 200.1.1.3/32     30.1.1.30                      0 933 800 100 300 i
*> 200.1.1.4/32     30.1.1.30                      0 933 800 100 400 i
*> 200.1.1.5/32     30.1.1.30                      0 933 800 100 500 i
*> 200.1.1.6/32     30.1.1.30                      0 933 800 100 300 200 600 i
*> 200.1.1.7/32     30.1.1.30                      0 933 800 100 700 i
*> 200.1.1.8/32     30.1.1.30                      0 933 800 i
*> 200.1.1.9/32     0.0.0.0                  0    32768 i
*> 200.1.1.10/32    30.1.1.30                      0 933 800 100 i
*> 200.1.1.20/32    30.1.1.30                      0 933 800 100 300 i
*> 200.1.1.30/32    30.1.1.30                0     0 933 i
*> 200.1.1.40/32    30.1.1.30                      0 933 944 i
```

Task 3

Configure R2 such that ASes 300 and 600 do not use AS 200 as a transit AS.

If R2 advertises the prefixes it receives from one AS to another, then, these ASes can go through AS 200 to reach the prefixes advertised by the other AS, which means that they can use AS 200 as a transit AS.

To avoid this from happening, AS 200 should ONLY advertise the prefixes that originated in its own AS, these are prefixes with an empty AS path.

Let's check R6's BGP table before implementing the solution:

On R6:

```
R6#Show ip bgp | B Network

   Network          Next Hop      Metric LocPrf Weight Path
*> 200.1.1.1/32     2.2.2.2                        0 200 300 100 i
*> 200.1.1.2/32     2.2.2.2                  0     0 200 i
*> 200.1.1.3/32     2.2.2.2                        0 200 300 i
*> 200.1.1.4/32     2.2.2.2                        0 200 300 100 400 i
*> 200.1.1.5/32     2.2.2.2                        0 200 300 500 i
*> 200.1.1.6/32     0.0.0.0                  0    32768 i
*> 200.1.1.7/32     2.2.2.2                        0 200 300 100 700 i
*> 200.1.1.8/32     2.2.2.2                        0 200 300 100 800 i
*> 200.1.1.9/32     2.2.2.2                        0 200 300 100 800 933 900 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 200.1.1.10/32    2.2.2.2                          0 200 300 100 i
*> 200.1.1.20/32    2.2.2.2                          0 200 300 i
*> 200.1.1.30/32    2.2.2.2                          0 200 300 100 800 933 i
*> 200.1.1.40/32    2.2.2.2                          0 200 300 100 800 933 944 i
```

On R2:

The following regular expression identifies the prefixes that do not have any AS number in their AS-Path, these are the prefixes that have originated in the local AS, let's check:

```
R2#Show ip bgp regexp ^$

BGP table version is 338, local router ID is 200.1.1.2
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
*> 200.1.1.2/32     0.0.0.0                  0         32768 i
```

NOTE: The AS-PATH is empty, meaning that it does not include any AS number. Let's translate the regular expression:

Regular Expression	Meaning
^	The carrot means the start of string
$	The dollar means the end of string

An AS-Path access-list is configured to reference all the prefixes that are originated locally:

```
R2(config)#IP as-path access-list 1 permit ^$
```

Since the as-path access-list can not be referenced directly in the Neighbor command, a route-map called "tst" is configured to reference the as-path access-list:

```
R2(config)#Route-map tst permit 10
R2(config-route-map)#match as-path 1
```

The route-map is referenced in the neighbor command outbound to the directly connected neighbors R6 in AS 600 and SW2 in AS 300:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 20.1.1.20 route-map tst out
R2(config-router)#Neighbor 6.6.6.6 route-map tst out
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The implement the changes:

```
R2#Clear ip bgp * out
```

To verify the configuration:

On R2:

```
R2#Show ip bgp neighbor 20.1.1.20 advertised-routes

BGP table version is 14, local router ID is 200.1.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-
external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.2/32     0.0.0.0                  0         32768 i

Total number of prefixes 1

R2#Show ip bgp neighbor 6.6.6.6 advertised-routes

BGP table version is 14, local router ID is 200.1.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-
external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.2/32     0.0.0.0                  0         32768 i

Total number of prefixes 1
```

On R6:

```
R6#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.2/32     2.2.2.2                  0             0 200 i
*> 200.1.1.6/32     0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

Configure SW4 such that it filters any prefixe/s that has originated or traversed AS 200. DO NOT use a route-map to accomplish this task.

In this case the "_200_" regexp can be used, let's translate the regular expression:

Regular Expression	Meaning
_	The underscore matches a space, comma, left brace ({) or right brace (}), beginning and the end of an input string

Therefore, the "_200_" regexp means any prefix/es that have passed through AS 200, REMEMBER THAT 200.1.1.6/32 WILL NOT BE THERE BECAUSE OF THE FILTER THAT WAS APPLIED IN THE PREVIOUS TASK.

Let's test it:

On SW4:

```
SW4#Show ip bgp regexp _200_

BGP table version is 16, local router ID is 200.1.1.40
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
*> 200.1.1.2/32     30.1.1.30                        0 933 800 100 300 200 i
```

To see the entire BGP table on SW4:

```
SW4#Show ip bgp | B Network

   Network          Next Hop        Metric LocPrf Weight Path
*> 200.1.1.1/32     30.1.1.30                        0 933 800 100 i
*> 200.1.1.2/32     30.1.1.30                        0 933 800 100 300 200 i
*> 200.1.1.3/32     30.1.1.30                        0 933 800 100 300 i
*> 200.1.1.4/32     30.1.1.30                        0 933 800 100 400 i
*> 200.1.1.5/32     30.1.1.30                        0 933 800 100 500 i
*> 200.1.1.7/32     30.1.1.30                        0 933 800 100 700 i
*> 200.1.1.8/32     30.1.1.30                        0 933 800 i
*> 200.1.1.9/32     30.1.1.30                        0 933 900 i
*> 200.1.1.10/32    30.1.1.30                        0 933 800 100 i
*> 200.1.1.20/32    30.1.1.30                        0 933 800 100 300 i
*> 200.1.1.30/32    30.1.1.30              0         0 933 i
*> 200.1.1.40/32    0.0.0.0                0     32768 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Now that we know that this regular expression works, let's implement the solution:

```
SW4(config)#IP as-path access-list 1 deny _200_
SW4(config)#IP as-path access-list 1 permit .*
```

Instead of configuring a route-map that references the as-path access-list, and a Neighbor command referencing the route-map inbound; in the following configuration, the "as-path access-list" configuration is referenced using the "filter-list" option:

```
SW4(config)#Router bgp 944
SW4(config-router)#Neighbor 30.1.1.30 filter-list 1 in
```

To implement the change:

```
SW4#Clear ip bgp * in
```

To verify the configuration:

On SW4:

```
SW4#Show ip bgp regexp _200_
SW4#

SW4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     30.1.1.30                             0 933 800 100 i
*> 200.1.1.3/32     30.1.1.30                             0 933 800 100 300 i
*> 200.1.1.4/32     30.1.1.30                             0 933 800 100 400 i
*> 200.1.1.5/32     30.1.1.30                             0 933 800 100 500 i
*> 200.1.1.7/32     30.1.1.30                             0 933 800 100 700 i
*> 200.1.1.8/32     30.1.1.30                             0 933 800 i
*> 200.1.1.9/32     30.1.1.30                             0 933 900 i
*> 200.1.1.10/32    30.1.1.30                             0 933 800 100 i
*> 200.1.1.20/32    30.1.1.30                             0 933 800 100 300 i
*> 200.1.1.30/32    30.1.1.30                0            0 933 i
*> 200.1.1.40/32    0.0.0.0                  0        32768 i
```

Task 5

Configure R9 to filter any prefix/es that has originated in AS 400.

© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

Step #1

Test the regular expression:

On R9:

```
R9#Show ip bgp regexp _400$

BGP table version is 18, local router ID is 200.1.1.9
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
*> 200.1.1.4/32     30.1.1.30                              0 933 800 100 400 i
```

Let's translate the regular expression:

Basically you are referencing the first AS number in the path; or another way of looking at this is to say that we don't care about the prefixes that have been advertised by any AS/es in the path, but we like to match the prefix that has been advertised by the last ASN in the path. This regular expression is different to the previous one, in the previous one we were saying that if you see AS 200 anywhere in the path (The prefixes that have originated or the ones that have traversed through AS 200) and we used the "_200_", in this case we are asking for prefixes that have originated and NOT traversed through AS 400. Let's see the difference, in the following example AS 800 is used:

```
R9#Show ip bgp regexp _800_

BGP table version is 17, local router ID is 200.1.1.9
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-
external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     30.1.1.30                              0 933 800 100 i
*> 200.1.1.2/32     30.1.1.30                              0 933 800 100 300 200 i
*> 200.1.1.3/32     30.1.1.30                              0 933 800 100 300 i
*> 200.1.1.4/32     30.1.1.30                              0 933 800 100 400 i
*> 200.1.1.5/32     30.1.1.30                              0 933 800 100 500 i
*> 200.1.1.7/32     30.1.1.30                              0 933 800 100 700 i
*> 200.1.1.8/32     30.1.1.30                              0 933 800 i
*> 200.1.1.10/32    30.1.1.30                              0 933 800 100 i
*> 200.1.1.20/32    30.1.1.30                              0 933 800 100 300 i

R9#Show ip bgp regexp _800$
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
BGP table version is 17, local router ID is 200.1.1.9
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-
external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.8/32     30.1.1.30                              0 933 800 i
```

Step #2

Implement the solution:

On R9:

```
R9(config)#IP as-path access-list 1 deny _400$
R9(config)#IP as-path access-list 1 permit .*

R9(config)#Router bgp 900
R9(config-router)#Neighbor 30.1.1.30 filter-list 1 in

R9#Clear ip bgp * in
```

To verify the configuration:

On R9:

```
R9#Show ip bgp

BGP table version is 19, local router ID is 200.1.1.9
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
*> 200.1.1.1/32     30.1.1.30                              0 933 800 100 i
*> 200.1.1.2/32     30.1.1.30                              0 933 800 100 300 200 i
*> 200.1.1.3/32     30.1.1.30                              0 933 800 100 300 i
*> 200.1.1.5/32     30.1.1.30                              0 933 800 100 500 i
*> 200.1.1.7/32     30.1.1.30                              0 933 800 100 700 i
*> 200.1.1.8/32     30.1.1.30                              0 933 800 i
*> 200.1.1.9/32     0.0.0.0                  0         32768 i
*> 200.1.1.10/32    30.1.1.30                              0 933 800 100 i
*> 200.1.1.20/32    30.1.1.30                              0 933 800 100 300 i
*> 200.1.1.30/32    30.1.1.30                0             0 933 i
*> 200.1.1.40/32    30.1.1.30                              0 933 944 i
```


© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

Task 6

Configure R5 to filter AS 800 as the second AS in their AS-Path.

Before we test and configure this task, let's see which prefixes have AS 222 as the second AS in their AS-Path:

On R5:

```
R5#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  200.1.1.1/32     3.3.3.3                                0 300 100 i
*                   4.4.4.4                                0 400 100 i
*>                  1.1.1.1                  0             0 100 i
*  200.1.1.2/32     1.1.1.1                                0 100 300 200 i
*>                  3.3.3.3                                0 300 200 i
*> 200.1.1.3/32     3.3.3.3                  0             0 300 i
*                   4.4.4.4                                0 400 100 300 i
*                   1.1.1.1                                0 100 300 i
*  200.1.1.4/32     3.3.3.3                                0 300 100 400 i
*>                  4.4.4.4                  0             0 400 i
*                   1.1.1.1                                0 100 400 i
*> 200.1.1.5/32     0.0.0.0                  0         32768 i
*  200.1.1.7/32     3.3.3.3                                0 300 100 700 i
*                   1.1.1.1                                0 100 700 i
*>                  4.4.4.4                                0 400 700 i
*  200.1.1.8/32     3.3.3.3                                0 300 100 800 i
*                   4.4.4.4                                0 400 100 800 i
*>                  1.1.1.1                                0 100 800 i
*  200.1.1.9/32     4.4.4.4                                0 400 100 800 933 900 i
*                   3.3.3.3                                0 300 100 800 933 900 i
*>                  1.1.1.1                                0 100 800 933 900 i
*  200.1.1.10/32    4.4.4.4                                0 400 100 i
*                   3.3.3.3                                0 300 100 i
*>                  1.1.1.1                                0 100 i
*  200.1.1.20/32    1.1.1.1                                0 100 300 i
*>                  3.3.3.3                                0 300 i
*  200.1.1.30/32    4.4.4.4                                0 400 100 800 933 i
*                   3.3.3.3                                0 300 100 800 933 i
*>                  1.1.1.1                                0 100 800 933 i
*  200.1.1.40/32    4.4.4.4                                0 400 100 800 933 944 i
*                   3.3.3.3                                0 300 100 800 933 944 i
*>                  1.1.1.1                                0 100 800 933 944 i
```

We can see that ONLY four prefixes have AS 800 as the second AS in their AS-Path. 200.1.1.8, 200.1.1.9, 200.1.1.30, and 200.1.1.40.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's test the regular expression:

```
R5#Show ip bgp regexp ^([0-9]+)_800

BGP table version is 25, local router ID is 200.1.1.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-
external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.8/32     1.1.1.1                              0 100 800 i
*> 200.1.1.9/32     1.1.1.1                              0 100 800 933 900 i
*> 200.1.1.30/32    1.1.1.1                              0 100 800 933 i
*> 200.1.1.40/32    1.1.1.1                              0 100 800 933 944 i
```

Based on the output of the above show command we can see that the regular expression is correct.
To translate the regexp:

Regular Expression	Meaning
^([0-9]+)	We know that carrot means the beginning of the string; the numbers within brackets specify a range of valid numbers, and the "+" means one or more instance of the numbers within the brackets. This means any combination and any number of digit/s. This can easily match any AS number.
_800	As previously explained, the underscore can also mean space, then 800.

Therefore, the regexp used in this task means any AS followed by AS 800.

Let's implement the solution:

On R5:

```
R5(config)#IP as-path access-list 1 deny ^([0-9]+)_800
R5(config)#IP as-path access-list 1 permit .*

R5(config)#Router bgp 500
R5(config-router)#Neighbor 4.4.4.4 filter-list 1 in
R5(config-router)#Neighbor 3.3.3.3 filter-list 1 in
R5(config-router)#Neighbor 1.1.1.1 filter-list 1 in
```

To enforce the policy:

```
R5#Clear ip bgp * in
```


© 2014 Narbik Kocharians, All rights reserved
https://t.me/learningnets

To verify the configuration:

On R5:

```
R5#Show ip bgp regexp ^([0-9]+)_800
R5#

R5#Show ip bgp | B Network
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	200.1.1.1/32	3.3.3.3			0	300 100 i
*		4.4.4.4			0	400 100 i
*>		1.1.1.1	0		0	100 i
*	200.1.1.2/32	1.1.1.1			0	100 300 200 i
*>		3.3.3.3			0	300 200 i
*>	200.1.1.3/32	3.3.3.3	0		0	300 i
*		4.4.4.4			0	400 100 300 i
*		1.1.1.1			0	100 300 i
*	200.1.1.4/32	3.3.3.3			0	300 100 400 i
*>		4.4.4.4	0		0	400 i
*		1.1.1.1			0	100 400 i
*>	200.1.1.5/32	0.0.0.0	0		32768	i
*	200.1.1.7/32	3.3.3.3			0	300 100 700 i
*		1.1.1.1			0	100 700 i
*>		4.4.4.4			0	400 700 i
*>	200.1.1.8/32	3.3.3.3			0	300 100 800 i
*		4.4.4.4			0	400 100 800 i
*	200.1.1.9/32	4.4.4.4			0	400 100 800 933 900 i
*>		3.3.3.3			0	300 100 800 933 900 i
*	200.1.1.10/32	4.4.4.4			0	400 100 i
*		3.3.3.3			0	300 100 i
*>		1.1.1.1			0	100 i
*	200.1.1.20/32	1.1.1.1			0	100 300 i
*>		3.3.3.3			0	300 i
*	200.1.1.30/32	4.4.4.4			0	400 100 800 933 i
*>		3.3.3.3			0	300 100 800 933 i
*	200.1.1.40/32	4.4.4.4			0	400 100 800 933 944 i
*>		3.3.3.3			0	300 100 800 933 944 i

We can see that the prefixes are still in the BGP table, but you can see that the AS 800 is NOT the second AS in the AS-Path.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 7

Configure R8 to ONLY allow prefixes from existing and future ASes that are directly connected to its AS (AS 800).

Before the solution is implemented, let's see the existing prefixes in the BGP table of R8:

On R8:

```
R8#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     10.1.1.10                             0 100 i
*> 200.1.1.2/32     10.1.1.10                             0 100 300 200 i
*> 200.1.1.3/32     10.1.1.10                             0 100 300 i
*> 200.1.1.4/32     10.1.1.10                             0 100 400 i
*> 200.1.1.5/32     10.1.1.10                             0 100 500 i
*> 200.1.1.7/32     10.1.1.10                             0 100 700 i
*> 200.1.1.8/32     0.0.0.0                  0         32768 i
*> 200.1.1.9/32     30.1.1.30                             0 933 900 i
*> 200.1.1.10/32    10.1.1.10                0            0 100 i
*> 200.1.1.20/32    10.1.1.10                             0 100 300 i
*> 200.1.1.30/32    30.1.1.30                0            0 933 i
*> 200.1.1.40/32    30.1.1.30                             0 933 944 i
```

Let's test the solution before it's implemented:

On R8

```
R8#Show ip bgp regexp ^[0-9]+$

BGP table version is 191, local router ID is 200.1.1.8
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
*> 200.1.1.1/32     10.1.1.10                             0 100 i
*> 200.1.1.10/32    10.1.1.10                0            0 100 i
*> 200.1.1.30/32    30.1.1.30                0            0 933 i
```

The output of the above show command proves that the regular expression ONLY shows the prefixes that have originated in the directly connected ASes.

Let's translate the regular expression:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Regular expression	Meaning
^	Start of string, in BGP it means directly connected
[0-9]	Range of characters, in BGP it means any value from 0 to 9 Inclusively
+	means one or more of the instance; in this configuration it means that one or more instance of [0-9].
$	End of string.

Therefore, the regexp used in this solution means any AS number as long as it's directly connected.

Let's implement the solution and test:

On R8:

The directly connected ASes are the ONLY ones permitted:

```
R8(config)#IP AS-path access-list 1 permit ^[0-9]+$
```

The AS-Path access-list is applied through all neighbors inbound:

```
R8(config)#Router bgp 800
R8(config-router)#Neighbor 10.1.1.10 filter-list 1 in
R8(config-router)#Neighbor 30.1.1.30 filter-list 1 in
```

To enforce the policy without tearing down the session:

```
R8#Clear ip bgp * in
```

To test the configuration:

On R8:

```
R8#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     10.1.1.10                              0 100 i
*> 200.1.1.8/32     0.0.0.0                  0         32768 i
*> 200.1.1.10/32    10.1.1.10                0             0 100 i
*> 200.1.1.30/32    30.1.1.30                0             0 933 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 8

Configure SW1 to filter prefix/es that originated in AS 200 and traversed through AS 300.

Let's verify the BGP table of SW1:

On SW1:

```
SW1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i200.1.1.1/32     1.1.1.1                  0    100      0 i
*>i200.1.1.2/32     1.1.1.1                  0    100      0 300 200 i
*>i200.1.1.3/32     1.1.1.1                  0    100      0 300 i
*  200.1.1.4/32     7.7.7.7                                0 700 400 i
*>i                 1.1.1.1                  0    100      0 400 i
*>i200.1.1.5/32     1.1.1.1                  0    100      0 500 i
*> 200.1.1.7/32     7.7.7.7                  0             0 700 i
*> 200.1.1.8/32     8.8.8.8                  0             0 800 i
*> 200.1.1.10/32    0.0.0.0                  0         32768 i
*>i200.1.1.20/32    1.1.1.1                  0    100      0 300 i
*> 200.1.1.30/32    8.8.8.8                                0 800 933 i
```

The 200.1.1.2/32 prefix is the ONLY prefix that has originated in AS 200 that has traversed through AS 300.

Let's test the solution:

```
SW1#Sh ip bgp reg _300_200$

BGP table version is 496, local router ID is 200.1.1.10
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
*>i200.1.1.2/32     1.1.1.1                  0    100      0 300 200 i
```

The solution seems to be correct, let's translate the solution:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Regular Expression	Meaning
_300_200$	This regexp identifies all the prefixes that have originated in AS 200 and have traversed through AS 300. NOTE: AS 300 is directly connected neighbor of AS 200.

To implement the solution:

On SW1:

```
SW1(config)#IP AS-Path access-list 1 permit _300_200$

SW1(config)#Route-map tst deny 10
SW1(config-route-map)#match as-path 1
SW1(config)#Route-map tst permit 90

SW1(config)#Router bgp 100
SW1(config-router)#Neighbor 1.1.1.1 route-map tst in
SW1(config-router)#Neighbor 7.7.7.7 route-map tst in

SW1#Clear ip bgp * in
```

To test the configuration:

On SW1:

```
SW1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i200.1.1.1/32     1.1.1.1                  0    100      0 i
*>i200.1.1.3/32     1.1.1.1                  0    100      0 300 i
*  200.1.1.4/32     7.7.7.7                                0 700 400 i
*>i                 1.1.1.1                  0    100      0 400 i
*>i200.1.1.5/32     1.1.1.1                  0    100      0 500 i
*> 200.1.1.7/32     7.7.7.7                  0             0 700 i
*> 200.1.1.8/32     8.8.8.8                  0             0 800 i
*> 200.1.1.10/32    0.0.0.0                  0         32768 i
*>i200.1.1.20/32    1.1.1.1                  0    100      0 300 i
*> 200.1.1.30/32    8.8.8.8                                0 800 933 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 9

Configure R7 to prepend its AS number 5 times when it advertises its Loopback 1 interface to its directly connected neighbors.

On R7:

```
R7(config)#Ip prefix-list NET700 seq 5 permit 200.1.1.7/32

R7(config)#Route-map Task9 permit 10
R7(config-route-map)#Match ip addr NET700
R7(config-route-map)#Set as-path prepend 700 700 700 700
R7(config)#Route-map Task9 permit 90

R7(config)#Router bgp 700
R7(config-router)#Neighbor 4.4.4.4 route-map Task9 out
R7(config-router)#Neighbor 1.1.1.1 route-map Task9 out

R7#Clear ip bgp * out
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  200.1.1.1/32     5.5.5.5                                0 500 100 i
*                   7.7.7.7                                0 700 100 i
*>                  1.1.1.1                  0             0 100 i
*  200.1.1.2/32     5.5.5.5                                0 500 300 200 i
*>                  1.1.1.1                                0 100 300 200 i
*  200.1.1.3/32     5.5.5.5                                0 500 300 i
*                   7.7.7.7                                0 700 100 300 i
*>                  1.1.1.1                                0 100 300 i
*> 200.1.1.4/32     0.0.0.0                  0         32768 i
*  200.1.1.5/32     7.7.7.7                                0 700 100 500 i
*                   1.1.1.1                                0 100 500 i
*>                  5.5.5.5                  0             0 500 i
*  200.1.1.7/32     5.5.5.5                                0 500 100 700 700 700 700 700 i
*                   1.1.1.1                                0 100 700 700 700 700 700 i
*>                  7.7.7.7                  0             0 700 700 700 700 700 i
*  200.1.1.8/32     5.5.5.5                                0 500 300 100 800 i
*                   7.7.7.7                                0 700 100 800 i
*>                  1.1.1.1                                0 100 800 i
*  200.1.1.10/32    5.5.5.5                                0 500 100 i
*>                  1.1.1.1                                0 100 i
*                   7.7.7.7                                0 700 100 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*  200.1.1.20/32    5.5.5.5                                0 500 300 i
*>                  1.1.1.1                                0 100 300 i
*                   7.7.7.7                                0 700 100 300 i
*  200.1.1.30/32    5.5.5.5                                0 500 300 100 800 933 i
*>                  1.1.1.1                                0 100 800 933 i
*                   7.7.7.7                                0 700 100 800 933 i
```

Task 10

Configure R4 to filter any prefix/es that have prepended their AS number multiple times.

Let's test the solution:

On R4:

```
R4#Show ip bgp regexp ^([0-9]+)(_\1)+$

BGP table version is 318, local router ID is 200.1.1.4
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
*> 200.1.1.7/32     7.7.7.7                  0             0 700 700 700 700 700 i
```

NOTE: The above regular expression filters the output based on the directly connected AS/es, but the task states to filter any prefix/es that have prepended their AS number multiple times, and it does NOT state directly connected. Let's try the following regular expression:

```
R4#Sh ip bgp regexp _([0-9]+)(_\1)+$

BGP table version is 55, local router ID is 200.1.1.4
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
*  200.1.1.7/32     5.5.5.5                                0 500 100 700 700 700 700 700 i
*                   1.1.1.1                                0 100 700 700 700 700 700 i
*>                  7.7.7.7                  0             0 700 700 700 700 700 i
```

We can see the difference, the previous regular expression started with "^" whereas, the above regular expression used the "_".

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Now that we know the regular expression is correct, let's translate it:

Regular Expression	Meaning
_([0-9]+)	We have seen this in the previous example which matches on ANY AS number.
(_\1)+$	The _\1 means one or more instance of the previously matched string. For example: ^(1)_\1_\1 matches a directly connected AS that has prepended its AS number 3 times or more. The "+" means one or more instance and the "$" means the end of string. If you were NOT sure if any of your directly connected AS/es have performed AS-Path Prepending, then ^([0-9]+)(_\1)*$ can be used, the difference between this regular expression and the one that we have used as the solution to this task is the "*" before the "$" sign, which means 0 or more instance. But if one of your directly connected ASes has indeed prepended their AS number multiple times, the regexp that is used in this task can be used, even though in this example they will both produce the same result.

The following AS-Path access-list denies all directly connected ASes that have prepended their AS number one or more times:

```
R4(config)#IP AS-Path access-list 4 deny _([0-9]+)(_\1)+$
R4(config)#IP AS-Path access-list 4 permit .*
```

The following filter-list is applied to neighbors R5, R1 and R7 respectively:

```
R4(config)#Router bgp 400
R4(config-router)#Neighbor 5.5.5.5 filter-list 4 in
R4(config-router)#Neighbor 1.1.1.1 filter-list 4 in
R4(config-router)#Neighbor 7.7.7.7 filter-list 4 in
```

To enforce the configured policy:

```
R4#Clear ip bgp * in
```

To test the configuration:

On R4:

```
R4#Show ip bgp regexp _([0-9]+)(_\1)+$
R4#
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  200.1.1.1/32     5.5.5.5                              0 500 100 i
*                   7.7.7.7                              0 700 100 i
*>                  1.1.1.1                   0          0 100 i
*  200.1.1.2/32     5.5.5.5                              0 500 300 200 i
*>                  1.1.1.1                              0 100 300 200 i
*  200.1.1.3/32     5.5.5.5                              0 500 300 i
*                   7.7.7.7                              0 700 100 300 i
*>                  1.1.1.1                              0 100 300 i
*> 200.1.1.4/32     0.0.0.0                   0      32768 i
*  200.1.1.5/32     7.7.7.7                              0 700 100 500 i
*                   1.1.1.1                              0 100 500 i
*>                  5.5.5.5                   0          0 500 i
*  200.1.1.8/32     5.5.5.5                              0 500 300 100 800 i
*                   7.7.7.7                              0 700 100 800 i
*>                  1.1.1.1                              0 100 800 i
*  200.1.1.10/32    5.5.5.5                              0 500 100 i
*>                  1.1.1.1                              0 100 i
*                   7.7.7.7                              0 700 100 i
*  200.1.1.20/32    5.5.5.5                              0 500 300 i
*>                  1.1.1.1                              0 100 300 i
*                   7.7.7.7                              0 700 100 300 i
*  200.1.1.30/32    5.5.5.5                              0 500 300 100 800 933 i
*>                  1.1.1.1                              0 100 800 933 i
*                   7.7.7.7                              0 700 100 800 933 i
```

Task 11

Configure R6 to filter all prefixes that have originated in AS 200.

Let's see the BGP table of R6:

On R6:

```
R6#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.2/32     2.2.2.2                  0             0 200 i
*> 200.1.1.6/32     0.0.0.0                  0         32768 i
```

To test the regular expression:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6#Show ip bgp regexp ^200$

BGP table version is 86, local router ID is 200.1.1.6
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
*> 200.1.1.2/32     2.2.2.2                  0             0 200 i
```

To translate the regular expression:

Since the "^" means the start of string and the "$" means the end of string, this can ONLY mean all the existing and future prefixes that have originated in directly connected AS (AS 200).

Let's implement the solution:

On R6:

The following AS-Path access-list denies all existing and future prefixes that have originated in AS 200, and permits all other prefixes.

```
R6(config)#IP AS-Path access-list 6 deny ^200$
R6(config)#IP AS-Path access-list 6 permit .*
```

The AS-Path access-list is referenced by the filter-list for neighbor R2 inbound:

```
R6(config)#Router bgp 600
R6(config-router)#Neighbor 2.2.2.2 filter-list 6 in
```

The policy is enforced:

```
R6#Clear ip bgp * in
```

To test the configuration:

On R6:

```
R6#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.6/32     0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 12

Configure SW4 to filter the prefixes that have originated from AS 933's directly connected neighbors. This should NOT override the previous policy implemented in task 4.

In task 4, we were asked to filter any prefix that had traversed through AS 200. To accomplish this task we need to add to the existing "IP AS-Path access-list 1" configured in task 4. To test the regular expression:

On SW4:

```
SW4#Sh ip bgp regexp ^933_[0-9]+$

BGP table version is 70, local router ID is 200.1.1.40
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
*> 200.1.1.8/32     30.1.1.30                              0 933 800 i
*> 200.1.1.9/32     30.1.1.30                              0 933 900 i
```

We can clearly see that ASes 800 and 900 are the directly connected ASes to AS 933, and prefixes 200.1.1.8/32 and 200.1.1.9/32 are the ones advertised by these two ASes. Since this regular expression gives us the desireable result let's implement it.

The previously configured IP AS-Path access-list 1 is removed, and the new IP AS-Path access-list is added:

```
SW4(config)#No ip as-path access-list 1
SW4(config)#IP AS-Path access-list 1 deny _200_
SW4(config)#IP AS-Path access-list 1 deny ^933_[0-9]+$
SW4(config)#IP AS-Path access-list 1 permit .*
```

NOTE: Since the filter-list is still configured to reference the "IP AS-Path access-list" there is no need to re-apply it.

To enforce the policy:

```
SW4#Clear ip bgp * in
```

To test the configuration:

On SW4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     30.1.1.30                             0 933 800 100 i
*> 200.1.1.10/32    30.1.1.30                             0 933 800 100 i
*> 200.1.1.30/32    30.1.1.30                0            0 933 i
*> 200.1.1.40/32    0.0.0.0                  0        32768 i
```

Task 13

Configure R7 such that it discards paths for any prefix/es that have more than two AS hops.

On R7:

```
R7#Sh ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     10.1.1.10                             0 100 i
*                   4.4.4.4                               0 400 100 i
*> 200.1.1.2/32     4.4.4.4                               0 400 100 300 200 i
*> 200.1.1.3/32     10.1.1.10                             0 100 300 i
*                   4.4.4.4                               0 400 100 300 i
*  200.1.1.4/32     10.1.1.10                             0 100 400 i
*>                  4.4.4.4                  0            0 400 i
*  200.1.1.5/32     4.4.4.4                               0 400 500 i
*>                  10.1.1.10                             0 100 500 i
*> 200.1.1.7/32     0.0.0.0                  0        32768 i
*  200.1.1.8/32     4.4.4.4                               0 400 100 800 i
*>                  10.1.1.10                             0 100 800 i
*  200.1.1.10/32    4.4.4.4                               0 400 100 i
*>                  10.1.1.10                0            0 100 i
*  200.1.1.20/32    4.4.4.4                               0 400 100 300 i
*>                  10.1.1.10                             0 100 300 i
*  200.1.1.30/32    4.4.4.4                               0 400 100 800 933 i
*>                  10.1.1.10                             0 100 800 933 i
```

To implement the solution:

```
R7(config)#Router bgp 700
R7(config-router)#BGP maxas-limit 2
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

The above command instructs the local router to discard routes that have more than two AS hops in their AS path segment. The default value is 75 AS-Segment lengths.

To enforce the policy:

```
R7#Clear ip bgp * in
```

To test the configuration:

On R7:

```
R7#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     10.1.1.10                             0 100 i
*                   4.4.4.4                               0 400 100 i
*> 200.1.1.3/32     10.1.1.10                             0 100 300 i
*  200.1.1.4/32     10.1.1.10                             0 100 400 i
*>                  4.4.4.4                  0            0 400 i
*  200.1.1.5/32     4.4.4.4                               0 400 500 i
*>                  10.1.1.10                             0 100 500 i
*> 200.1.1.7/32     0.0.0.0                  0        32768 i
*> 200.1.1.8/32     10.1.1.10                             0 100 800 i
*  200.1.1.10/32    4.4.4.4                               0 400 100 i
*>                  10.1.1.10                0            0 100 i
*> 200.1.1.20/32    10.1.1.10                             0 100 300 i
```

Task 14

Configure R3 to advertise Network 30.3.3.0 /24.

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 30.3.3.0 mask 255.255.255.0
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip route | Inc 30.3.3.0

B        30.3.3.0 [20/0] via 3.3.3.3, 00:13:12
```

Task 15

R1 should be configured such that if the number of routes received from R3 exceeds 10, it should shutdown the neighbor (R3). R1 should generate a console warning message when 80% of this threshold is reached. If the adjacency goes down because of this policy, R1 should restart the adjacency after 1 minutes and check the number of routes, and if they still exceed the threshold, the adjacency should once again go down and the cycle is repeated.

On R1:

```
R1(config)#Router bgp 100
```

The following command sets the maximum number of prefixes received from R3 (Neighbor 3.3.3.3) to 10, and it also instructs the router to send a console message when %80 of the configured threshold is reached, and to restart/recover the peer session after a minute.

```
R1(config-router)#Neighbor 3.3.3.3 maximum-prefix 10 80 restart 1
```

To verify the configuration:

On R1:

```
R1#Show ip bgp neighbor 3.3.3.3 | Inc Maximum|restart

  Maximum prefixes allowed 10
  Threshold for warning message 80%, restart interval 1 min
Maximum output segment queue size: 50
```

Let's see the number of prefixes advertised by R3:

```
R1#Show ip bgp neighbor 3.3.3.3 routes | Inc Total

Total number of prefixes 5

R1#Show ip bgp neighbor 3.3.3.3 routes | B Network
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 30.3.3.0/24      3.3.3.3                  0             0 300 i
*> 200.1.1.2/32     3.3.3.3                                0 300 200 i
*> 200.1.1.3/32     3.3.3.3                  0             0 300 i
*  200.1.1.5/32     3.3.3.3                                0 300 500 i
*> 200.1.1.20/32    3.3.3.3                                0 300 i

Total number of prefixes 5
```

Let's test this scenario by configuring R3 to advertise the following four networks, this will add up to nine networks which exceeds the threshold set by R1 for R3:

31.3.3.0/24, 32.3.3.0/24, 33.3.3.0/24, and 34.3.3.0/24

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 31.3.3.0 mask 255.255.255.0
R3(config-router)#Network 32.3.3.0 mask 255.255.255.0
R3(config-router)#Network 33.3.3.0 mask 255.255.255.0
R3(config-router)#Network 34.3.3.0 mask 255.255.255.0
```

On R1:

You should see the following console message stating that %80 of the configured threshold is reached:

```
*Dec  9 10:44:44.189: %BGP-4-MAXPFX: Number of prefixes received from
3.3.3.3 (afi 0) reaches 9, max 10

R1#Show ip bgp summary | Inc Neighbor|3.3.3.3

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4   300    1597    1600       48    0    0 23:53:12        9
```

Now let's advertise one more network and verify the result:

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 35.3.3.0 mask 255.255.255.0
```

On R1:

You should see the following console messages:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*Dec  9 10:49:27.637: %BGP-4-MAXPFX: Number of prefixes received from
3.3.3.3 (afi 0) reaches 10, max 10
```

The above console message states that the maximum number of prefixes has reached the set threshold and if R3 advertises one more network, R1 will tear down the adjacency by shuting down R3 as a neighbor, let's test:

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#Network 36.3.3.0 mask 255.255.255.0
```

You should see the following console messages on R3:

```
*Feb  1 21:25:55.104: %BGP-5-ADJCHANGE: neighbor 1.1.1.1 Down Peer closed
the session
*Feb  1 21:25:55.104: %BGP_SESSION-5-ADJCHANGE: neighbor 1.1.1.1 IPv4
Unicast topology base removed from session  Peer closed the session
```

On R1:

You should have received the following console messages:

```
*Dec  9 11:44:18.568: %BGP-3-MAXPFXEXCEED: Number of prefixes received
from 3.3.3.3 (afi 0): 11 exceeds limit 10

*Dec  9 11:44:18.568: %BGP-5-ADJCHANGE: neighbor 3.3.3.3 Down BGP
Notification sent

*Dec  9 11:44:18.568: %BGP-3-NOTIFICATION: sent to neighbor 3.3.3.3 3/1
(update malformed) 0 bytes
```

After a minute you should receive the following:

```
*Feb  1 21:26:59.036: %BGP-5-ADJCHANGE: neighbor 1.1.1.1 Up
*Feb  1 21:26:59.148: %BGP-5-ADJCHANGE: neighbor 1.1.1.1 Down Peer closed
the session
```

And the cycle repeats.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 16

Configure SW4 to advertise networks 40.0.0.0/8 – 49.0.0.0/8, these networks were pre-configured in SW4's routing table as static routes pointing to Null 0.

On SW4:

```
SW4#Show ip route static

S     49.0.0.0/8 is directly connected, Null0
S     48.0.0.0/8 is directly connected, Null0
S     42.0.0.0/8 is directly connected, Null0
S     43.0.0.0/8 is directly connected, Null0
S     40.0.0.0/8 is directly connected, Null0
S     41.0.0.0/8 is directly connected, Null0
S     46.0.0.0/8 is directly connected, Null0
S     47.0.0.0/8 is directly connected, Null0
S     44.0.0.0/8 is directly connected, Null0
S     45.0.0.0/8 is directly connected, Null0
```

To advertise the prefixes:

```
SW4(config)#Router bgp 944
SW4(config-router)#Network 40.0.0.0
SW4(config-router)#Network 41.0.0.0
SW4(config-router)#Network 42.0.0.0
SW4(config-router)#Network 43.0.0.0
SW4(config-router)#Network 44.0.0.0
SW4(config-router)#Network 45.0.0.0
SW4(config-router)#Network 46.0.0.0
SW4(config-router)#Network 47.0.0.0
SW4(config-router)#Network 48.0.0.0
SW4(config-router)#Network 49.0.0.0
```

Task 17

Since SW3 is running low on system resources it can not handle networks advertised by SW4 in the previous task. Configure the appropriate device/s such that networks 40.0.0.0/8 – 49.0.0.0/8 are filtered.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The obvious thing to do is to filter these networks inbound on SW3. **To see that this is the wrong solution**, let's configure "Soft reconfiguration inbound" on SW3:

On SW3:

```
SW3(config)#Router bgp 933
SW3(config-router)#Neighbor 40.1.1.40 soft-reconfiguration inbound

SW3#Clear ip bgp * in
```

NOTE: Contrary to popular belief the "Show ip bgp received-routes" displays the unprocessed packets, meaning that these prefixes have not gone through filtering or any policy at all; they have been received from a given neighbor and they are in "Adj-RIBS-IN" and NOT the "Local-RIBS" which is the actual BGP table.

```
SW3#Show ip bgp neighbor 40.1.1.40 received-routes

BGP table version is 91, local router ID is 200.1.1.30
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
*> 40.0.0.0         40.1.1.40                0             0 944 i
*> 41.0.0.0         40.1.1.40                0             0 944 i
*> 42.0.0.0         40.1.1.40                0             0 944 i
*> 43.0.0.0         40.1.1.40                0             0 944 i
*> 44.0.0.0         40.1.1.40                0             0 944 i
*> 45.0.0.0         40.1.1.40                0             0 944 i
*> 46.0.0.0         40.1.1.40                0             0 944 i
*> 47.0.0.0         40.1.1.40                0             0 944 i
*> 48.0.0.0         40.1.1.40                0             0 944 i
*> 49.0.0.0         40.1.1.40                0             0 944 i
*> 200.1.1.40/32    40.1.1.40                0             0 944 i

Total number of prefixes 11
```

Let's configure an inbound filter on SW3 to filter prefixes 40.0.0.0/8 – 49.0.0.0/8 and verify:

```
SW3(config)#IP Prefix-list NET deny 40.0.0.0/8
SW3(config)#IP Prefix-list NET deny 41.0.0.0/8
SW3(config)#IP Prefix-list NET deny 42.0.0.0/8
SW3(config)#IP Prefix-list NET deny 43.0.0.0/8
SW3(config)#IP Prefix-list NET deny 44.0.0.0/8
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
SW3(config)#IP Prefix-list NET deny 45.0.0.0/8
SW3(config)#IP Prefix-list NET deny 46.0.0.0/8
SW3(config)#IP Prefix-list NET deny 47.0.0.0/8
SW3(config)#IP Prefix-list NET deny 48.0.0.0/8
SW3(config)#IP Prefix-list NET deny 49.0.0.0/8
SW3(config)#IP Prefix-list NET permit 0.0.0.0/0 LE 32

SW3(config)#Router bgp 933
SW3(config-router)#Neighbor 40.1.1.40 prefix-list NET in

SW3#Clear ip bgp * in
```

Let's check the BGP table of SW3:

```
SW3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.1/32     8.8.8.8                                0 800 100 i
*> 200.1.1.8/32     8.8.8.8                  0             0 800 i
*> 200.1.1.9/32     9.9.9.9                  0             0 900 i
*> 200.1.1.10/32    8.8.8.8                                0 800 100 i
*> 200.1.1.30/32    0.0.0.0                  0         32768 i
*> 200.1.1.40/32    40.1.1.40                0             0 944 i
```

So you may think that this solution worked, but let's see if SW3 truly consumes resources, let's check SW3's "Adj-RIBs-In" table:

On SW3:

```
SW3#Show ip bgp neighbor 40.1.1.40 received-routes | B Network

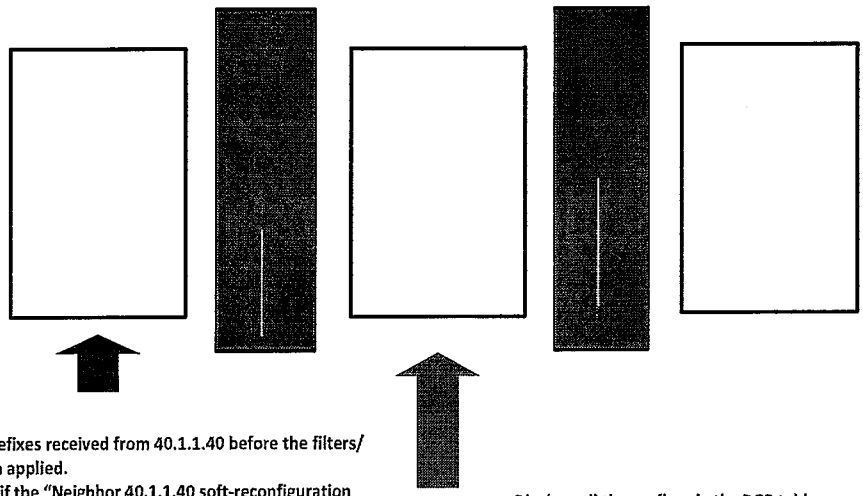
   Network          Next Hop            Metric LocPrf Weight Path
*  40.0.0.0         40.1.1.40                0             0 944 i
*  41.0.0.0         40.1.1.40                0             0 944 i
*  42.0.0.0         40.1.1.40                0             0 944 i
*  43.0.0.0         40.1.1.40                0             0 944 i
*  44.0.0.0         40.1.1.40                0             0 944 i
*  45.0.0.0         40.1.1.40                0             0 944 i
*  46.0.0.0         40.1.1.40                0             0 944 i
*  47.0.0.0         40.1.1.40                0             0 944 i
*  48.0.0.0         40.1.1.40                0             0 944 i
*  49.0.0.0         40.1.1.40                0             0 944 i
*> 200.1.1.40/32    40.1.1.40                0             0 944 i

Total number of prefixes 11
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The output of the above show command reveals that the prefixes that were filtered are still consuming memory/resources on SW3; when a BGP speaker filters certain prefix/es the actual filtering is performed in an area between the "Adj-Ribs-in" and "Local-Ribs", the following diagram demonstrates this process.



Displays all the prefixes received from 40.1.1.40 before the filters/ policies have been applied.
BUT it only works if the "Neighbor 40.1.1.40 soft-reconfiguration inbound" is configured for that neighbor.

- Displays all the prefixes in the BGP table
- Displays all the prefixes received from neighbor 40.1.1.40 after the policies/filters have been applied
- Displays all the prefixes advertised to 40.1.1.40 neighbor.

So we can see how the solution did NOT work. The correct solution is to configure Outbound Route Filtering (ORF).

Prefix based ORF feature uses ORF send and receive capabilities to minimize the number of routes that are exchanged between two BGP peers.

This feature can reduce the amount of system resources to process the routes. In this case when SW4 advertises networks 40.0.0.0/8 – 49.0.0.0/8 to SW3, even though SW3 has configured an inbound prefix-list to filter these routes, it still has to receives the routes before it can apply the filter. The routes are received and are placed in the "Adj-RIBs-In" table and the actual configured filter resides between this table and the "Local-Ribs" (The BGP Table). Therefore, as the routes go from "Adj-Ribs-In" into "Local-Ribs" they are filtered.

Let's configure ORF and test:

On SW3:

This device will be configured to send all the prefixes that it has filtered inbound to SW4:

```
SW3(config)#Router bgp 933
SW3(config-router)#Neighbor 40.1.1.40 capability orf prefix-list send
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On SW4:

SW4 is configured to receive the prefixes from SW3 and filter them outbound which occurs between "Local-Ribs" and "Adj-Ribs-Out", this means that SW3 won't even receive these prefixes in its "Adj-Ribs-In".

```
SW4(config)#Router bgp 944
SW4(config-router)#Neighbor 30.1.1.30 capability orf prefix-list receive
```

To verify the configuration:

On SW3:

```
SW3#Show ip bgp neighbor 40.1.1.40 received-routes | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.40/32    40.1.1.40                0             0 944 i
```

The output of the above command reveals that ONLY prefix 200.1.1.40/32 was advertised to SW3.

On SW4:

```
SW4#Show ip bgp neighbor 30.1.1.30 advertised-routes | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 200.1.1.40/32    0.0.0.0                  0         32768 i

Total number of prefixes 1
```

We can clearly see that SW4 filtered the routes and it is no longer advertising the routes to SW3, and since SW3 no longer receives the routes it does not use additional resources.

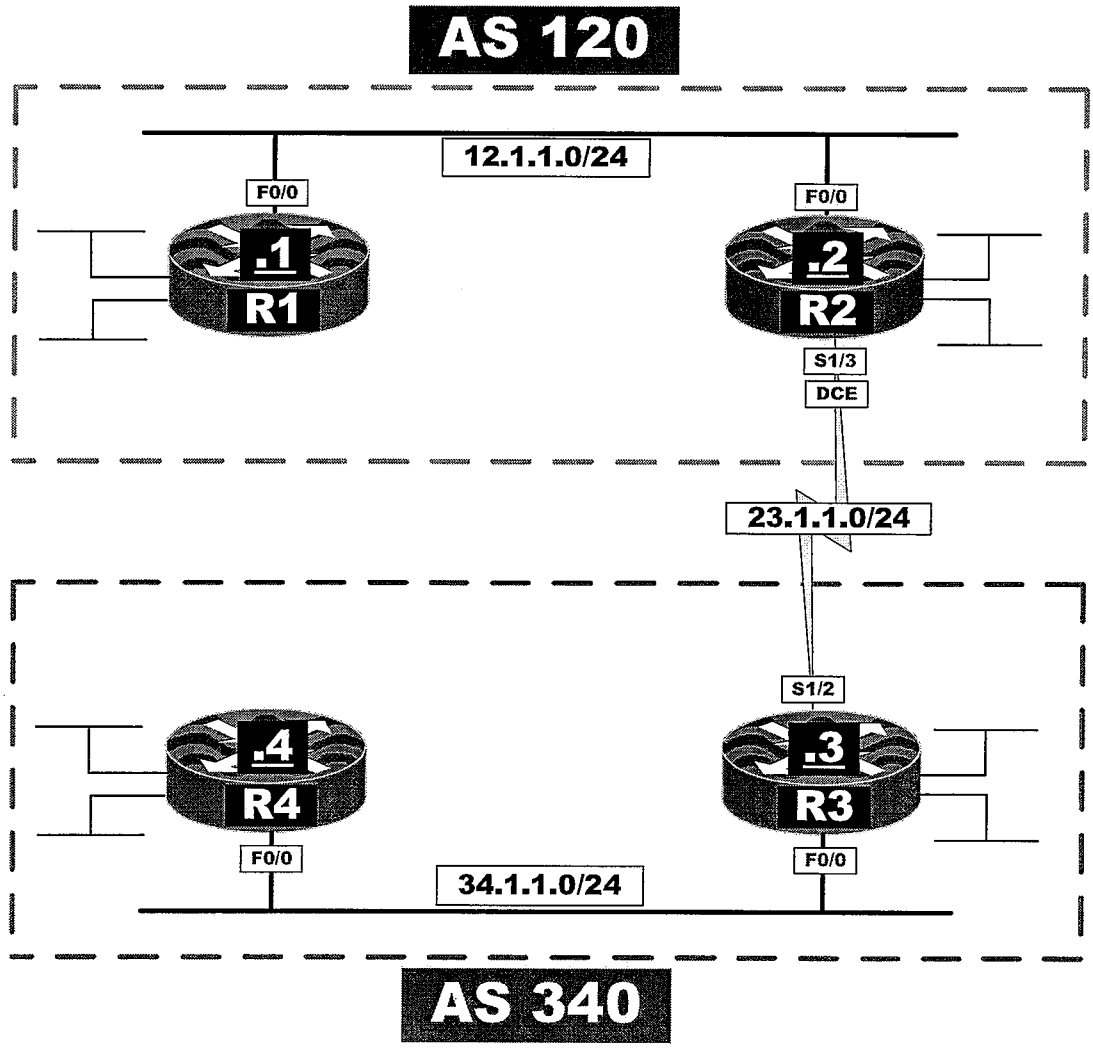
Task 18

Erase the startup configuration on all routers and switches, delete vlan.dat on all switches and reload all devices before proceeding to the next lab

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 15 - Administrative Distance



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-15".

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 1

Configure the routers based on the above diagram; these routers should ONLY advertise their Loopback interfaces in BGP, the peering between the routers should be established based on the following policy:

R1 in AS 120 should establish an IBGP peer session with R2 in AS 120.
R2 should establish an IBGP peer session with R1 in AS 120 and an EBGP peer session with R3 in AS 340.
R3 should establish an EBGP peer session with R2 in AS 120 and an IBGP peer session with R4 in AS 340.
R4 should establish an IBGP peer sessions with R3 in AS 340.

On R1:

```
R1(config)#Router bgp 120
R1(config-router)#No au
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Netw 11.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 120
```

On R2:

```
R2(config)#Router bgp 120
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Netw 22.0.0.0
R2(config-router)#No au
R2(config-router)#Neighbor 12.1.1.1 remote-as 120
R2(config-router)#Neighbor 12.1.1.1 next-hop-self
R2(config-router)#Neighbor 23.1.1.3 remote-as 340
```

On R3:

```
R3(config)#Router bgp 340
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Netw 33.0.0.0
R3(config-router)#Neighbor 23.1.1.2 remote-as 120
R3(config-router)#Neighbor 34.1.1.4 remote-as 340
R3(config-router)#Neighbor 34.1.1.4 next-hop-self
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#Router bgp 340
R4(config-router)#No au
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Netw 44.0.0.0
R4(config-router)#Neighbor 10.1.34.3 remote-as 340
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	0.0.0.0	0		32768	i
*>i2.0.0.0	12.1.1.2	0	100	0	i
*>i3.0.0.0	12.1.1.2	0	100	0	340 i
*>i4.0.0.0	12.1.1.2	0	100	0	340 i
*> 11.0.0.0	0.0.0.0	0		32768	i
*>i22.0.0.0	12.1.1.2	0	100	0	i
*>i33.0.0.0	12.1.1.2	0	100	0	340 i
*>i44.0.0.0	12.1.1.2	0	100	0	340 i

On R2:

```
R2#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i1.0.0.0	12.1.1.1	0	100	0	i
*> 2.0.0.0	0.0.0.0	0		32768	i
*> 3.0.0.0	23.1.1.3	0		0	340 i
*> 4.0.0.0	23.1.1.3			0	340 i
*>i11.0.0.0	12.1.1.1	0	100	0	i
*> 22.0.0.0	0.0.0.0	0		32768	i
*> 33.0.0.0	23.1.1.3	0		0	340 i
*> 44.0.0.0	23.1.1.3			0	340 i

On R3:

```
R3#Show ip bgp | b Network
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	23.1.1.2			0	120 i
*> 2.0.0.0	23.1.1.2	0		0	120 i
*> 3.0.0.0	0.0.0.0	0		32768	i

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*>i4.0.0.0          34.1.1.4                 0    100      0 i
*> 11.0.0.0         23.1.1.2                              0 120 i
*> 22.0.0.0         23.1.1.2                 0            0 120 i
*> 33.0.0.0         0.0.0.0                  0        32768 i
*>i44.0.0.0         34.1.1.4                 0    100      0 i
```

On R4:

```
R4#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          34.1.1.3                 0    100      0 120 i
*>i2.0.0.0          34.1.1.3                 0    100      0 120 i
*>i3.0.0.0          34.1.1.3                 0    100      0 i
*> 4.0.0.0          0.0.0.0                  0        32768 i
*>i11.0.0.0         34.1.1.3                 0    100      0 120 i
*>i22.0.0.0         34.1.1.3                 0    100      0 120 i
*>i33.0.0.0         34.1.1.3                 0    100      0 i
*> 44.0.0.0         0.0.0.0                  0        32768 i
```

Task 2

Configure R2 such that it changes the administrative distance of all prefixes received from R3 to 150.

In the following command the "23.1.1.3" is the neighbor's IP address followed by the inverse mask of 23.1.1.3. The distance command can be used to filter prefixes by setting the administrative distance to 255, prefer one path/Neighbor over another one, or setup primary and backup path for a given prefix.

On R2:

```
R2(config)#Router bgp 120
R2(config-router)#Distance 150 23.1.1.3 0.0.0.0
```

To verify the configuration:

On R2:

```
R2#Clear ip route *
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

B      1.0.0.0/8 [200/0] via 12.1.1.1, 00:00:24
B      3.0.0.0/8 [150/0] via 23.1.1.3, 00:00:24
B      4.0.0.0/8 [150/0] via 23.1.1.3, 00:00:24
B      11.0.0.0/8 [200/0] via 12.1.1.1, 00:00:24
B      33.0.0.0/8 [150/0] via 23.1.1.3, 00:00:24
B      44.0.0.0/8 [150/0] via 23.1.1.3, 00:00:24
```

Task 3

Remove the configuration commands from the previous task before proceeding to the next task.

On R2:

```
R2(config)#Router bgp 120
R2(config-router)#No Distance 23.1.1.3 0.0.0.0
```

To verify the configuration

On R2:

```
R2#Clear ip route *

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

B      1.0.0.0/8 [200/0] via 12.1.1.1, 00:02:22
B      3.0.0.0/8 [20/0] via 23.1.1.3, 00:02:22
B      4.0.0.0/8 [20/0] via 23.1.1.3, 00:02:22
B      11.0.0.0/8 [200/0] via 12.1.1.1, 00:02:22
B      33.0.0.0/8 [20/0] via 23.1.1.3, 00:02:22
B      44.0.0.0/8 [20/0] via 23.1.1.3, 00:02:22
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 4

Configure R2 such that it **only** changes the administrative distance of prefix 33.0.0.0 /8 to 150.

On R2:

```
R2(config)#Access-list 33 permit 33.0.0.0

R2(config)#Router bgp 120
R2(config-router)#Distance 150 23.1.1.3 0.0.0.0 33
```

With the above configuration we are saying "Permit network 33.0.0.0/8 coming from 23.1.1.3 to have an administrative distance of 150".

To verify the configuration:

On R2:

```
R2#Clear ip route *

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

B      1.0.0.0/8 [200/0] via 12.1.1.1, 00:00:10
B      3.0.0.0/8 [20/0] via 23.1.1.3, 00:00:10
B      4.0.0.0/8 [20/0] via 23.1.1.3, 00:00:10
B      11.0.0.0/8 [200/0] via 12.1.1.1, 00:00:10
B      33.0.0.0/8 [150/0] via 23.1.1.3, 00:00:10
B      44.0.0.0/8 [20/0] via 23.1.1.3, 00:00:10
```

Task 5

Using minimum number of commands change the administrative distance of all IBGP prefixes to 90, EBGP prefixes to 60 and locally generated prefixes to 20. This should be performed on the routers in AS 340.

On R3 and R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

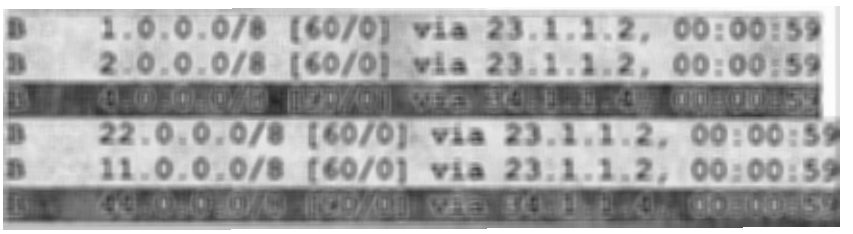
```
Rx(config)#Router bgp 340
Rx(config-router)#Distance bgp 60 90 20

Rx#Clear ip route *
```

To verify the configuration:

On R3:

```
R3#Show ip route bgp
```



The output of the above show command verifies the administrative distance of IBGP prefixes (60) and EBGP prefixes (90). To see the administrative distance of self generated prefix/es, Loopback 100 – 103 inerfaces are configured and then aggregated:

On R3:

```
R3(config)#Int lo100
R3(config-if)#Ip address 100.1.0.3 255.255.255.0

R3(config)#Int lo101
R3(config-if)#Ip address 100.1.1.3 255.255.255.0

R3(config)#Int lo102
R3(config-if)#Ip address 100.1.2.3 255.255.255.0

R3(config)#Int lo103
R3(config-if)#Ip address 100.1.3.3 255.255.255.0

R3(config-if)#Router bgp 340
R3(config-router)#Netw 100.1.0.0 mask 255.255.255.0
R3(config-router)#Netw 100.1.1.0 mask 255.255.255.0
R3(config-router)#Netw 100.1.2.0 mask 255.255.255.0
R3(config-router)#Netw 100.1.3.0 mask 255.255.255.0
R3(config-router)#Aggregate-address 100.1.0.0 255.255.252.0 summary-only
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R3:

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

B      1.0.0.0/8 [60/0] via 23.1.1.2, 00:14:53
B      2.0.0.0/8 [60/0] via 23.1.1.2, 00:14:53
B      4.0.0.0/8 [90/0] via 34.1.1.4, 00:14:53
B      11.0.0.0/8 [60/0] via 23.1.1.2, 00:14:53
B      22.0.0.0/8 [60/0] via 23.1.1.2, 00:14:53
B      44.0.0.0/8 [90/0] via 34.1.1.4, 00:14:53
       100.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
B         100.1.0.0/22 [20/0] via 0.0.0.0, 00:00:22, Null0
```

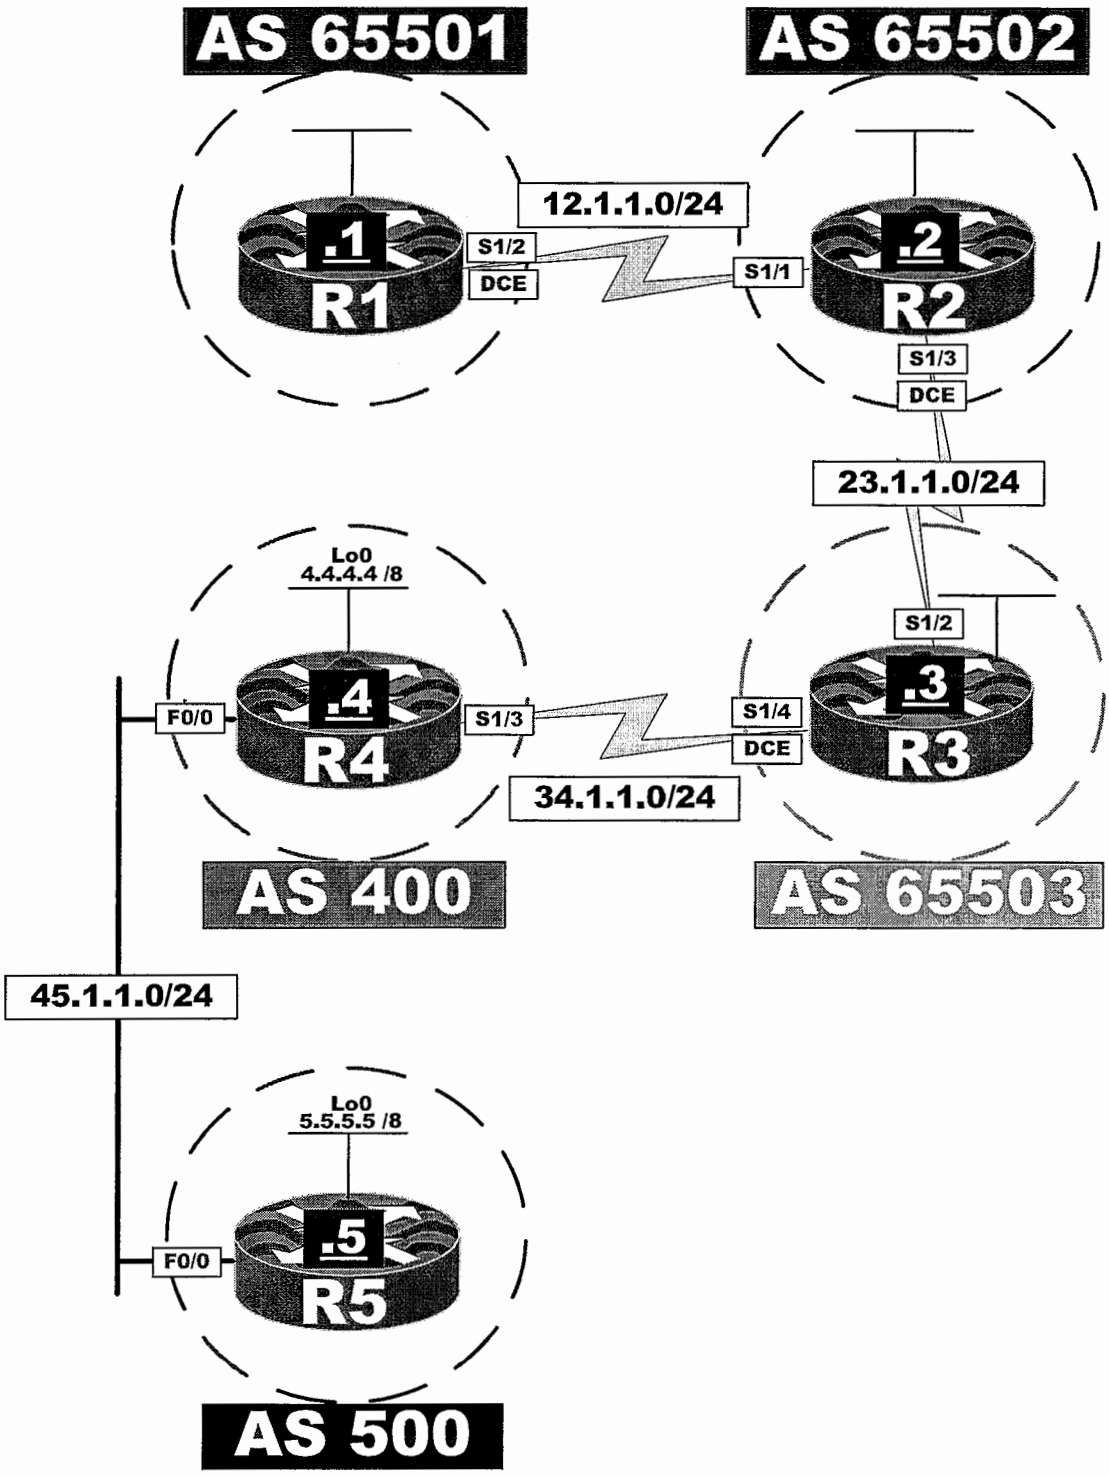
Task 6

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 16 - Advanced BGP Configurations



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-16".

Task 1

Configure BGP peering on the routers as follows:

R1 in AS 65501, should establish an EBGP peer session with R2 in AS 65502.
R2 in AS 65502, should establish EBGP peer sessions with R1 and R3 in AS 65501 and 65503 respectively.
R3 in AS 65503, should establish EBGP peer sessions with R2 and R4 in AS 65502 and AS 400 respectively.
R4 in AS 400, should establish EBGP peer sessions with R3 and R5 in AS 65503 and AS 500 respectively.
R5 in AS 500 should establish an EBGP peer session with R4 in AS 400.
These routers should advertise their loopback interface/s in their assigned AS.

On R1:

```
R1(config)#Router bgp 65501
R1(config-router)#No au
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 65502
```

On R2:

```
R2(config)#Router bgp 65502
R2(config-router)#No au
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 65501
R2(config-router)#Neighbor 23.1.1.3 remote-as 65503
```

On R3:

```
R3(config)#Router bgp 65503
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Neighbor 23.1.1.2 remote-as 65502
R3(config-router)#Neighbor 34.1.1.4 remote-as 400
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#No au
R4(config-router)#Netw 4.0.0.0
R4(config-router)#Neighbor 34.1.1.3 remote-as 65503
R4(config-router)#Neighbor 45.1.1.5 remote-as 500
```

On R5:

```
R5(config-if)#Router bgp 500
R5(config-router)#No au
R5(config-router)#Netw 5.0.0.0
R5(config-router)#Neighbor 45.1.1.4 remote-as 400
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 65502 i
*> 3.0.0.0          12.1.1.2                               0 65502 65503 i
*> 4.0.0.0          12.1.1.2                               0 65502 65503 400 i
*> 5.0.0.0          12.1.1.2                               0 65502 65503 400 500 i
```

Task 2

Configure R2 such that if any of its links to any of its external peers go down, the routes from that particular neighbor are **not** removed from the BGP table. The routes should only be removed if a given neighbor is down for longer than 2 minutes.

On R2:

```
R2(config-if)#Router bgp 65502
R2(config-router)#No bgp fast-external-fallover
R2(config-router)#Neighbor 12.1.1.1 timers 60 120
R2(config-router)#Neighbor 23.1.1.3 timers 60 120
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1(config-if)#Router bgp 65501
R1(config-router)#Neighbor 12.1.1.2 timers 60 120
```

On R3:

```
R3(config-if)#Router bgp 65503
R3(config-router)#Neighbor 23.1.1.2 timers 60 120
```

This feature only supports the directly connected EXTERNAL peers. If "BGP fast external-fallover" is disabled, the BGP routing process will wait until the configured hold timer expires before the peer session is reset. If the other neighbors are not configured with the timers, the lowest value will be negotiated. The default value is 180 seconds.

To verify and test the configuration:

On R2:

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 65501 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          23.1.1.3                 0             0 65503 i
*> 4.0.0.0          23.1.1.3                               0 65503 400 i
*> 5.0.0.0          23.1.1.3                               0 65503 400 500 i
```

Let's shutdown R2's link to R1:

On R2:

```
R2(config)#Int S1/1
R2(config-if)#Shut
```

You should see the following console message stating that the link is down:

```
*Feb  2 02:55:42.435: %LINK-5-CHANGED: Interface Serial1/1, changed state
to administratively down

*Feb  2 02:55:43.435: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Serial1/1, changed state to down

R2#Show ip bgp | B Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          12.1.1.1                 0             0 65501 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          23.1.1.3                 0             0 65503 i
*> 4.0.0.0          23.1.1.3                               0 65503 400 i
*> 5.0.0.0          23.1.1.3                               0 65503 400 500 i
```

We can see that R1 still holds network 1.0.0.0/8 in its BGP table, the greater than sign (>) is no longer there and it's because the local router (R2)does NOT have NLRI to the next hop IP address because the link to R1 is down. The following show command displays that R2 still recognizes R1 as a neighbor:

```
R2#Show ip bgp summary | B Neigh

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
12.1.1.1        4        65501      26      31        1    0    0 00:20:57        1
23.1.1.3        4        65503      29      35        8    0    0 00:20:02        3
```

After two minutes you should see the following console message:

```
*Feb  2 02:57:57.299: %BGP-5-ADJCHANGE: neighbor 12.1.1.1 Down BGP
Notification sent

*Feb  2 02:57:57.299: %BGP-3-NOTIFICATION: sent to neighbor 12.1.1.1 4/0
(hold time expired) 0 bytes

R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          23.1.1.3                 0             0 65503 i
*> 4.0.0.0          23.1.1.3                               0 65503 400 i
*> 5.0.0.0          23.1.1.3                               0 65503 400 500 i
```

Let's "No shutdown" the S1/1 interface of R2 and proceed to the next task.

On R2:

```
R2(config)#Int S1/1
R2(config-if)#No shut
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 3

Configure R5 in AS 500 so it sets the hello and hold timer values to double its default value for its neighboring router R4.

On R5:

```
R5(config)#Router bgp 500
R5(config-router)#Neighbor 45.1.1.4 timer 120 360
```

The default hello interval is set to 60 seconds and the hold timer is set to 180 seconds.

Task 4

Configure R4 such that it removes the private AS numbers when it advertises prefixes to R5.

Before configuring R4, we should display the BGP table of R5.

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          45.1.1.4                               0 400 65503 65502 65501 i
*> 2.0.0.0          45.1.1.4                               0 400 65503 65502 i
*> 3.0.0.0          45.1.1.4                               0 400 65503 i
*> 4.0.0.0          45.1.1.4                 0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

When a customer network is large, they may require to have a registered AS number, this can be done in two different ways:

One way to accomplish this task is to acquire an AS number in the range of 1 – 64511. This is done when the customer network connects to two different ISPs, in a multihoming environment.

The second method is done by assigning a private AS number which are the last 1024 AS numbers, or 64512 – 65535, this is done when the customer network connects to a single ISP or is multihomed to the same ISP and the reason this happens is because of the AS number conservation.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

When a Private AS number is used by the customer, the BGP updates from that given customer will include the Private AS number in its AS-Path list, and when the ISP advertises the customer network/s to its peers it should remove the private AS numbers, or else the updates will include the private AS numbers. The following command is one way to remove the private AS number/s.

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#Neighbor 45.1.1.5 remove-private-as

R4#Clear ip bgp * out
```

To verify the configuration:

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          45.1.1.4                               0 400 i
*> 2.0.0.0          45.1.1.4                               0 400 i
*> 3.0.0.0          45.1.1.4                               0 400 i
*> 4.0.0.0          45.1.1.4                 0             0 400 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

Task 5

Configure R1 such that it sets the minimum time between sending BGP advertisement for its EBGP neighbors to 60 seconds.

On R1:

```
R1(config)#Router bgp 65501
R1(config-router)#Neighbor 12.1.1.2 advertisement-interval 60
```

The default values are as follows:

- External peers – 30 seconds.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

- **Internal peers – 5 seconds.**

This command sets the minimum interval between the sending of BGP routing updates for internal and external neighbors.

Task 6

Configure the following loopback1 interfaces on R4 and R5:

R4 – Lo1 – 44.4.4.4/8
R5 – Lo1 – 55.5.5.5/8

Ensure that R4 and R5 establish their EBGP peer session based on their Lo1 interface. **Do not** use "**neighbor ebgp multihop**" command to establish this peer session. You should use RIPv2 to accomplish this task.

On R4:

```
R4(config)#Int lo1
R4(config-if)#Ip addr 44.4.4.4 255.0.0.0

R4(config)#Router rip
R4(config-router)#No au
R4(config-router)#Ver 2
R4(config-router)#Network 44.0.0.0
R4(config-router)#Network 45.0.0.0
```

On R5:

```
R5(config)#Int lo1
R5(config-if)#Ip addr 55.5.5.5 255.0.0.0

R5(config)#Router rip
R5(config-router)#No au
R5(config-router)#Ver 2
R5(config-router)#Network 55.0.0.0
R5(config-router)#Network 45.0.0.0
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config-if)#Router bgp 400
R4(config-router)#No neighbor 45.1.1.5 remote-as 500
R4(config-router)#Neighbor 55.5.5.5 remote-as 500
R4(config-router)#Neighbor 55.5.5.5 update-source lo1
R4(config-router)#Neighbor 55.5.5.5 disable-connected-check
```

On R5:

```
R5(config-if)#Router bgp 500
R5(config-router)#No neighbor 45.1.1.4 remote-as 400
R5(config-router)#Neighbor 44.4.4.4 remote-as 400
R5(config-router)#Neighbor 44.4.4.4 update-source lo1
R5(config-router)#Neighbor 44.4.4.4 disable-connected-check
```

The default behavior of BGP:

A BGP speaker will verify the connection of a single-hop EBGP peering session to determine if the EBGP peer is directly connected to the same network segment. If the peer is not directly connected to the same network segment, the connection verification will fail and it will prevent the peering session from being established.

The "neighbor disable-connected-check" command will disable the connection verification process for EBGP peering session that are reachable by a single hop but are configured on a loopback interface or configured with a non-directly connected IP address.

This solution is typically used when the two peers try to mitigate man in the middle attacks.

Task 7

Configure the following loopback1 interfaces on R1 and R2:

R1 – Lo1 – 11.1.1.1/8
R2 – Lo1 – 22.2.2.2/8

Ensure that R1 and R2 establish their EBGP peer session based on their Lo1 interface. **Do not** use "**neighbor ebgp multihop**"or "**neighbor disable-connected-check**" command to establish this peer session. You should use RIPv2 to accomplish this task.

On R1:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config)#Int lo1
R1(config-if)#IP addr 11.1.1.1 255.0.0.0

R1(config-if)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 11.0.0.0
R1(config-router)#Network 12.0.0.0

R1(config)#Router bgp 65501
R1(config-router)#Neigh 22.2.2.2 remote-as 65502
R1(config-router)#Neigh 22.2.2.2 update-source Lo1
R1(config-router)#Neigh 22.2.2.2 ttl-security hops 2
```

On R2:

```
R2(config)#Int Lo1
R2(config-if)#IP addr 22.2.2.2 255.0.0.0

R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 22.0.0.0
R2(config-router)#Network 12.0.0.0

R2(config)#Router bgp 65502
R2(config-router)#Neigh 11.1.1.1 remote-as 65501
R2(config-router)#Neigh 11.1.1.1 update-source lo1
R2(config-router)#Neigh 11.1.1.1 ttl-security hops 2
```

This feature is enabled in 12.3(7)T and it provides a lightweight security for the BGP routers against CPU utilization attacks. These types of attacks flood the network with IP packets that contain forged source and destination IP addresses in the packet headers. This configuration accepts only IP packets with a TTL count that is equal to or greater than the locally configured value.

Task 8

Configure R3 to inject a default route to neighbor 34.1.1.4 ONLY if there is a route to 2.0.0.0 /8 in R3's routing table.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3(config)#Access-list 2 permit 2.0.0.0

R3(config)#Route-map TEST permit 10
R3(config-route-map)#match ip addr 2

R3(config)#Router bgp 65503
R3(config-router)#Neighbor 34.1.1.4 default-originate route-map TEST

R3#Clear ip bgp *
```

To verify the configuration:

On R4

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          34.1.1.3                                0 65503 i
*> 1.0.0.0          34.1.1.3                                0 65503 65502 65501 i
*> 2.0.0.0          34.1.1.3                                0 65503 65502 i
*> 3.0.0.0          34.1.1.3                 0              0 65503 i
*> 4.0.0.0          0.0.0.0                  0          32768 i
*> 5.0.0.0          55.5.5.5                 0              0 500 i
```

To test the configuration:

On R2

```
R2(config)#Int lo0
R2(config-if)#Shut
```

NOTE: Network 2.0.0.0/8 is no longer in the BGP table of R3:

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
   0.0.0.0          0.0.0.0                                 0 i
*> 1.0.0.0          23.1.1.2                                0 65502 65501 i
*> 3.0.0.0          0.0.0.0                  0          32768 i
*> 4.0.0.0          34.1.1.4                 0              0 400 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 5.0.0.0          34.1.1.4                               0 400 500 i
```

Since network 2.0.0.0/8 is no longer in the BGP table of R2, R2 will NOT inject a default route.

On R4:

```
R4#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          34.1.1.3                               0 65503 65502 65501 i
*> 3.0.0.0          34.1.1.3                 0             0 65503 i
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          55.5.5.5                 0             0 500 i
```

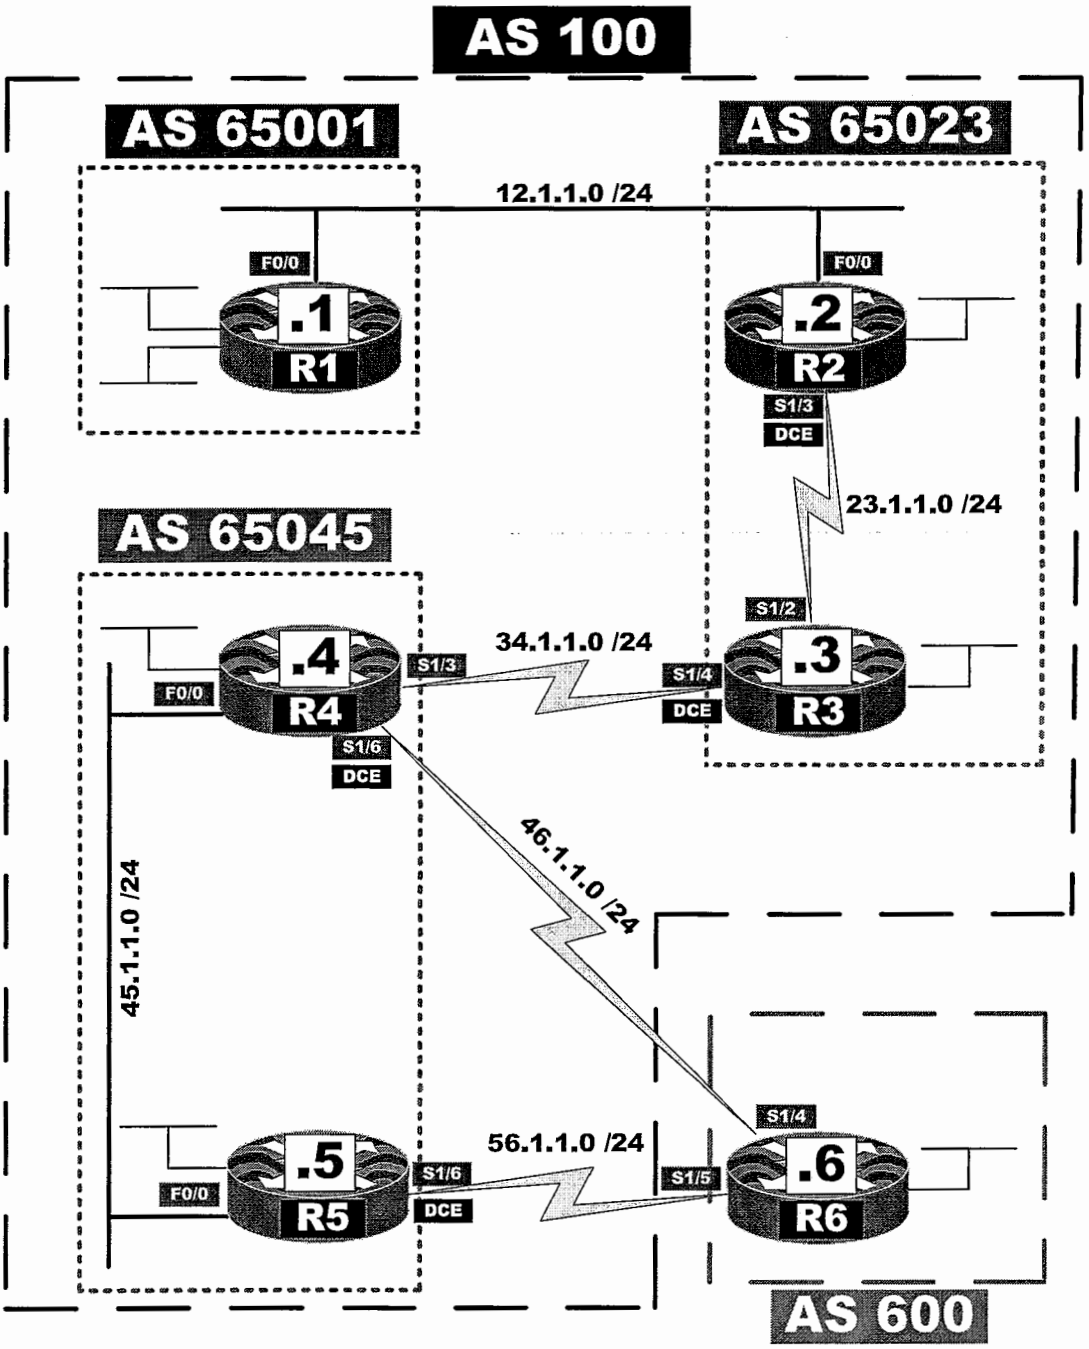
Task 11

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 17 - BGP Confederation



© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-17".

Task 1

Configure BGP peering on the routers as follows:

- R1 in AS 65001, should establish an EBGP peer session with R2 in AS 65023.
- R2 in AS 65023, should establish an IBGP peer session with R3.
- R3 in AS 65023, should establish an EBGP peer sessions with R4 in AS 65045.
- R4 in AS 65045, should establish an IBGP peer sessions with R5.
- R4 and R5 should establish an EBGP peer session with R6 in AS 600.
- Provide NLRI to the links in AS 100 using RIPv2.
- These routers should advertise their Loopback interface/s in BGP.
- Routers R1, R2, R3, R4 and R5 should be configured in AS 100.

The implementation of BGP confederation reduces the iBGP mesh inside an AS. The trick is to divide an AS into multiple ASes and assign the whole group to a single confederation. Within each Sub-AS the routers are fully meshed and they have connection/s to other ASes inside the confederation. Even though these ASes have an EBGP peers to ASes within the confederation, the ASes exchange routing as if they used IBGP. In this way, the confederation preserves next hop, metric, and local preference information. To the outside world, the confederation appears to be a single AS.

On R1:

```
R1(config)#Router bgp 65001
R1(config-router)#No au
R1(config-router)#Network 1.0.0.0
R1(config-router)#Network 11.0.0.0
```

The following command is used to identify the real AS. This command must be configured on all the routers that are in the same confederation. The "BGP confederation identifier" command is used to configure a single AS number to identify a group of smaller ASes (Sub-ASes) as a single confederation.

```
R1(config-router)#bgp confederation identifier 100
```

The following command is used to identify the directly connected EBGP sub-confed peers. If the local router does NOT have an EBGP Sub-confed peer/s then this command is NOT required and should NOT be configured.

```
R1(config-router)#bgp confederation peers 65023
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

We can see that within the confederation the routers establish a peer session using the private AS numbers:

```
R1(config-router)#Neighbor 12.1.1.2 remote-as 65023

R1(config-router)#Router rip
R1(config-router)#No au
R1(config-router)#Ver 2
R1(config-router)#Network 12.0.0.0
```

On R2:

```
R2(config)#Router bgp 65023
R2(config-router)#No au
R2(config-router)#Network 2.0.0.0
R2(config-router)#Bgp confederation identifier 100
R2(config-router)#Bgp confederation peers 65001
R2(config-router)#Neighbor 12.1.1.1 remote-as 65001
R2(config-router)#Neighbor 23.1.1.3 remote-as 65023

R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Network 12.0.0.0
R2(config-router)#Network 23.0.0.0
```

On R3:

```
R3(config)#Router bgp 65023
R3(config-router)#No au
R3(config-router)#Network 3.0.0.0
R3(config-router)#Bgp confederation identifier 100
R3(config-router)#Bgp confederation peers 65045
R3(config-router)#Neighbor 23.1.1.2 remote-as 65023
R3(config-router)#Neighbor 34.1.1.4 remote-as 65045

R3(config)#Router rip
R3(config-router)#No au
R3(config-router)#Ver 2
R3(config-router)#Network 34.0.0.0
R3(config-router)#Network 23.0.0.0
```

On R4:

```
R4(config)#Router bgp 65045
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R4(config-router)#No au
R4(config-router)#Network 4.0.0.0
R4(config-router)#Bgp confederation identifier 100
R4(config-router)#Bgp confederation peers 65023
R4(config-router)#Neighbor 45.1.1.5 remote-as 65045
R4(config-router)#Neighbor 34.1.1.3 remote-as 65023
R4(config-router)#Neighbor 46.1.1.6 remote-as 600

R4(config)#Router rip
R4(config-router)#No au
R4(config-router)#Ver 2
R4(config-router)#Network 45.0.0.0
R4(config-router)#Network 34.0.0.0
R4(config-router)#Network 46.0.0.0
```

The following command must be configured so the EBGP peer through this link can NOT see your internal networks. IF the following command was NOT configured and R6 runs RIPv2 on its S1/4 interface, it will start exchanging routes with R4. Another way to accomplish this task is to run RIP Authentication, but the passive-interface is a better choice.

```
R4(config-router)#Passive-interface S1/6
```

On R5:

Since R5 does not have any EBGP Sub-confed peer/s; it is NOT configured with the "bgp confederation peers" command.

```
R5(config)#Router bgp 65045
R5(config-router)#No au
R5(config-router)#Network 5.0.0.0
R5(config-router)#Bgp confederation identifier 100
R5(config-router)#Neighbor 56.1.1.6 remote-as 600
R5(config-router)#Neighbor 45.1.1.4 remote-as 65045

R5(config)#Router rip
R5(config-router)#No au
R5(config-router)#Ver 2
R5(config-router)#Network 45.0.0.0
R5(config-router)#Network 56.0.0.0
R5(config-router)#Passive-interface S1/6
```

On R6:

```
R6(config)#Router bgp 600
R6(config-router)#No au
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6(config-router)#Network 6.0.0.0
R6(config-router)#Neighbor 56.1.1.5 remote-as 100
R6(config-router)#Neighbor 46.1.1.4 remote-as 100
```

Note: From R6's perspective, R4 and R5 are in AS 100 and NOT 65045; AS 65045 is ONLY known within the confederation and NOT to any EBGP peer outside of the confederation.

On R5:

```
R5#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          12.1.1.1                 0    100      0 (65023 65001) i
*>i2.0.0.0          23.1.1.2                 0    100      0 (65023) i
*>i3.0.0.0          34.1.1.3                 0    100      0 (65023) i
*>i4.0.0.0          45.1.1.4                 0    100      0 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
* i6.0.0.0          46.1.1.6                 0    100      0 600 i
*>                  56.1.1.6                 0             0 600 i
*>i11.0.0.0         12.1.1.1                 0    100      0 (65023 65001) i
```

Note: Since R5 received all the prefixes from the other Sub-ASes through an IBGP Sub-confed, the origin code remained as IGP, this is the letter (i) in the Network colume. The AS-Path/s in the parenthesis are the private AS numbers within the confederation, ASes outside of the confederation do not see the private AS numbers, they ONLY see the real AS number. We can see that AS 600 is NOT part of this confederation.

On R6:

```
R6#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          56.1.1.5                               0 100 i
*>                  46.1.1.4                               0 100 i
*  2.0.0.0          56.1.1.5                               0 100 i
*>                  46.1.1.4                               0 100 i
*  3.0.0.0          56.1.1.5                               0 100 i
*>                  46.1.1.4                               0 100 i
*  4.0.0.0          56.1.1.5                               0 100 i
*>                  46.1.1.4                 0             0 100 i
*  5.0.0.0          56.1.1.5                 0             0 100 i
*>                  46.1.1.4                               0 100 i
*> 6.0.0.0          0.0.0.0                  0         32768 i
*  11.0.0.0         56.1.1.5                               0 100 i
*>                  46.1.1.4                               0 100 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

We can see that R6 in AS 600 sees all the routes in AS 100 and the private AS numbers are all automatically replaced with the real AS, AS 100.

On R1:

```
R1#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0    100      0 (65023) i
*> 3.0.0.0          23.1.1.3                 0    100      0 (65023) i
*> 4.0.0.0          34.1.1.4                 0    100      0 (65023 65045) i
*> 5.0.0.0          45.1.1.5                 0    100      0 (65023 65045) i
*> 6.0.0.0          46.1.1.6                 0    100      0 (65023 65045) 600 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

By looking at network 6.0.0.0 we can see that the next-hop IP address did NOT change, even though it traversed through many different AS numbers (65045, and 65023), when it comes to the next-hop attribute within the confederation, the next-hop attribute treats the confederation as one.

The next-hop attribute reads the "bgp confederation identifier" and treats the entire confederation as one AS, the AS specified in the "bgp confederation identifier" command.

Task 2

Change the default local-preference attribute on R5 to 500, and R4 to 400.

On R4:

```
R4(config)#Router bgp 65045
R4(config-router)#BGP Default local-preference 400

R4#Clear ip bgp * in
```

On R5:

```
R5(config)#Router bgp 65045
R5(config-router)#BGP Default local-preference 500

R5#Clear ip bgp * in
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0    100      0 (65023) i
*> 3.0.0.0          23.1.1.3                 0    100      0 (65023) i
*> 4.0.0.0          34.1.1.4                 0    400      0 (65023 65045) i
*> 5.0.0.0          45.1.1.5                 0    500      0 (65023 65045) i
*> 6.0.0.0          56.1.1.6                 0    500      0 (65023 65045) 600 i
*> 11.0.0.0         0.0.0.0                  0         32768 i
```

To test the configuration:

On R4:

```
R4(config)#Router bgp 65045
R4(config-router)#Neighbor 46.1.1.6 default-originate

R4#Clear ip bgp *
```

On R5:

```
R5(config)#Router bgp 65045
R5(config-router)#Neighbor 56.1.1.6 default-originate

R5#Clear ip bgp *
```

Wait for both neighbors to come up before you traceroute.

On R1:

```
R1#Traceroute 6.6.6.6

Type escape sequence to abort.
Tracing the route to 6.6.6.6
VRF info: (vrf in name/id, vrf out name/id)
  1 12.1.1.2 0 msec 4 msec 0 msec
  2 23.1.1.3 16 msec 16 msec 16 msec
  3 34.1.1.4 28 msec 28 msec 32 msec
  4 45.1.1.5 32 msec 28 msec 32 msec
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
  5 56.1.1.6 40 msec *  40 msec
```

We can see that the local-preference leaks through the entire confederation.

Task 3

Configure R6 such that when it advertises its network 6.0.0.0/8 to routers R4 and R5 in AS 65045, the routers in AS 65045 do not advertise or export this network to any of their EBGP peer/s.

Well, this is an easy task to configure, let's configure the community attribute of "No-Export" and test:

On R6:

```
R6(config)#IP Prefix-list NET6 permit 6.0.0.0/8

R6(config)#Route-map tst permit 10
R6(config-route-map)#Match ip addr prefix NET6
R6(config-route-map)#Set community no-export
R6(config)#Route-map tst permit 90

R6(config)#Router bgp 600
R6(config-router)#Neighbor 46.1.1.4 route-map tst out
R6(config-router)#Neighbor 56.1.1.5 route-map tst out
R6(config-router)#Neighbor 56.1.1.5 Send-community
R6(config-router)#Neighbor 46.1.1.4 Send-community

R6#Clear ip bgp *
```

To verify the configuration:

On R4:

```
R4#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 26
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     12          13
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  600
    56.1.1.6 (metric 1) from 45.1.1.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 500, valid, confed-internal, best
  600
    46.1.1.6 from 46.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 400, valid, external
      Community: no-export
```

On R5:

```
R5#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 21
Paths: (1 available, best #1, table default, not advertised to EBGP peer)
  Advertised to update-groups:
     6
  600
    56.1.1.6 from 56.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 500, valid, external, best
      Community: no-export
```

We can see that the "No-export" community attribute is assigned to the prefix. Let's check this on R3:

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          12.1.1.1                 0    100      0 (65001) i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          34.1.1.4                 0    400      0 (65045) i
*> 5.0.0.0          45.1.1.5                 0    500      0 (65045) i
*> 6.0.0.0          56.1.1.6                 0    500      0 (65045) 600 i
*>i11.0.0.0         12.1.1.1                 0    100      0 (65001) i
```

R3 does NOT see the "No-export" attribute, but why did R4 advertise the route to R3 to begin with?

The reason is that the "No-export" attribute leaks through the entire confederation and therefore, it does NOT work. So what is the solution?
The solution is the "Local-as" community attribute.

Let's configure the "Local-as" and test it:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R6:

```
R6(config)#Route-map tst permit 10
R6(config-route-map)#No set community no-export
R6(config-route-map)#set community local-as

R6#Clear ip bgp * out
```

On R5:

```
R5#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 22
Paths: (1 available, best #1, table default, not advertised outside local
AS)
  Advertised to update-groups:
     6
  600
    56.1.1.6 from 56.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 500, valid, external, best
      Community: local-AS
```

On R4:

```
R4#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 26
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     12          13
  600
    56.1.1.6 (metric 1) from 45.1.1.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 500, valid, confed-internal, best
  600
    46.1.1.6 from 46.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 400, valid, external
      Community: local-AS
```

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          12.1.1.1                 0    100      0 (65001) i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*>i2.0.0.0              23.1.1.2                 0    100      0 i
*> 3.0.0.0              0.0.0.0                  0         32768 i
*> 4.0.0.0              34.1.1.4                 0    400      0 (65045) i
*> 5.0.0.0              45.1.1.5                 0    500      0 (65045) i
*> 6.0.0.0              56.1.1.6                 0    500      0 (65045) 600 i
*>i11.0.0.0             12.1.1.1                 0    100      0 (65001) i
```

Question:
Why is R3 getting the route?

Answer:
What is the BEST path from R4's perspective? Let's check:

On R4:

```
R4#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 16
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     10          13
  600
    56.1.1.6 (metric 1) from 45.1.1.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 500, valid, confed-internal, best
  600
    46.1.1.6 from 46.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 400, valid, external
      Community: local-AS
```

We can see R4 has two paths to 6.0.0.0/8, but the best path is through R5, BUT WHY?

Remember the local preference from the previous task? R5's default local-preference was changed to 500, and R4's default local-preference was changed to 400, so R4 and the rest of the routers in AS 100 will traverse through R5 to connect to network 6.0.0.0/8.
So in this case we need to configure R5 to send community to R4, let's configure this and verify and test the result:

On R5:

```
R5(config)#Router bgp 65045
R5(config-router)#Neighbor 45.1.1.4 send-community

R5#Clear ip bgp *
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

To verify the configuration:

On R4:

```
R4#Show ip bgp 6.0.0.0

BGP routing table entry for 6.0.0.0/8, version 21
Paths: (2 available, best #1, table default, not advertised outside local
AS)
  Not advertised to any peer
  600
    56.1.1.6 (metric 1) from 45.1.1.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 500, valid, confed-internal, best
      Community: local-AS
  600
    46.1.1.6 from 46.1.1.6 (6.6.6.6)
      Origin IGP, metric 0, localpref 400, valid, external
      Community: local-AS
```

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.0.0.0          12.1.1.1                 0    100      0 (65001) i
*>i2.0.0.0          23.1.1.2                 0    100      0 i
*> 3.0.0.0          0.0.0.0                  0         32768 i
*> 4.0.0.0          34.1.1.4                 0    400      0 (65045) i
*> 5.0.0.0          45.1.1.5                 0    500      0 (65045) i
*>i11.0.0.0         12.1.1.1                 0    100      0 (65001) i
```

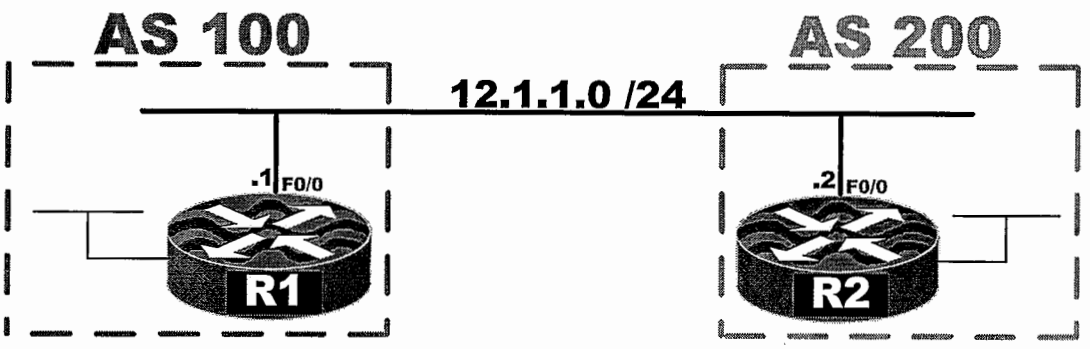
Task 4

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 18 – BGP Hiding Local AS number



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-18".

Task 1

Configure R1 in AS 100 to establish an EBGP session with R2 in AS 200. These routers should advertise their loopback 0 interface. Disable auto summarization.

On R1:

```
R1(config)#Router bgp 100
R1(config-router)#No auto-summary
R1(config-router)#Network 1.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No auto-summary
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R2(config-router)#Network 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 100
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note: From R2's perspective prefix 1.0.0.0/8 was originated & advertised by AS 100, and from R1's perspective, prefix 2.0.0.0/8 was originated and advertised by AS 200.

Task 2

Configure R1 in AS 111 to establish an EBGP session with R2 in AS 200 such that the output of the "**Show ip bgp**" command on these two routers will be identical to the follows:

On R1:

```
R1#Sh ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 100 200 i
```

On R2:

```
R2#Sh ip bgp | b Network
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1:

```
R1(config)#No router bgp 100

R1(config)#Router bgp 111
R1(config-router)#No auto-summary
R1(config-router)#Network 1.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Neighbor 12.1.1.2 local-as 100
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 12.1.1.2 Up
```

The "Neighbor 12.1.1.2 local-as 100" command on R1 tells R2 that it's in AS 100 even though R1 is configured with "Router bgp 111"; R2 will see R1's real AS of "111" originating the route and then AS 100 was the AS that advertised it.

NOTE: On R1, it shows that prefix 2.0.0.0/8 was originated by AS 200 but the advertising AS to R1 was AS 100. They both see this invisible AS (AS 100).

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 100 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Both routers see this fake AS 100.

Task 3

Configure R1 such that when R2 advertises network 2.0.0.0/8, the output of the "**Show ip bgp**" command on R1 resembles to the following. DO NOT remove any command/s to accomplish this task.

On R1:

```
R1#Sh ip bgp b Network
```

```
     Network          Next Hop            Metric LocPrf Weight Path
 *> 1.0.0.0          0.0.0.0                  0         32768 i
 *> 2.0.0.0          12.1.1.2                 0             0 200 i
```

On R1:

```
R1(config)#Router bgp 111
R1(config-router)#Neighbor 12.1.1.2 local-as 100 no-prepend
```

Note: The "No-prepend" keyword tells the router not to prepend the fake AS number (AS 100) to the advertised prefixes. This will only affect R1's BGP table.

On R1:

```
R1#Show ip bgp | b Network
```

```
     Network          Next Hop            Metric LocPrf Weight Path
 *> 1.0.0.0          0.0.0.0                  0         32768 i
 *> 2.0.0.0          12.1.1.2                 0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network
```

```
     Network          Next Hop            Metric LocPrf Weight Path
 *> 1.0.0.0          12.1.1.1                 0             0 100 111 i
 *> 2.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Note: The above configuration only affected R1.

Task 4

Configure R1 such that the output of the "**Show ip bgp**" command on R2 is identical to the following. DO NOT remove any command/s to accomplish this task.

On R2:

```
R2#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1:

```
R1(config)#Router bgp 111
R1(config-router)#Neighbor 12.1.1.2 local-as 100 no-prepend replace-as
```

Note: The "replace-as" keyword instructs the local router not to prepend the real AS number.

On R2:

```
R2#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

On R1:

```
R1#Show ip bgp | b Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
```

Note: The above configuration did not affect R1 at all.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Task 5

Configure R1 such that R2 can establish an EBGP peer session with R1 using AS 111 or 100.

On R1:

```
R1(config)#Router bgp 111
R1(config-router)#Neighbor 12.1.1.2 local-as 100 no-prepend replace-as dual-as
```

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 200 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 100 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

To see R2's configuration:

On R2:

```
R2#Show run | s router bgp 200

router bgp 200
 no synchronization
 bgp log-neighbor-changes
 network 2.0.0.0
 neighbor 12.1.1.1 remote-as 100
 no auto-summary
```

To see R1's configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show run | s router bgp 111
router bgp 111
 no synchronization
 bgp log-neighbor-changes
 network 1.0.0.0
 neighbor 12.1.1.2 remote-as 200
 neighbor 12.1.1.2 local-as 100 no-prepend replace-as dual-as
 no auto-summary
```

To see R2's configuration:

On R2:

```
R2#Show run | S bgp
router bgp 200
 bgp log-neighbor-changes
 network 2.0.0.0
 neighbor 12.1.1.1 remote-as 100
```

The output of the above show command on R2 reveals that R2 is establishing a peer session with R1 using AS 100; the following verifies that R2 can also establish a peer session with R1 using AS 111:

To test the configuration:

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#Neighbor 12.1.1.1 remote-as 111
```

To verify the configuration:

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          12.1.1.1                 0             0 111 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
```

Note: The only difference here is the AS that originated and advertised 1.0.0.0/8 prefix.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

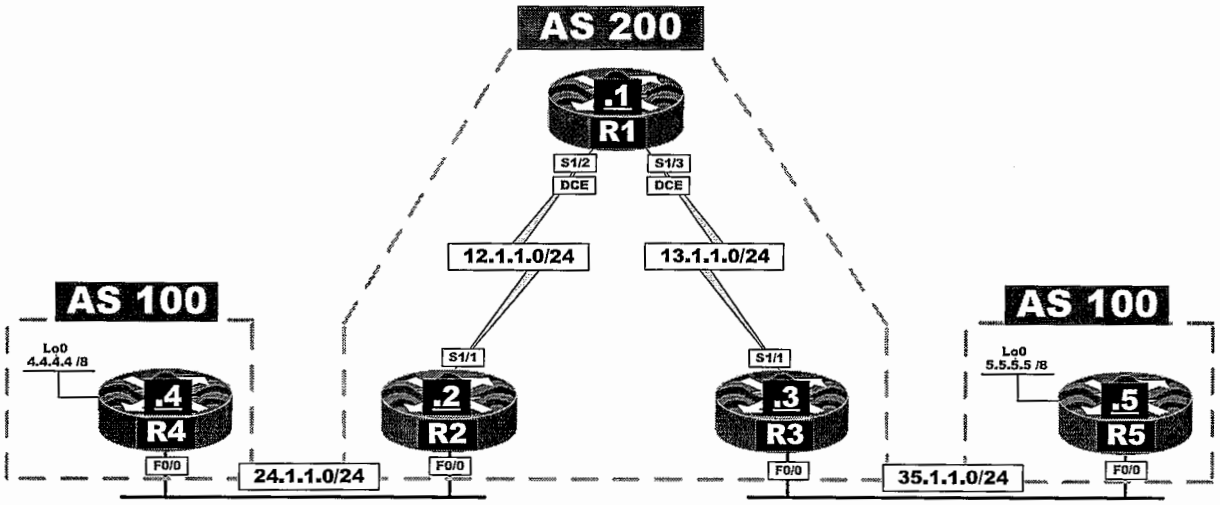
Task 6

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 19 – BGP Allowas-In



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→ "Advanced-init"→ "BGP"→ "Lab-19".

Task 1

Configure the routers according to the above diagram, if this configuration is performed successfully, all routers should have network 4.0.0.0 /8 and 5.0.0.0 /8 in their BGP and routing table. You should use OSPF in AS 200 to provide reachability.

On R1:

```
R1(config)#Router bgp 200
R1(config-router)#No au
R1(config-router)#Neighbor 12.1.1.2 remote-as 200
R1(config-router)#Neighbor 13.1.1.3 remote-as 200

R1(config)#Router ospf 1
R1(config-router)#Router-id 0.0.0.1
R1(config-router)#Network 12.1.1.1 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
R1(config-router)#Network 13.1.1.1 0.0.0.0 area 0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Neighbor 12.1.1.1 remote-as 200
R2(config-router)#Neighbor 12.1.1.1 Next-hop-self
R2(config-router)#Neighbor 13.1.1.3 remote-as 200
R2(config-router)#Neighbor 13.1.1.3 Next-hop-self
R2(config-router)#Neighbor 24.1.1.4 remote-as 100

R2(config)#Router ospf 1
R2(config-router)#Router-id 0.0.0.2
R2(config-router)#Network 12.1.1.2 0.0.0.0 area 0
R2(config-router)#Network 24.1.1.2 0.0.0.0 area 0
R2(config-router)#Passive-interface F0/0
```

On R3:

```
R3(config)#Router bgp 200
R3(config-router)#No au
R3(config-router)#Neighbor 13.1.1.1 remote-as 200
R3(config-router)#Neighbor 35.1.1.5 remote-as 100
R3(config-router)#Neighbor 13.1.1.1 Next-hop-self
R3(config-router)#Neighbor 12.1.1.2 remote-as 200
R3(config-router)#Neighbor 12.1.1.2 Next-hop-self

R3(config)#Router ospf 1
R3(config-router)#Router-id 0.0.0.3
R3(config-router)#Network 13.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 35.1.1.3 0.0.0.0 area 0
R3(config-router)#Passive-interface F0/0
```

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#No au
R4(config-router)#Network 4.0.0.0
R4(config-router)#Neighbor 24.1.1.2 remote-as 200
```

On R5:

```
R5(config)#Router bgp 100
R5(config-router)#No au
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config-router)#Network 5.0.0.0
R5(config-router)#Neighbor 35.1.1.3 remote-as 200
```

To verify the configuration:

On R1:

```
R1#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i4.0.0.0          12.1.1.2                 0    100      0 100 i
*>i5.0.0.0          13.1.1.3                 0    100      0 100 i
```

On R2:

```
R2#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path

*> 4.0.0.0          24.1.1.4                 0             0 100 i
*>i5.0.0.0          13.1.1.3                 0    100      0 100 i
```

On R3:

```
R3#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i4.0.0.0          12.1.1.2                 0    100      0 100 i
*> 5.0.0.0          35.1.1.5                 0             0 100 i
```

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          0.0.0.0                  0         32768 i
```

On R5:

```
R5#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path

*> 5.0.0.0          0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Not:e R4 is missing prefix 5.0.0.0 /8 and R5 is missing prefix 4.0.0.0 /8, the question is did R2 advertise prefix 5.0.0.0 /8 to R4? Or did R3 advertise prefix 4.0.0.0 /8 to R5?

The output of the following show command can be used to verify:

On R2:

```
R2#Show ip bgp neighbor 24.1.1.4 advertised-routes | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*>i5.0.0.0          13.1.1.3                 0    100      0 100 i

Total number of prefixes 1
```

Note: R2 has advertised prefix 5.0.0.0 /8 to R4 but this prefix came from AS 100 and since R4 is in the same AS (AS 100), it discarded that prefix, this is done by BGP as loop avoidance mechanism.

There are multiple ways to fix this problem, but one common way to accomplish this task is to use the "Allow-as" feature in BGP:

Allowas-in – This command can be configured on both R4 and R5, this command tells R4 and R5 that they should process and accept prefix/es that have matching AS number in the AS-path. Let's configure and test this option:

On R4:

```
R4(config)#Router bgp 100
R4(config-router)#Neighbor 24.1.1.2 allowas-in
```

On R5:

```
R5(config)#Router bgp 100
R5(config-router)#Neighbor 35.1.1.3 allowas-in
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | b Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          24.1.1.2                               0 200 100 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R5:

```
R5#Show ip bgp | b Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 4.0.0.0          35.1.1.3                               0 200 100 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
```

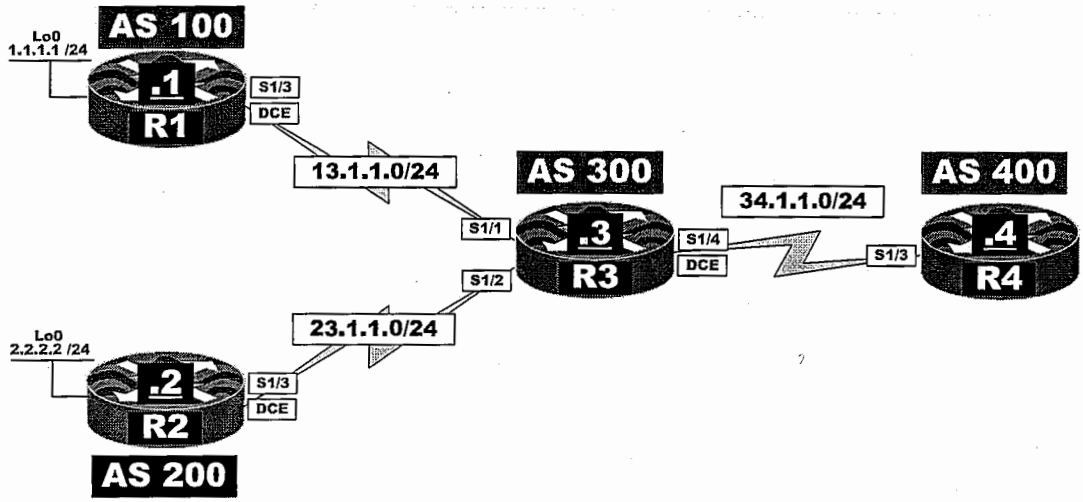
Note: Prefix 5.0.0.0 /8 is in R4's BGP table with AS number that matches its own. The same applies for prefix 4.0.0.0 /8 in R5's BGP table.

Task 2

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

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Lab 20 – BGP Automatic-tag, AS-path tag and Table-map



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-20".

Task 1

Configure BGP based on the following requirements:

- Configure R1 in AS 100 to establish a BGP peer session with R3 in AS 300. R1 should be configured to advertise its loopback 0 interface in this AS.
- Configure R2 in AS 200 to establish a BGP peer session with R3 in AS 300. R2 should be configured to advertise its loopback 0 interface in this AS.
- Configure R3 to advertise a default route to R1 in AS 100 and R2 in AS 200.
- Configure R4 in AS 400. This router should NOT establish a peer session with any of the routers in this topology. The router-id of this router should be configured to be 4.4.4.4.

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Router bgp 100
R1(config-router)#No au
R1(config-router)#Neighbor 13.1.1.3 remote-as 200
R1(config-router)#Network 1.1.1.0 mask 255.255.255.0
```

On R2:

```
R2(config)#Router bgp 200
R2(config-router)#No au
R2(config-router)#Neighbor 23.1.1.3 remote-as 300
R2(config-router)#Network 2.2.2.0 mask 255.255.255.0
```

On R3:

```
R3(config)#Router bgp 300
R3(config-router)#No au
R3(config-router)#Neighbor 13.1.1.1 remote-as 100
R3(config-router)#Neighbor 13.1.1.1 default-originate
R3(config-router)#Neighbor 23.1.1.2 remote-as 200
R3(config-router)#Neighbor 23.1.1.2 default-originate
```

On R4:

```
R4(config)#Router bgp 400
R4(config-router)#No au
R4(config-router)#bgp router-id 4.4.4.4
```

To verify the configuration:

On R3:

```
R3#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       13.1.1.1                 0             0 100 i
*> 2.2.2.0/24       23.1.1.2                 0             0 200 i
```

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          13.1.1.3                 0             0 300 i
*> 1.1.1.0/24       0.0.0.0                  0         32768 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 2.2.2.0/24        13.1.1.3                               0 300 200 i
```

On R2:

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 0.0.0.0          23.1.1.3                 0             0 300 i
*> 1.1.1.0/24       23.1.1.3                               0 300 100 i
*> 2.2.2.0/24       0.0.0.0                  0         32768 i
```

Task 2

Configure OSPF based on the following requirements:

- Configure OSPF area 0 on R3's S1/4 interface. The router-id of this router should be set to 0.0.0.3.
- Configure OSPF area 0 on R4's S1/3 interface. The router-id of this router should be set to 0.0.0.4.
- R3 should be configured to redistribute BGP into OSPF.
- R4 should be configured to redistribute OSPF into BGP.
- R4 should have networks 1.1.1.0/24 and 2.2.2.0/24 in its BGP table with the AS numbers in which these two prefixes originated in.

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Router-id 0.0.0.3
R3(config-router)#Network 34.1.1.3 0.0.0.0 area 0
R3(config-router)#Redistribute bgp 300 subnets
```

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#Router-id 0.0.0.4
R4(config-router)#Network 34.1.1.4 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#Router bgp 400
R4(config-router)#Redistribute ospf 1 match internal external
```

To verify the configuration:

On R3:

```
R3#Show ip ospf database

            OSPF Router with ID (0.0.0.3) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link
count
0.0.0.3         0.0.0.3         173         0x80000004 0x00AA9F 2
0.0.0.4         0.0.0.4         179         0x80000003 0x009AB0 2

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
1.1.1.0         0.0.0.3         265         0x80000001 0x000247 100
2.2.2.0         0.0.0.3         265         0x80000001 0x00ECF4 200
```

When BGP is redistributed into an IGP, the BGP AS number is copied as a route tag in that given IGP. By looking at the external routes, we can see that network 1.1.1.0 has a Tag of 100, and network 2.2.2.0 is tagged with 200.

Let's verify the BGP table on R4:

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       34.1.1.3                 1             32768 ?
*> 2.2.2.0/24       34.1.1.3                 1             32768 ?
*> 34.1.1.0/24      0.0.0.0                  0             32768 ?
```

What happened to the AS-Path?

NOTE: When OSPF was redistributed into BGP, the AS-Path was lost and the origin code was set to unknown.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

This AS-Path can be retrieved on R4 using the tag that was set when the BGP was redistributed into OSPF; this can be configured using the "AS-Path tag", let's configure this and verify:

On R4:

```
R4(config)#Route-map tst permit 10
R4(config-route-map)#Set as-path tag

R4(config)#Router bgp 400
R4(config-router)#Redistribute ospf 1 match internal external route-map tst
```

To verify the configuration:

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       34.1.1.3                 1             32768 100 ?
*> 2.2.2.0/24       34.1.1.3                 1             32768 200 ?
*> 34.1.1.0/24      0.0.0.0                  0             32768 ?
```

The output of the above show command reveals that the route tags which were the BGP AS numbers are copied into the AS-Path. But the second problem that we are facing is the origin code, which is set to unknown.
R3 can be configured to fix this problem, in this case R3 can be configured to use the "Automatic-tag and the Table-map" as follows:

On R3:

The following configures a route-map and sets the "automatic-tag":

```
R3(config)#Route-map tst permit 10
R3(config-route-map)#Set automatic-tag
```

The route-map is referenced under router BGP using the "Table-map":

```
R3(config)#Router bgp 300
R3(config-router)#Table-map tst
```

To implement the changes:

```
R3#Clear ip bgp *
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R4:

```
R4#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       34.1.1.3                 1          32768 3489661028 i
*> 2.2.2.0/24       34.1.1.3                 1          32768 3489661128 i
*> 34.1.1.0/24      0.0.0.0                  0          32768 ?
```

So what happened to the AS number?
The AS numbers are correct but we have to go through a calculation to figure them, let's do this:

Take the number and convert it to Hex:
3489661028 → D0000064
Let's focus on the last 4 characters "0064" and convert it back to decimal → 100

Let's convert the AS number for network 2.2.2.0/24:
3489661128 → D00000C8
Let's convert the last 4 characters "00C8" to decimal → 200

There MUST be an easier/better way to do this, how would you handle this kind of issue in the real world?

Let's remove the "route-map tst" that was configured on R4:

On R4:

```
R4(config)#No route-map tst

R4(config)#Router bgp 400
R4(config-router)#No redis ospf 1 match inter exte 1 exte 2 route-map tst
R4(config-router)#Redistr ospf 1 match inter exter

R4#Clear ip bgp *
```

NOW, let's check the OSPF database and verify the tag for networks 1.1.1.0/24 and 2.2.2.0/24:

```
R4#Show ip ospf database

            OSPF Router with ID (0.0.0.4) (Process ID 1)
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link
count
0.0.0.3         0.0.0.3         1607        0x80000004 0x009B0F 2
0.0.0.4         0.0.0.4         1606        0x80000002 0x008D1F 2

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum Tag
1.1.1.0         0.0.0.3         1326        0x80000003 0x003A3C 3489661028
2.2.2.0         0.0.0.3         1326        0x80000003 0x0025E9 3489661128
```

Based on the output of the above show command we can see that the "Automatic tag" is the feature that changed the tag to the two highlighted in yellow and blue. Let's verify the tags in the BGP table:

```
R4#Show ip bgp | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       34.1.1.3                 1             32768 100 i
*> 2.2.2.0/24       34.1.1.3                 1             32768 200 i
*> 34.1.1.0/24      0.0.0.0                  0             32768 ?
```

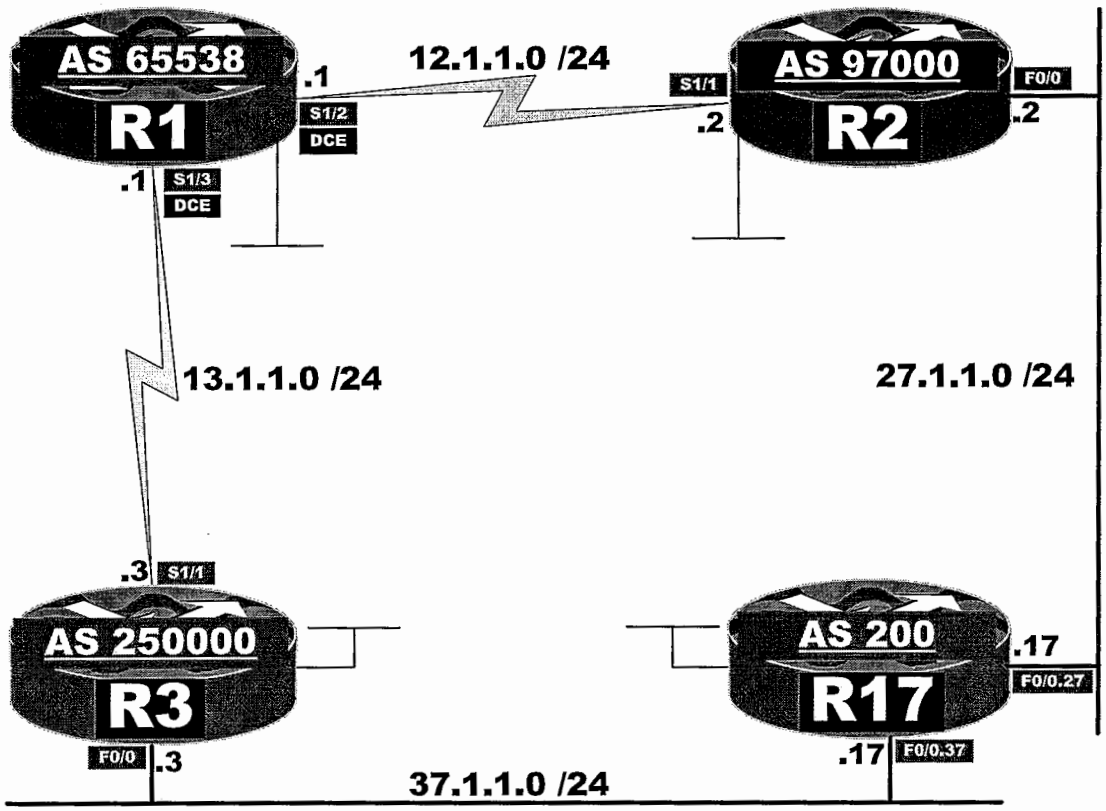
Task 3

Erase the startup configuration on all routers and switches, delete vlan.dat on all switches and reload all devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 21 – 4 Byte AS Numbers



Lab Setup:

To copy and paste the initial configurations, go to http://micronics.nl"→"Advanced-init"→ "BGP"→ "Lab-21".

Task 1

Configure BGP based on the above diagram, the IOS on R17 does NOT support 4 Byte ASNs. Advertise the loopback0 interfaces in BGP and disable auto-summarization.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

Besides the IP addressing the AS numbers are also depleting so they added another 2 Bytes to the AS numbers to increase the number of ASes.

16 bits or 2 Bytes gives you 65536 or "0 – 65535". With 4 Bytes allocated to the AS numbers, it increased the 65536 possible numbers to 4,294,967,296 possible AS Numbers.

The following identifies the ranges:

- The range for 16-bit ASNs are: 0 – 65535
- The range for 32-bit ASNs ONLY are: 65536 - 4,294,967,296
- The range for 32-bit ASNs are: 0 - 4,294,967,296

These AS Numbers include the previous range, and this helps for interoperability between 2 Bytes and 4 Byte ASNs.

But how does it all work?
The 32 bits AS Number can be written in four possible ways:

1. Asplain – This is a simple decimal representation and the AS Numbers can be written in plain decimal and the choices are 0 – 4,294,967,296.

2. Asdot+ - This technique divides the entire number into two equal sections and in each section there are 16 bits or 2 Bytes. So we end up with 2 Bytes of high order and 2 Bytes of low order chunks of 16 bits, separated by a dot, hence the name "asdot". NOT a lot of support. If the high order bits are all set to zero, the AS number will look like "0.65535", so if we increment the number by one, it will be 65536, correct? But 65536 is outside of the range, with the range being 0-65535 in the low order bit value, so how do we write this AS number? We can write it like "1.0" which is 65536, and the next value 65537 can be written as 1.1, or 68000 can be written as "1.2464". The way we sort this out is by subtracting the 65536 from 68000, this gives us a difference of 2464. Since 65536 goes into 68000 once, the high order is set to "1" and the difference of 2464 goes into the second section. Therefore, the high order 16 bits are in increments of 65536.

3. ASdot – This is a mixture of both. In this case an AS number within the 2 Bytes range can be written in ASplain format, and any number outside of that range can be written in ASdot+ format. But you can write them in any way you like, for example if your AS number is "100", this can be written as simply "100" which is in Asplain format or in ASdot which is "0.100" format. An ASN of 65537 can simply be written as 65537 or "1.1", this format is widely used.

4. AScolon – In this case instead of the dot a colon is used and this is NOT widely used at all and ONLY used if you have Redback routers.

What if you have a mixture of BGP speakers in which case one is using the new 4 Byte ASN and the other

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

BGP speaker is using a 2 Byte ASN and it does NOT support the 4 Byte ASN?

Well, we know that this capability will be advertised in the BGP capability Advertisements, to be accurate, in the OPEN message. But we have to configure BGP to get to that stage, the question is how and what do we configure on the BGP speaker that ONLY supports the 2 Byte ASN? Let's say it's peering up with a neighbor that is in AS 69000?

This is when the use of "AS-TRANS" comes to rescue. The "AS-TRANS BGP ASN" which is a reserved 2 Byte ASN of "23456". Let's configure this and test it.

On R1:

```
R1(config)#Router bgp 65538
```

When configuring BGP with 4 Byte ASNs the router id must be configured:

```
R1(config-router)#bgp router-id 11.1.1.1
R1(config-router)#No au
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Neighbor 12.1.1.2 remote-as 97000
R1(config-router)#Neighbor 13.1.1.3 remote-as 250000
```

On R2:

```
R2(config)#Router bgp 97000
R2(config-router)#bgp router-id 22.2.2.2
R2(config-router)#No au
R2(config-router)#Netw 2.0.0.0
R2(config-router)#Neighbor 12.1.1.1 remote-as 65538
```

NOTE: You can have a mixture of neighbors where one is using the 4 Byte ASN and the other using the 2 Byte ASN.

```
R2(config-router)#Neighbor 27.1.1.17 remote-as 200
```

On R3:

```
R3(config)#Router bgp 250000
R3(config-router)#bgp router-id 33.3.3.3
R3(config-router)#No au
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Neighbor 13.1.1.1 remote-as 65538
R3(config-router)#Neighbor 37.1.1.17 remote-as 200
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

On R1:

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 2.0.0.0          12.1.1.2                 0             0 97000 i
*> 3.0.0.0          13.1.1.3                 0             0 250000 i

R1#Show ip bgp neighbors | s Neighbor capabilities

  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Multisession Capability:
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Multisession Capability:
```

Let's configure R17, remember that this router does NOT have the capability to handle 4 Byte ASN:

On R17:

Since R17 does not have the capability to run 4 Byte ASN, it will use the "AS-TRANS" ASN which is 23456 instead of the ASN of 97000 and 250000 for routers R2 and R3 respectively.

```
R17(config)#Router bgp 200
R17(config-router)#No au
R17(config-router)#Netw 7.0.0.0
R17(config-router)#Neighbor 27.1.1.2 remote-as 23456
R17(config-router)#Neighbor 37.1.1.3 remote-as 23456
```

You should see the following console messages stating that a Neighbor adjacency has been formed with neighbors 27.1.1.2 and 37.1.1.3.

```
*Sep 27 10:45:57.663: %BGP-5-ADJCHANGE: neighbor 27.1.1.2 Up
*Sep 27 10:46:20.471: %BGP-5-ADJCHANGE: neighbor 37.1.1.3 Up

R17#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          37.1.1.3                               0 23456 23456 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*>                  27.1.1.2                                 0 23456 23456 i
*  2.0.0.0          37.1.1.3                                 0 23456 23456 23456 i
*>                  27.1.1.2                 0               0 23456 i
*> 3.0.0.0          37.1.1.3                 0               0 23456 i
*                   27.1.1.2                                 0 23456 23456 23456 i
*> 7.0.0.0          0.0.0.0                  0           32768 i
```

Even though the peer sessions have been established the problem here is that from R17's perspective all the prefixes were originated from AS 23456. This can cause some serious problems.

```
R17#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 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

Task 2

Configure R1 such that the 32 bit ASNs are displayed in "asdot" format.

On R1

```
R1(config)#Router bgp 65538
```

The following command changes the format of the 32 bits ASNs from "asplain" to "asdot":

```
R1(config-router)#bgp asnotation dot

R1#Clear ip bgp *
```

To verify the configuration:

On R1

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*  2.0.0.0          13.1.1.3                               0 3.53392 200 1.31464 i
*>                  12.1.1.2                 0             0 1.31464 i
```


© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets

```
*> 3.0.0.0          13.1.1.3                 0             0 3.53392 i
*                   12.1.1.2                               0 1.31464 200 3.53392 i
*  7.0.0.0          13.1.1.3                               0 3.53392 200 i
*>                  12.1.1.2                               0 1.31464 200 i
```

We can see that the 32 bit ASNs are converted from "asplain" to "asdot" format. Let's see how the IOS calculated the ASNs, let's look at "3.53392":

We know that the high order numbers are in increments of 65536, therefore, we must multiply 3 by 65536:
3 X 65,536 = 196608

Then we need to add the 53,392 to 196,608:
196,608 + 53392 = 250,000

We can see that "3.53392" represents AS 250,000 which is the AS that R3 belongs to. Let's calculate "1.31464":

1 X 65,536 = 65,536 + 31,464 = 97,000
This is the ASN that R3 belongs to.

Let's change the format on R2 and verify the output of the "Show ip bgp" command:

On R2

NOTE: We can see that the asdot format can also be used:

```
R2(config)#Router bgp 1.31464
R2(config-router)#bgp asnotation dot

R2#Clear ip bgp *
```

To verify the configuration:

On R2

```
R2#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          27.1.1.17                              0 200 3.53392 1.2 i
*>                  12.1.1.1                 0             0 1.2 i
*> 2.0.0.0          0.0.0.0                  0         32768 i
*> 3.0.0.0          27.1.1.17                              0 200 3.53392 i
*                   12.1.1.1                               0 1.2 3.53392 i
*> 7.0.0.0          27.1.1.17                0             0 200 i
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
*                    12.1.1.1                                0 1.2 3.53392 200 i
```

Let's calculate the ASN for "1.2":

65536 + 2 = 65538

Task 3

Erase the startup configuration on all routers and switches, delete vlan.dat on all switches and reload all devices before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved
https://t.me/learningnets