azarpara.vahid@gmail.com

[redacted]

CCIE Foundation

5.0

www.MicronicsTraining.com

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

VOL-I

© 2014 Narbik Kocharians. All rights reserved

[redacted]

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

azarpara.vahid@gmail.com

Table of Content:

Subject	Page
Topology	4

Section One:
Logical or Physical

Subject	Page
Lab 1 – Physical to Logical Topology I	9
Lab 2 – Physical to Logical Topology II	18
Lab 3 – Physical to Logical Topology III	33

Section Two:
3560 Switching

Subject	Page
Lab 1 – Basic 3560 configuration	56
Lab 2 – Spanning-tree 802.1d	92

Section Three:
DMVPN

Subject	Page
Lab 1 – DMVPN Phase #1 with Static Mappings	108
Lab 2 - DMVPN Phase #1 with Dynamic Mappings	118
Lab 3 – DMVPN Phase #2 and 3 with Static Mappings	126
Lab 4 – DMVPN Phase #2 and 3 with Dynamic Mappings	134
Lab 5 – Running Routing Protocols on DMVPN Phase #1	141
Lab 6 – Running Routing Protocols on DMVPN Phase #2 and 3	165

Section Four:
RIPv2

Subject	Page
Lab 1 – Configuring RIPv2	191
Lab 2 – RIPv2 Authentication (Clear text and MD5)	199
Lab 3 – Configuring different RIPv2 Update methods	205
Lab 4 – Injection of Default routes in RIPv2	211
Lab 5 – Filtering RIPv2 routes	221

© 2014 Narbik Kocharians. All rights reserved

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

Section five:

Eigrp

Subject	Page
Lab 1 – Configuring Eigrp and Adjusting the Timers	230
Lab 2 – Eigrp Metric	239
Lab 3 – Eigrp Summarization	243
Lab 4 – Eigrp Authentication & Advanced Configuration	254
Lab 5 – Eigrp Stub	260

Section Six:

OSPF

Subject	Page
Lab 1 – Advertising Networks	274
Lab 2 – OSPF Broadcast Networks	291
Lab 3 – OSPF Non-Broadcast Networks	307
Lab 4 – OSPF Point-to-point Networks	318
Lab 5 – OSPF "Point-to-Multipoint" & "Point-to-Multipoint" Networks	322
Lab 6 – OSPF Cost	329
Lab 7 – OSPF Authentication	337
Lab 8 – OSPF Summarization	370
Lab 9 – OSPF Filtering	382
Lab 10 – Virtual-links and GRE Tunnels	413
Lab 11 – OSPF Stub, T/Stubby, NSSA, NSS-Stub, NSS-T/Stub	427

Section Seven:

Redistribution

Subject	Page
Lab 1 – Redistribution Basics	451

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

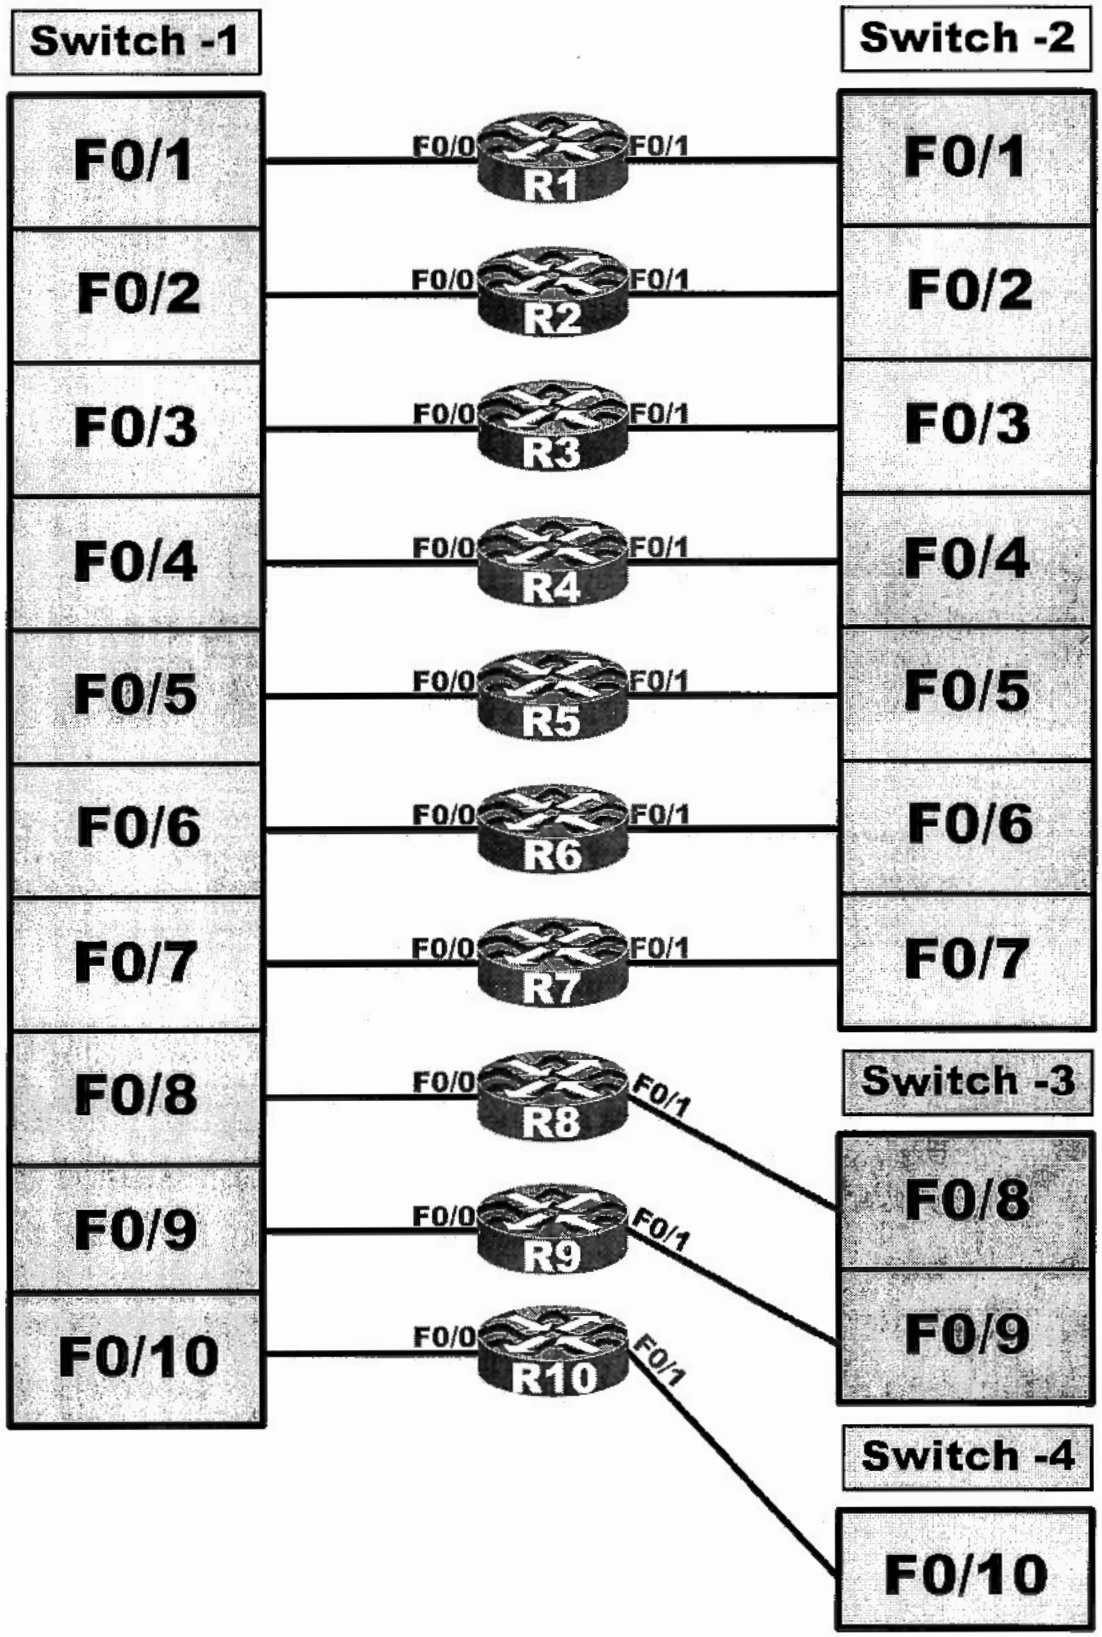

© 2014 Narbik Kocharians. All rights reserved

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

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

R6	S1/5	R5	S1/6
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

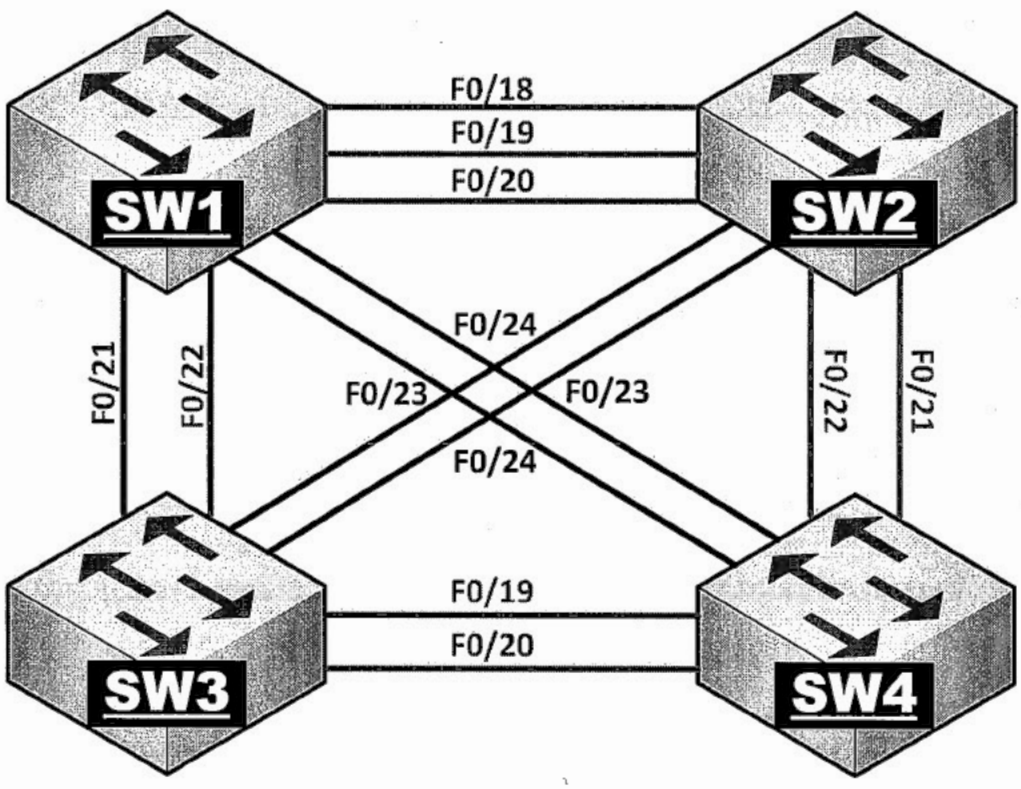

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

CCIE Foundation

v5.0

www.MicronicsTraining.com

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

Physical or Logical

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

LAB 1-
Physical to Logical Topology - I

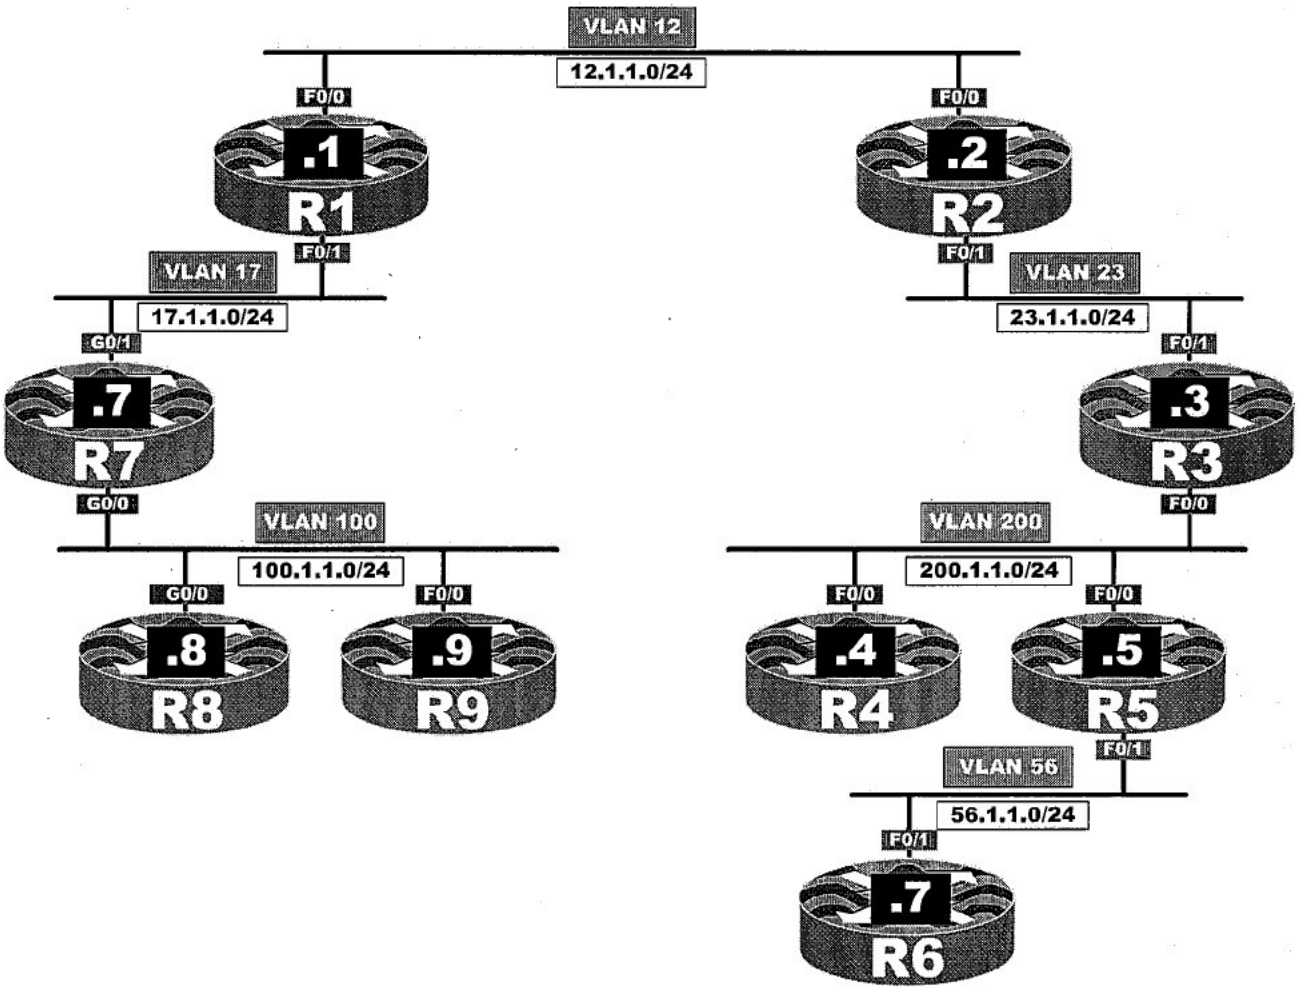

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Shutdown all ports on all switches.

On All Switches:

```
SWx(config)#Int range f0/1-24
SWx(config-if-range)#Shut
```

Task 2

Configure the above topology, if this configuration is performed successfully, every router should be able to ping its neighboring router/s in the same subnet.

Let's start with R1 and R2's connection in VLAN 12, we can see that these two routers are connected via their F0/0 interfaces, and their F0/1 interfaces are connected to other routers in another VLAN.

If the physical topology is checked, you can easily see that the F0/0 interfaces of these two routers are connected to SW1's F0/1 and F0/2 for R1 and R2 respectively, so let's configure these two ports on SW1 in VLAN 12 and verify.

On SW1:

```
SW1(config)#Int range f0/1-2
SW1(config-if-range)#Switchport mode access
SW1(config-if-range)#switchport access vlan 12
SW1(config-if-range)#No shut
```

Let's verify:

On SW1:

```
SW1#Show vlan brief | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    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
                                                Fa0/19, Fa0/20, Gi0/1, Gi0/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
12   VLAN0012                         active    Fa0/1, Fa0/2
```

Let's configure the F0/0 interfaces of R1 and R2:

On R1:

```
R1(config)#Int F0/0
R1(config-if)#Ip addr 12.1.1.1 255.255.255.0
R1(config-if)#No shut
```

On R2:

```
R2(config)#Int F0/0
R2(config-if)#Ip addr 12.1.1.2 255.255.255.0
R2(config-if)#No shut
```

To test and verify the configuration:

On R1:

```
R1#Ping 12.1.1.2

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

We can configure R2's connection to R3 or R1's connection to R7, the following configures R1's connection to R7.
Before we assign an IP address to the interfaces of these routers, let's configure the F0/1 interface of R1 and the G0/1 interface of R7 in VLAN 17 on SW2, and then configure the routers with the assigned IP addresses.
We can see that these interfaces are connected to SW2's F0/1 and F0/7 for R1 and R7 respectively; therefore, these two ports on SW2 should be configured in VLAN 17.

On SW2:

```
W2(config)#Int Range f0/1,f0/7
SW2(config-if-range)#Swi mode acc
SW2(config-if-range)#Swi acc v 17
SW2(config-if-range)#No shut
```

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#Int F0/1
R1(config-if)#Ip address 17.1.1.1 255.255.255.0
R1(config-if)#No shut
```

On R7:

```
R7(config)#Int G0/1
R7(config-if)#Ip addr 17.1.1.7 255.255.255.0
R7(config-if)#No shut
```

To verify the configuration:

On R1:

```
R1#Ping 17.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 17.1.1.7, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

NOW…let's configure the R2 and R3's F0/1 interfaces in VLAN 23, we can see that these two interfaces are connected to SW2's F0/2 and F0/3 for R2's F0/1 and R3's F0/1 interfaces respectively.

On SW2:

```
SW2(config)#Int Range F0/2-3
SW2(config-if-range)#Swi mode acc
SW2(config-if-range)#swi acc v 23
SW2(config-if-range)#No shut
```

On R2:

```
R2(config)#Int F0/1
R2(config-if)#Ip addr 23.1.1.2 255.255.255.0
R2(config-if)#No shut
```

On R3:

```
R3(config)#Int F0/1
R3(config-if)#Ip addr 23.1.1.3 255.255.255.0
R3(config-if)#No shut
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Let's move on to R7, R8 and R9's configuration in VLAN 100. In this case we can see that R7's G0/0 interface is connected to SW1's port F0/7, R8's G0/0 interface is connected to SW1's F0/8 interface, and R9's F0/0 interface is connected to SW1's F0/9 interface.

On SW1:

```
SW1(config)#Int Range f0/7-9
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc v 100
SW1(config-if-range)#No shut
```

Let's verify the VLAN configuration:

On SW1:

```
SW1#Show vlan br | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/3, Fa0/4, Fa0/5, Fa0/6
                                                Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                                Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                                Fa0/18, Fa0/19, Fa0/20, Fa0/21
                                                Fa0/22, Fa0/23, Fa0/24, Gi0/1
                                                Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2
100  VLAN0100                         active    Fa0/7, Fa0/8, Fa0/9
```

Let's configure the routers:

On R7:

```
R7(config)#Int G0/0
R7(config-if)#Ip addr 100.1.1.7 255.255.255.0
R7(config-if)#No shut
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R8:

```
R8(config)#Int G0/0
R8(config-if)#Ip addr 100.1.1.8 255.255.255.0
R8(config-if)#No shut
```

On R9:

```
R9(config)#Int F0/0
R9(config-if)#IP addr 100.1.1.9 255.255.255.0
R9(config-if)#No shut
```

To test the configuration:

On R7:

```
R7#Ping 100.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.8, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R7#Ping 100.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.9, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

The second last VLAN left to be configured is VLAN 200, by looking at the interfaces of the routers used in this VLAN we can see that R3, R4 and R5 are all using their F0/0 interfaces in this VLAN, let's configure the switch and then the routers.

On SW1:

```
SW1(config)#int range f0/3-5
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc v 200
SW1(config-if)#No shut
```

To verify the configuration:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1#Show vlan br | exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/6, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2
100  VLAN0100                         active    Fa0/7, Fa0/8, Fa0/9
200  VLAN0200                         active    Fa0/3, Fa0/4, Fa0/5
```

On R3:

```
R3(config)#Int F0/0
R3(config-if)#Ip addr 200.1.1.3 255.255.255.0
R3(config-if)#No shut
```

On R4:

```
R4(config)#Int F0/0
R4(config-if)#Ip addr 200.1.1.4 255.255.255.0
R4(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/0
R5(config-if)#Ip addr 200.1.1.5 255.255.255.0
R5(config-if)#No shut
```

To verify the configuration:

On R3:

```
R3#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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

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:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

The last VLAN is VLAN 56, since both routers are using their F0/1 interfaces in this VLAN, we should configure SW2 to assign its F0/5 and F0/6 interfaces to VLAN 56.

On SW2:

```
SW2(config)#Int range F0/5-6
SW2(config-if)#Swi mode acc
SW2(config-if)#Swi acc v 56
SW2(config-if)#No shu
```

To verify the configuration:

On SW2:

```
SW2#Show vlan br | exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                                Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                                Fa0/18, Fa0/19, Fa0/20, Fa0/21
                                                Fa0/22, Fa0/23, Fa0/24, Gi0/1
                                                Gi0/2
12   VLAN0012                         active    Fa0/1, Fa0/2
56   VLAN0056                         active    Fa0/5, Fa0/6
100  VLAN0100                         active    Fa0/7, Fa0/8, Fa0/9
200  VLAN0200                         active    Fa0/3, Fa0/4
```

On R6:

```
R6(config)#Int F0/1
R6(config-if)#IP addr 56.1.1.6 255.255.255.0
R6(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/1
R5(config-subif)#IP addr 56.1.1.5 255.255.255.0
R5(config-if)#No shut
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6#Ping 56.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 56.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Task 3

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

LAB 2 -
Physical to Logical Topology - II

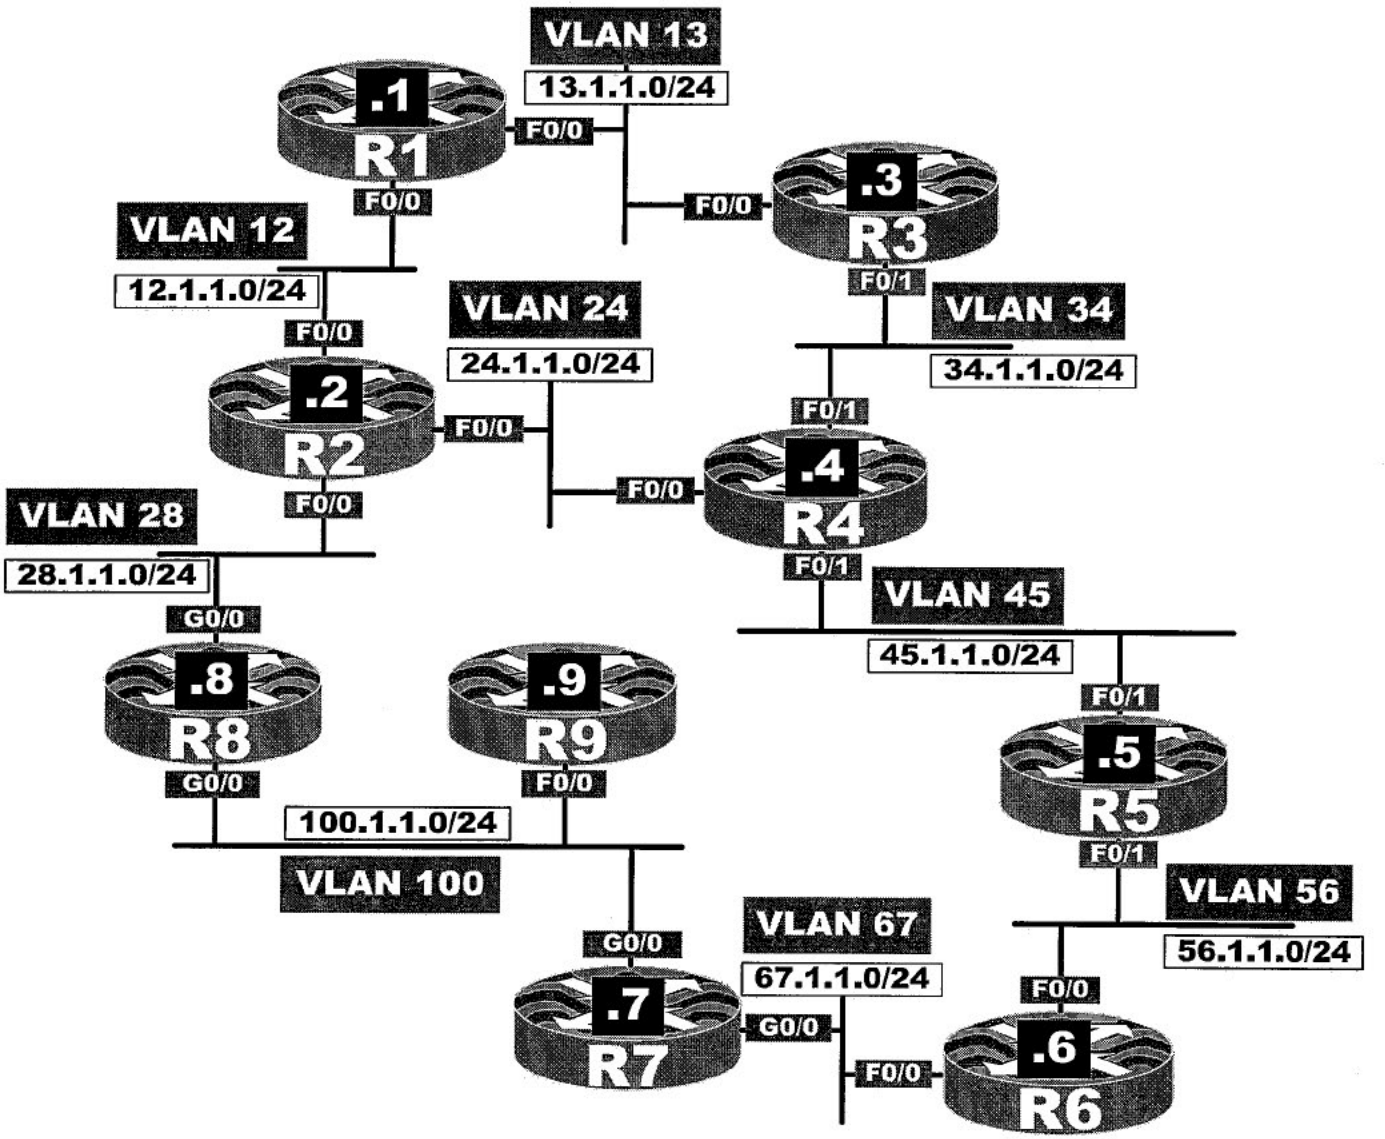

Task 1

Shutdown all ports on all switches.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On All Switches:

```
SWx(config)#Int range f0/1-24
SWx(config-if-range)#Shut
```

Task 2

Configure the above topology, if this configuration is performed successfully, every router should be able to ping its neighboring router/s in the same subnet.

Let's do a top down configuration starting from VLAN 13 all the way to VLAN 67.

NOTE: The F0/0 interface of R3 is configured in this VLAN, and the other Ethernet interface of this router is configured in another VLAN, whereas, the F0/0 interface of R1 is configured in two VLANs, VLAN 13 and VLAN 12; since this is Physically impossible, logical interfaces must be configured to accomplish this task; to accomplish this task, on SW1, a trunk is configured with different DOT1q VLAN tags, 12 for VLAN 12 and 13 for VLAN 13.

Since the F0/0 interface of all routers are connected to SW1, let's configure SW1 for these routers:

On SW1:

```
SW1(config)#Int F0/3
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc vlan 13
SW1(config-if)#No shut
```

NOTE: Since the F0/1 interface of SW1 is connected to R1's F0/0 interface, and R1's F0/0 interface must be configured in different VLANs, the F0/1 interface of this switch MUST be configured as a trunk.

```
SW1(config)#Int F0/1
SW1(config-if)#Swi trunk encap dot1q
SW1(config-if)#Swi mode trunk
SW1(config-if)#No shut
```

Let's configure the routers starting with R3:

On R3:

```
R3(config)#Int F0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-if)#IP addr 13.1.1.3 255.255.255.0
R3(config-if)#No shut
```

On R1:

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

R1(config-if)#Int F0/0.13
R1(config-subif)#Encap dot1q 13
R1(config-subif)#Ip addr 13.1.1.1 255.255.255.0
```

To verify the configuration:

On SW1:

```
SW1#Show interface trunk

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

Port        Vlans allowed on trunk
Fa0/1       1-4094

Port        Vlans allowed and active in management domain
Fa0/1       1,13

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1       1,13
```

On R1:

```
R1#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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

NOW….let's configure VLAN 34 connecting R3 to R4:
We need some configuration on the switch to which these routers are connected to, let's begin with the Switch configuration. Since the F0/1 interface of R3 is connected to SW2, the F0/3 interface of SW2 must be configured in VLAN 34:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2(config)#Int F0/3
SW2(config-if)#Swi mode acc
SW2(config-if)#Swi acc vlan 34
SW2(config-if)#No shut
```

NOTE: R4's F0/1 interface is also connected to SW2, but this interface is also configured in another VLAN (VLAN 45), so we know that the F0/1 interface of R4 must be configured as a trunk and the port on the Switch (SW2) to which it is connected should also be configured as trunk.

On SW2:

```
SW2(config)#int F0/4
SW2(config-if)#Swi trun encap dot1q
SW2(config-if)#Swi mode trunk
SW2(config-if)#No shut
```

Since the Switch is configured, let's move on to the routers starting with R3. This router's configuration is very basic and all we need to do is assign an IP address and "No Shut" the F0/1 interface.

On R3:

```
R3(config)#Int F0/1
R3(config-if)#Ip addr 34.1.1.3 255.255.255.0
R3(config-if)#No shut
```

Let's configure R4; this interface must be configured with sub-interfaces.

On R4:

```
R4(config)#Int F0/1
R4(config-if)#No shut

R4(config)#int F0/1.34
R4(config-subif)#Encap dot1q 34
R4(config-subif)#Ip addr 34.1.1.4 255.255.255.0
```

To verify and test the configuration:

On SW2:

```
SW2#Show interface trunk
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/4       on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/4       1-4094

Port        Vlans allowed and active in management domain
Fa0/4       1,34

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/4       1,34

R4#Ping 34.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

So we can see that when a Physical Ethernet interface is configured in multiple VLANs, the interface of the router MUST be configured with sub-interfaces and the port on the switch to which it is connected to MUST also be configured as a trunk.

Let's configure VLAN 12. Just like any VLAN configuration we have some configuration to perform on the switch/es and some configuration on the router/s.

In this VLAN, R1's F0/0 interface must be configured with another sub-interface, remember earlier the F0/0 interface of R1 was configured with a sub-interface for VLAN 13; we also know that the F0/1 interface of the SW1 is already configured as a trunk, let's verify this information:

On SW1:

```
SW1#Show interface trunk

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

Port        Vlans allowed on trunk
Fa0/1       1-4094

Port        Vlans allowed and active in management domain
Fa0/1       1,13

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1       1,13
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Let's configure SW1 for R2, but once again we can see that the F0/0 interface of R2 is configured in two different VLANs, this means that the F0/0 interface of R2 should be configured with two sub-interfaces, and the port to which it is connected to MUST also be configured as trunk.

On SW1:

```
SW1(config)#Int F0/2
SW1(config-if)#Swi trunk encap dot1q
SW1(config-if)#Swi mode trunk
SW1(config-if)#No shut
```

On R1:

```
R1(config)#Int F0/0.12
R1(config-subif)#Encap dot1q 12
R1(config-subif)#Ip address 12.1.1.1 255.255.255.0
```

On R2:

```
R2(config)#Int F0/0
R2(config-if)#No shut

R2(config)#Int F0/0.12
R2(config-subif)#Encap dot1q 12
R2(config-subif)#Ip addr 12.1.1.2 255.255.255.0
```

To verify the configuration:

On R1:

```
R1#Ping 12.1.1.2

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

What went wrong?
Let's verify and see if the VLAN is allowed to traverse over the trunk links:

On SW1:

```
SW1#Show interface trunk
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/1       on           802.1q         trunking      1
Fa0/2       on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/1       1-4094
Fa0/2       1-4094

Port        Vlans allowed and active in management domain
Fa0/1       1,13
Fa0/2       1,13

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1       1,13
Fa0/2       1,13
```

ONLY VLAN 13 is allowed over the trunk, but WHY? Let's see all the configured VLANs:

On SW1:

```
SW1#Show vlan brie | Exc unsup

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    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, Fa0/19
                                                Fa0/20, Fa0/21, Fa0/22, Fa0/23
                                                Fa0/24, Gi0/1, Gi0/2
13   VLAN0013                         active    Fa0/3
```

VLAN 13 was created when the F0/3 interface of SW1 was placed in VLAN 13, since none of the interfaces of SW1 is implicitly configured in VLAN 12 this VLAN was never created. Let's configure VLAN 12 on SW1:

On SW1:

```
SW1(config)#VLAN 12
SW1(config-vlan)#Exit
```

To test and verify the configuration:

On R1:

You may have to wait for Spanning-tree to converge before the ping is successful.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Ping 12.1.1.2

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

Perfect…..Let's configure VLAN 24:

On SW1:

NOTE: Since by placing the F0/4 interface of SW1 in VLAN 24, the IOS will auto-create this VLAN we won't run into the previous problem.

```
SW1(config)#int F0/4
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc vlan 24
SW1(config-if)#No shut
```

On R2:

Another sub-interface is configured in VLAN 24:

```
R2(config)#Int F0/0.24
R2(config-subif)#Encap dot1q 24
R2(config-subif)#Ip addr 24.1.1.2 255.255.255.0
```

On R4:

```
R4(config)#Int F0/0
R4(config-if)#Ip addr 24.1.1.4 255.255.255.0
R4(config-if)#No shut
```

To verify the configuration:

On R2:

```
R2#Ping 24.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24.1.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Next VLAN is VLAN 28. We can easily see that another sub-interface must be configured on R2.
The F0/2 interface of SW1 is already configured as trunk. R8's G0/0 interface is in two different VLANs, so a sub-interface must be configured on R8 and the port to which the interface is connected to must be configured as a trunk.

Let's start with SW1's configuration:

On SW1:

The port that R8's F0/0 interface is connected is configured as a trunk to allow VLANs 22 and 123 to traverse through:

```
SW1(config)#Int F0/8
SW1(config-if)#Swi tru encap dot1q
SW1(config-if)#SWi mode trunk
SW1(config-if)#No shut
```

VLAN 28 MUST be configured on the switch.

```
SW1(config)#Vlan 28
SW1(config-vlan)#exit
```

Let's configure another sub-interface for VLAN 28 on R2:

On R2:

```
R2(config)#Int F0/0.28
R2(config-subif)#Encap dot1q 28
R2(config-subif)#Ip addr 28.1.1.2 255.255.255.0
```

On R8:

```
R8(config)#Int G0/0
R8(config-if)#No shut

R8(config)#Int G0/0.28
R8(config-subif)#Encap dot1q 28
R8(config-subif)#Ip addr 28.1.1.8 255.255.255.0
```

To verify the configuration:

On R2:

```
R2#Ping 28.1.1.8
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 22.1.1.22, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Before going further into the configuration of this topology, let's summarize what we have covered so far in this lab:

When configuring routers in a VLAN we MUST pay attention to the following:

If the router's interface is in ONE VLAN, then, configure the VLAN on the switch and assign the interface to which the router is connected to in that VLAN.

If the router's interface is configured in multiple VLANs, then configure the interface of the router as a trunk, remember that ISL encapsulation is only available on the older IOS and routers and no longer in the CCIE Routing and Switching blueprint, therefore the encapsulation is configured as DOT1q, and this means we configure multiple sub-interfaces on the router.

Each sub-interface should be configured in the appropriate VLAN as identified in the topology. The switchport to which the router is connected to must also be configured as a trunk, YOU MUST ENSURE THAT THE VLAN IS CONFIGURED AND IT IS ALLOWED TO TRAVERSE THE TRUNK.

Let's configure VLAN 45. R4 needs another sub-interface configuration; R5's F0/1 interface should also be configured with sub-interfaces because it is in two different VLANs, and the F0/5 interface of SW2 should also be configured as a trunk and VLAN 45 MUST be configured/created on SW2.

On SW2:

```
SW2(config)#Int F0/5
SW2(config-if)#Swi trunk encap dot1q
SW2(config-if)#Swi mode trunk
SW2(config-if)#No shut

SW2(config)#Vlan 45
SW2(config-vlan)#exit
```

On R4:

```
R4(config)#Int F0/1.45
R4(config-subif)#encap dot1q 45
R4(config-subif)#Ip addr 45.1.1.4 255.255.255.0
```

On R5:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config)#Int F0/1
R5(config-if)#No shut

R5(config)#Int F0/1.45
R5(config-subif)#Encap dot1q 45
R5(config-subif)#Ip addr 45.1.1.5 255.255.255.0
```

To verify the configuration:

On R4:

```
R4#Ping 45.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 45.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Let's configure VLAN 100. We know that the following must be configured:

- The F0/0 interface of R9 must be configured in VLAN 100
- The F0/9 interface of SW1 must be configured in VLAN 100, this is the interface that R9's F0/0 interface is connected to
- R7's G0/0 must be configured as a sub-interface, since it is a member of multiple VLANs, VLAN 100, and VLAN 67.
- The interface of the switch to which R7 is connected to must also be configured as a trunk.
- Another sub-interface must be configured on R8.

On SW1:

```
SW1(config)#Int F0/9
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc vlan 100
SW1(config-if)#No shut
```

On R9:

```
R9(config)#Int F0/0
R9(config-if)#Ip addr 100.1.1.9 255.255.255.0
R9(config-if)#No shut
```

On R7:

```
R7(config)#Int G0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R7(config-if)#No shut

R7(config-if)#Int G0/0.100
R7(config-subif)#Encap dot1q 100
R7(config-subif)#Ip addr 100.1.1.7 255.255.255.0
```

On SW1:

```
SW1(config)#Int F0/7
SW1(config-if)#Swi tru encap dot1q
SW1(config-if)#Swi mode trunk
SW1(config-if)#No shu
```

On R8:

```
R8(config)#Int G0/0.100
R8(config-subif)#Encap dot1q 100
R8(config-subif)#Ip addr 100.1.1.8 255.255.255.0
```

To verify the configuration:

On R8:

```
R8#Ping 100.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.7, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms

R8#Ping 100.1.1.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.9, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
```

Let's look at the second to last VLAN which is VLAN 67. To configure this VLAN we must configure the following:

- The F0/0 interface of R6 should be configured as a sub-interface, because it is connected to two different VLANs, VLAN 67 and VLAN 56.
- The F0/6 interface of SW1 must be configured as a trunk; this is the interface to which R6's F0/0 interface is connected to.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

- VLAN 67 must be configured on SW1.
- Another sub-interface must be configured on R7 for VLAN 67.

On R6:

```
R6(config)#Int F0/0
R6(config-if)#No shut

R6(config)#Int F0/0.67
R6(config-subif)#Encap dot1q 67
R6(config-subif)#Ip addr 67.1.1.6 255.255.255.0
```

On SW1:

```
SW1(config)#Int F0/6
SW1(config-if)#Swi trunk encap dot1q
SW1(config-if)#Swi mode trunk
SW1(config-if)#No shut

SW1(config)#VLAN 67
SW1(config-vlan)#Exit
```

On R7:

```
R7(config)#Int G0/0.67
R7(config-subif)#Encap dot1q 67
R7(config-subif)#Ip addr 67.1.1.7 255.255.255.0
```

To test and verify the configuration:

On R7:

```
R7#Ping 67.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 67.1.1.6, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

NOW, let's configure the last VLAN in this topology, VLAN 56.

- In this case we can see that R5 is using its F0/1 and R6 is using its F0/0 interface, this means that they are connected to two different switches, therefore, a trunk must be configured to connect these two switches and the trunk must allow the VLAN to traverse through this trunk.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

- A sub-interface must be configured on R5 for this VLAN
- A sub-interface must be configured on R6 for this VLAN
- VLAN 56 must be configured on BOTH SWITCHES, or VTP messages must be configured to propagate the VLAN.

On SW1:

```
SW1(config)#Vlan 56
SW1(config-vlan)#exit
```

On SW2:

```
SW2(config)#Vlan 56
SW2(config-vlan)#exit
```

To configure a trunk link between SW1 and SW2. In this case the F0/18 interfaces of these two switches are configured as trunk.

On SW1 and SW2:

```
SWx(config)#Int F0/18
SWx(config-if)#Swi tru enc dot
SWx(config-if)#Swi mode trunk
SWx(config-if)#No shu
```

On R5:

```
R5(config)#Int F0/1.56
R5(config-subif)#Encap dot 56
R5(config-subif)#Ip addr 56.1.1.5 255.255.255.0
```

On R6:

```
R6(config)#Int F0/0.56
R6(config-subif)#Encap dot 56
R6(config-subif)#Ip addr 56.1.1.6 255.255.255.0
```

To verify and test the configuration

On SW1:

```
SW1#Show inter F0/18 trunk

Port        Mode         Encapsulation  Status        Native vlan
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/18      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/18      1-4094

Port        Vlans allowed and active in management domain
Fa0/18      1,12-13,24,28,56,67,100

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/18      1,12-13,24,28,56,67,100
```

On SW2:

```
SW2#Show interface f0/18 trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/18      on           802.1q         trunking      1

Port        Vlans allowed on trunk
Fa0/18      1-4094

Port        Vlans allowed and active in management domain
Fa0/18      1,34,45,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/18      1,34,45,56
```

On R5:

```
R5#Ping 56.1.1.6

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

Task 3

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

LAB 3-
Physical to Logical Topology - III

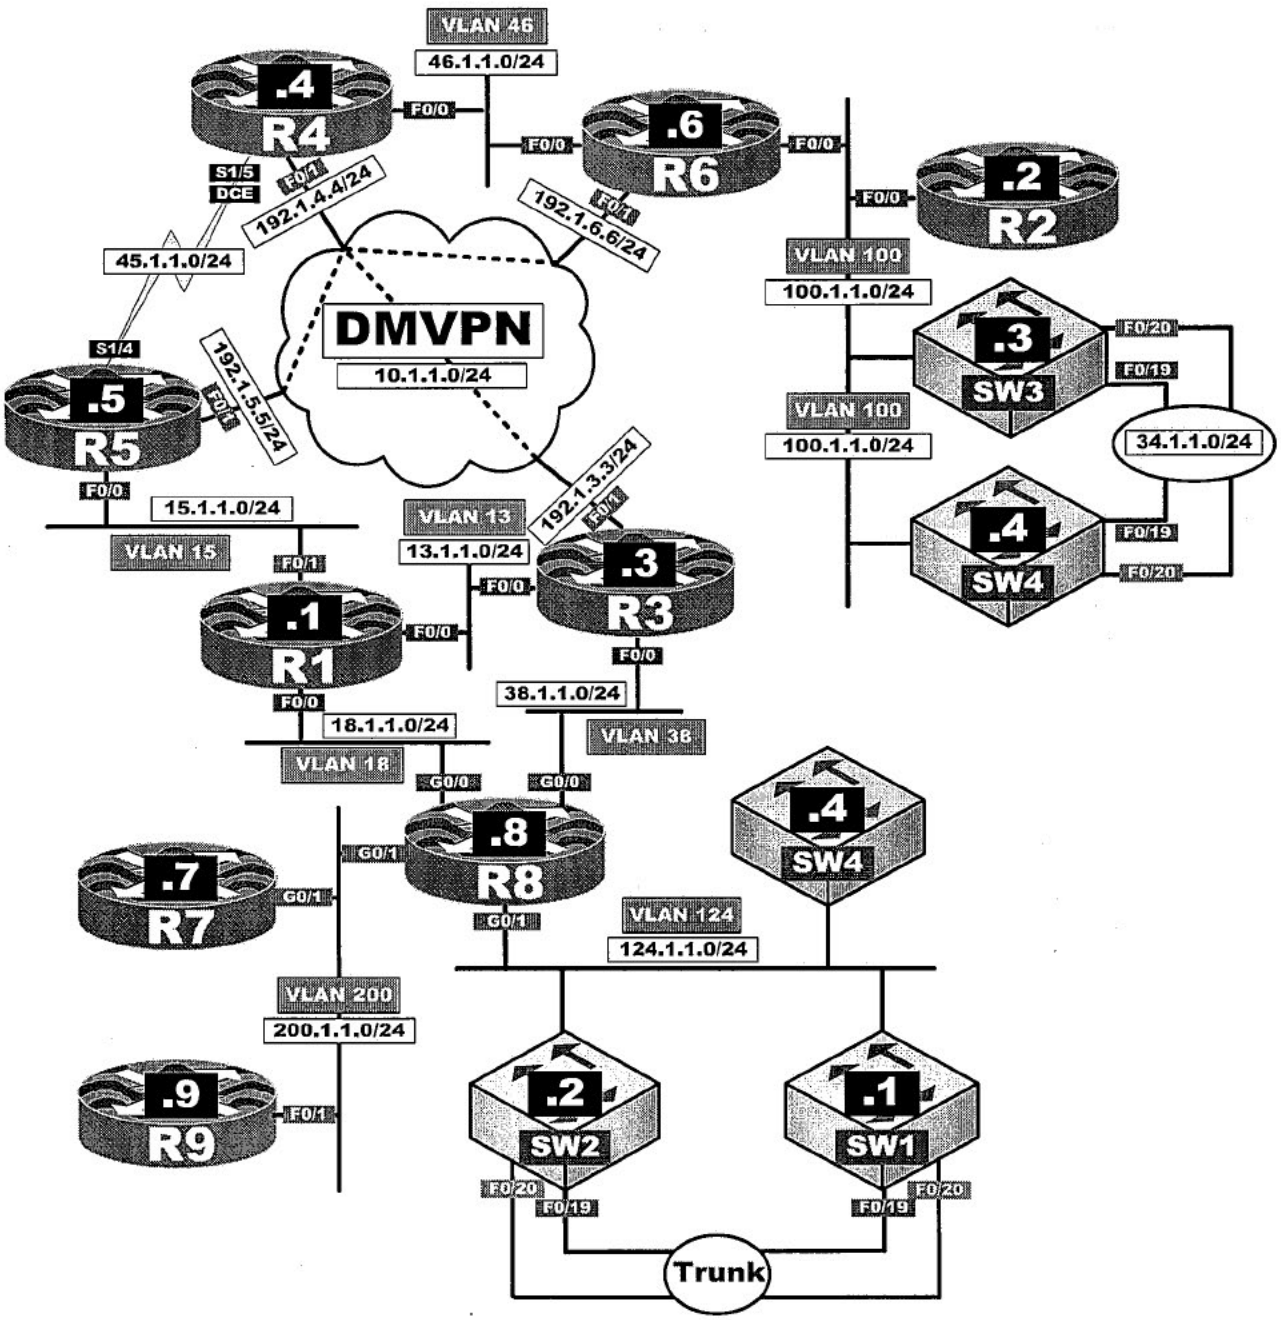

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Configure the above topology, if this configuration is performed successfully, every router should be able to ping its neighboring router/s in the same subnet.

The DMVPN network must be configured based on the following policy:

- R4 should be configured as the hub router.
- R3, R5 and R6 should be configured in a point-to-point manner.
- DO NOT configure any NHRP mapping on the hub.
- You should allow EIGRP over the tunnel interfaces using 224.0.0.9.

To configure this topology we can start from the top, or buttom. In this case let's configure the DMVPN network first.
Based on the above diagram we can see that R3 – R6 are all using their F0/1 interfaces, this means that the they are connected to SW2, since the lab did not specify an IP address on the switch, we will use the same network with the host portion of ".10", and the routers should have a default route pointing to that IP address; in this topology the switch is representing the Internet.
Once the routers and the switch are configured and they have full reachability to each other's IP address, we will then configure the tunnel interfaces using the 10.1.1.0/24 IP address.

On R3:

```
R3(config)#int f0/1
R3(config-if)#ip addr 192.1.3.3 255.255.255.0
R3(config-if)#No shu

R3(config)#ip route 0.0.0.0 0.0.0.0 192.1.3.10
```

On R4:

```
R4(config)#int f0/1
R4(config-if)#ip addr 192.1.4.4 255.255.255.0
R4(config-if)#No shu

R4(config)#ip route 0.0.0.0 0.0.0.0 192.1.4.10
```

On R5:

```
R5(config)#int f0/1
R5(config-if)#ip addr 192.1.5.5 255.255.255.0
R5(config-if)#No shu

R5(config)#ip route 0.0.0.0 0.0.0.0 192.1.5.10
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R6:

```
R6(config)#int f0/1
R6(config-if)#ip addr 192.1.6.6 255.255.255.0
R6(config-if)#No shu

R6(config)#ip route 0.0.0.0 0.0.0.0 192.1.6.10
```

On SW2:

```
SW2(config)#int f0/3
SW2(config-if)#No swi
SW2(config-if)#ip addr 192.1.3.10 255.255.255.0
SW2(config-if)#No shut

SW2(config)#int f0/4
SW2(config-if)#No swi
SW2(config-if)#IP addr 192.1.4.10 255.255.255.0
SW2(config-if)#No shut

SW2(config)#Int f0/5
SW2(config-if)#No swi
SW2(config-if)#ip addr 192.1.5.10 255.255.255.0
SW2(config-if)#No shut

SW2(config)#Int f0/6
SW2(config-if)#No swi
SW2(config-if)#ip addr 192.1.6.10 255.255.255.0
SW2(config-if)#No shut

SW2(config)#ip routing
```

To verify and test the configuration:

On R3:

```
R3#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

R3#Ping 192.1.5.5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R3#Ping 192.1.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R4:

```
R4#Ping 192.1.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R4#Ping 192.1.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R5:

```
R5#Ping 192.1.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Let's configure the tunnel interfaces:

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 F0/1
R4(config-if)#tunnel mode gre multipoint
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config-if)#IP nhrp network-id 444
R4(config-if)#ip nhrp map multicast dynamic
```

On R5:

```
R5(config)#Int tunnel 1
R5(config-if)#IP addr 10.1.1.5 255.255.255.0
R5(config-if)#tunnel source F0/1
R5(config-if)#tunnel destination 192.1.4.4
R5(config-if)#IP nhrp network-id 555
R5(config-if)#ip nhrp map 10.1.1.4 192.1.4.4
R5(config-if)#ip nhrp nhs 10.1.1.4
```

On R6:

```
R6(config)#int tunnel 1
R6(config-if)#IP addr 10.1.1.6 255.255.255.0
R6(config-if)#tunnel source f0/1
R6(config-if)#tunnel destination 192.1.4.4
R6(config-if)#ip nhrp network-id 666
R6(config-if)#ip nhrp map 10.1.1.4 192.1.4.4
R6(config-if)#ip nhrp nhs 10.1.1.4
```

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 f0/1
R3(config-if)#tunnel destination 192.1.4.4
R3(config-if)#ip nhrp network-id 333
R3(config-if)#ip nhrp map 10.1.1.4 192.1.4.4
R3(config-if)#ip nhrp nhs 10.1.1.4
```

To verify and test the configuration:

On R4:

```
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 = 1/3/4 ms

R4#Ping 10.1.1.5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R4#Ping 10.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

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#Ping 10.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R3#Ping 10.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R5:

```
R5#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

R5#Ping 10.1.1.4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R5#Ping 10.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

On R6:

```
R6#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

R6#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

R6#Ping 10.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Next, we should configure the serial connection that connects R4 to R5. When configuring a back to back serial connection it is VERY important to find out which router has the DCE end of the cable, the reason this is important is because the clocking, the "Clock rate" command MUST be configured on the router that has the DCE end of the cable connected to it. Let's find out:

On R4:

```
R4#Show controller S0/1 | Inc clock
DCE V.35, clock rate 2000000
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R5:

```
R5#Show controller S0/1/0 | Inc clock
DTE V.35 clocks stopped.
```

We can see that R4 has the DCE end of the cable, therefore, R4 should be configured with the clock rate:

On R4:

```
R4(config)#Int S1/5
R4(config-if)#Clock rate 64000
R4(config-if)#Ip addr 45.1.1.4 255.255.255.0
R4(config-if)#No shut
```

On R5:

```
R5(config)#Int S1/4
R5(config-if)#Ip address 45.1.1.5 255.255.255.0
R5(config-if)#No shut
```

To verify the configuration:

On R4:

```
R4#Ping 45.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 45.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
```

Since the serial connection and the DMVPN networks are configured, let's move on to the VLANs, let's start from the top and configure R4 and R6 for VLAN 46. Before we start this configuration, we can see that the F0/0 interface of R6 is configured in two different VLANs; therefore, we know that it must be configured with two sub-interfaces, one in VLAN 46 and the second one in VLAN 26. Since based on the physical diagram, the F0/0 interfaces of R4 and R6 are connected to SW1's F0/4 and F0/6 respectively, we should start with SW1's configuration and then configure the routers:

On SW1:

```
SW1(config)#Int F0/4
SW1(config-if)#Switchport mode access
SW1(config-if)#Switchport access vlan 46
% Access VLAN does not exist. Creating vlan 46
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1(config-if)#No shut
```

Since VLAN 46 did not exist, once the interface was configured in VLAN 46, the switch created that VLAN automatically.

The following configures the F0/6 interface of SW1 as a trunk:

```
SW1(config)#Int F0/6
SW1(config-if)#Switchport trunk encapsulation dot1q
SW1(config-if)#Switchport mode trunk
SW1(config-if)#No shut
```

Now, let's configure R6:

On R6:

```
R6(config)#Int F0/0
R6(config-if)#No shut

R6(config)#Int F0/0.46
R6(config-subif)#Encapsulation dot1q 46
R6(config-subif)#Ip addr 46.1.1.6 255.255.255.0
```

Let's configure R4:

On R4:

```
R4(config)#Int F0/0
R4(config-if)#Ip addr 46.1.1.4 255.255.255.0
R4(config-if)#No shut
```

To verify the configuration:

On R6:

```
R6#Ping 46.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 46.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

The next VLAN is VLAN 100:
To configure this VLAN, another sub-interface must be configured on R6.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R2's F0/0 interface should be configured in this VLAN.

We can see that SW3 and SW4 are also in this VLAN. The question is which interface of these switches should be configured in this VLAN?
Since these switches are NOT connected to another switch that interconnects these devices, the ONLY choice here is an SVI configuration on SW3 and SW4.
To configure an SVI, the VLAN that the SVI references must exist on the switch and the local switch should also have an interface in this VLAN, in this case this interface should be a trunk interface that connects SW3 and SW4 to SW1 so they can have a connection to each other.

The last connection in this VLAN is the Etherchannel connection between SW3 and SW4.

Let's start with SW1's configuration:

On SW1:

```
SW1(config)#Int F0/2
SW1(config-if)#Switchport mode access
SW1(config-if)#Switchport access vlan 100
% Access VLAN does not exist. Creating vlan 100
SW1(config-if)#No shut
```

Let's configure the SVIs on SW3 and SW3, and the trunk links between SW1 and SW3, and another one that connects SW1 to SW4:

On SW3:

```
SW3(config)#Int VLAN 100
SW3(config-if)#Ip addr 100.1.1.3 255.255.255.0

SW3(config)#Vlan 100
SW3(config-vlan)#Exit
```

On SW4:

```
SW4(config)#VLAN 100
SW4(config-vlan)#Exit

SW4(config)#Int VLAN 100
SW4(config-if)#IP addr 100.1.1.4 255.255.255.0
```

Let's configure the trunk links between SW1 and SW3, and SW1 and SW4:

On SW1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
SW1(config)#Int range F0/21,f0/23
SW1(config-if-range)#Swi trunk enc dot1q
SW1(config-if-range)#Swi mode trunk
SW1(config-if-range)#No shut
```

On SW3:

```
SW3(config)#Int F0/21
SW3(config-if)#swi tru enc dot
SW3(config-if)#swi mode tr
SW3(config-if)#No shut
```

On SW4:

```
SW4(config)#Int F0/23
SW4(config-if)#swi trunk enc dot
SW4(config-if)#swi mode trunk
SW4(config-if)#No shut
```

On R6:

```
R6(config)#Int F0/0.100
R6(config-subif)#Encapsulation dot1q 100
R6(config-subif)#Ip addr 100.1.1.6 255.255.255.0
```

On R2:

```
R2(config)#Int F0/0
R2(config-if)#Ip addr 100.1.1.2 255.255.255.0
R2(config-if)#No shut
```

To verify the configuration:

On SW4:

```
SW4#Ping 100.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/200/1000 ms

SW4#Ping 100.1.1.6

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 100.1.1.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/201/1000 ms

SW4#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 = 1/1/1 ms
```

That is perfect. Let's configure the Etherchannel between SW3 and SW4. By looking at he physical diagram in the beginning of this book, these two switches are connected to each other via their F0/19 and F0/20 interfaces, so let's configure the EtherChannel using these two ports:

On SW3 and SW4:

To configure a layer three EtherChannel you should go through the following process:

1- "Default interface" the interfaces that will be used for the EtherChannel.
2- Configure the interface port-channel
3- Configure the port-channel as a layer three interface by configuring the "No Switchport" command and configure the IP address.
4- Configure the physical ports as layer three as well, using the "No Switchport" command.
5- Assign the port-channel ID to the interfaces using the "Channel-group" interface command.

On SW3 and SW4:

```
SWx(config)#Default interface F0/19
SWx(config)#Default interface F0/20

SWx(config)#Int port-channel 34
SWx(config-if)#No switchport
```

On SW3:

```
SW3(config)#Int port-channel 34
SW3(config-if)#IP address 34.1.1.3 255.255.255.0
```

On SW4:

```
SW4(config)#Int port-channel 34
SW4(config-if)#IP address 34.1.1.4 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW3 and SW4:

```
SWx(config)#int range f0/19-20
SWx(config-if-range)#No switch
SWx(config-if-range)#Channel-group 34 mode on
SWx(config-if-range)#No shut
```

To verify the configuration:

On SW3:

```
SW3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/4 ms
```

The next VLAN is VLAN 15. We can see that R1's F0/1 interface should be configured in this VLAN, but R1's F0/1 interface is connected to SW2, therefore, the Switchport to which R1's F0/1 interface is connected to must be configured in VLAN 15, and a trunk link MUST be configured to connect SW1 to SW2 or else these two routers will NOT be able to commnunicate.

On SW1:

```
SW1(config)#Int F0/5
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc v 15
% Access VLAN does not exist. Creating vlan 15
SW1(config-if)#No shut
```

The F0/18 interfaces of SW1 and SW2 are configured as a trunk link, you can use any interface that connects these two switches:

On SW1 and SW2:

```
SWx(config)#Int F0/18
SWx(config-if)#Swi tru enc dot
SWx(config-if)#Swi mode trunk
SWx(config-if)#No shut
```

The following configures the F0/1 interface of SW2 in VLAN 15, and VLAN 15 is statically configured on SW2:

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

On SW2:

```
SW2(config)#Int F0/1
SW2(config-if)#Swi mode acc
SW2(config-if)#Swi acc v 15
% Access VLAN does not exist. Creating vlan 15
SW2(config-if)#No shu
```

Last step is to configure R1's F0/1 and R5's F0/0 interfaces:

On R1:

```
R1(config)#Int F0/1
R1(config-if)#Ip addr 15.1.1.1 255.255.255.0
R1(config-if)#No shu
```

On R5:

```
R5(config)#Int F0/0
R5(config-if)#Ip addr 15.1.1.5 255.255.255.0
R5(config-if)#No shut
```

To verify the configuration:

On R1:

```
R1#Ping 15.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 15.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

The next VLAN is 13. We can see that R1 and R3's F0/0 interfaces are configured in multiple VLANs, therefore, the switchport that these routers are connected to must be configured as a trunk link, and the VLAN must be configured statically:

On SW1:

```
SW1(config)#Int range F0/1, F0/3
SW1(config-if)#Swi tru enc dot
SW1(config-if)#Swi mode trunk
SW1(config-if)#No shu
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

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

R1(config)#Int F0/0.13
R1(config-subif)#Encap dot1 13
R1(config-subif)#Ip addr 13.1.1.1 255.255.255.0
```

On R3:

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

R3(config)#Int F0/0.13
R3(config-subif)#Encap dot1q 13
R3(config-subif)#Ip addr 13.1.1.3 255.255.255.0
```

On SW1:

```
SW1(config)#VLAN 13
SW1(config-vlan)#Exit
```

To verify the configuration:

On R1:

```
R1#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 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Next VLAN is 38. Perform the following steps to configure this VLAN:

- Configure a sub-interface on R8
- Configure another sub-interface on R3
- Configure VLAN 38 on SW1
- Configure the F0/8 interface of SW1 as a trunk

On R8:

```
R8(config)#Int G0/0
```


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

```
R8(config-if)#No shut

R8(config)#Int G0/0.38
R8(config-subif)#Encap dot1q 38
R8(config-subif)#Ip addr 38.1.1.8 255.255.255.0
```

On R3:

```
R3(config)#Int F0/0.38
R3(config-subif)#Encap dot1q 38
R3(config-subif)#Ip addr 38.1.1.3 255.255.255.0
```

On SW1:

```
SW1(config)#VLAN 38
SW1(config-vlan)#Exit

SW1(config)#Int F0/8
SW1(config-if)#Swi trunk enc dot
SW1(config-if)#Swi mode tru
SW1(config-if)#No shut
```

To verify the configuration:

On R8:

```
R8#Ping 38.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 38.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

The next VLAN is 18. Perform the following steps to configure this VLAN:

- Configure another sub-interface on R1
- Configure another sub-interface on R8
- Configure VLAN 18 on SW1

On R1:

```
R1(config)#Int F0/0.18
R1(config-subif)#Encap dot1q 18
R1(config-subif)#Ip addr 18.1.1.1 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R8:

```
R8(config)#Int G0/0.18
R8(config-subif)#Encap dot 18
R8(config-subif)#Ip addr 18.1.1.8 255.255.255.0
```

On SW1:

```
SW1(config)#VLAN 18
SW1(config-vlan)#Exit
```

To verify the configuration:

On R8:

```
R8#Ping 18.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 18.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Let's configure the next VLAN, VLAN 200. Perform the following steps to configure this VLAN:

- The G0/1 interface of R8 must be configured with a sub-interface
- The F0/8 interface of SW3 must be configured as a trunk
- On SW2, the switchport that the G0/1 interface of R7 is connected to must be configured in this VLAN.
- On SW3, the switchport that the F0/1 interface of R9 is connected to must be configured in this VLAN.
- SW3 and SW2 MUST have a trunk link so these devices can communicate.

On R8:

```
R8(config)#Int G0/1
R8(config-if)#No shut

R8(config)#Int G0/1.200
R8(config-subif)#Encap dot 200
R8(config-subif)#Ip addr 200.1.1.8 255.255.255.0
```

On SW3:

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

```
SW3(config)#Int F0/8
SW3(config-if)#Swi trunk enc dot
SW3(config-if)#Swi mode tru
SW3(config-if)#No shut
```

On SW2:

```
SW2(config)#Int F0/7
SW2(config-if)#Swi mode acc
SW2(config-if)#Swi acc v 200
% Access VLAN does not exist. Creating vlan 200
SW2(config-if)#No shut
```

On SW3:

```
SW3(config)#Int F0/9
SW3(config-if)#Swi mode acc
SW3(config-if)#Swi acc v 200
% Access VLAN does not exist. Creating vlan 200
SW3(config-if)#No shu
```

On SW2 and SW3:

```
SWx(config)#Int F0/23
SWx(config-if)#Swi tru enc do
SWx(config-if)#Swi mode tr
SWx(config-if)#No shu
```

On R9:

```
R9(config)#Int F0/1
R9(config-if)#Ip addr 200.1.1.9 255.255.255.0
R9(config-if)#No shut
```

On R7:

```
R7(config)#Int G0/1
R7(config-if)#Ip addr 200.1.1.7 255.255.255.0
R7(config-if)#No shut
```

To verify the configuration:

On R9:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

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

The last and the final VLAN is VLAN 124. Perform the following steps to configure this VLAN:

- Configure another sub-interface on R8
- Configure three SVIs, one on SW1, one on SW2 and another one on SW4
- Configure a trunk between SW2 and SW4 using their F0/19, and F0/20 interfaces
- Configure VLAN 124 on SW1, SW2 and SW4

Finally, the last configuration is the layer two EtherChannel between SW1 and SW2.

On R8:

```
R8(config)#Int G0/1.124
R8(config-subif)#Encap dot 124
R8(config-subif)#Ip addr 124.1.1.8 255.255.255.0
```

On SW1:

```
SW1(config)#Int Vlan 124
SW1(config-if)#Ip addr 124.1.1.1 255.255.255.0
```

On SW2:

```
SW2(config)#Int Vlan 124
SW2(config-if)#Ip addr 124.1.1.2 255.255.255.0
```

On SW4:

```
SW4(config)#Int Vlan 124
SW4(config-if)#Ip addr 124.1.1.4 255.255.255.0
```


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

On SW4:

```
SW4(config)#Int F0/21
SW4(config-if)#Swi tru enc dot
SW4(config-if)#Swi mode trunk
SW4(config-if)#No shut
```

On SW2:

```
SW2(config)#Int F0/21
SW2(config-if)#Swi tru enc dot
SW2(config-if)#Swi mode tr
SW2(config-if)#No shut
```

On SW1, SW2, SW3 and SW4:

```
SWx(config)#VLAN 124
SWx(config-vlan)#Exit
```

To verify the configuration:

On SW1:

```
SW1#Ping 124.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 124.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

SW1#Ping 124.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 124.1.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

SW1#Ping 124.1.1.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 124.1.1.8, 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

Perform the following steps to configure the layer two EtherChannel:

- Default interfaces F0/19 and F0/20
- Assign a Channel-group to both interfaces
- Configure the Port-channel as a trunk

On SW1 and SW2:

```
SWx(config)#Default inter F0/19
SWx(config)#Default inter F0/20

SWx(config)#Int Range F0/19-20
SWx(config-if-range)#Channel-group 12 mode on
```

You should get the following console message stating that the port-channel 12 is being created:

```
Creating a port-channel interface Port-channel 12
%LINK-3-UPDOWN: Interface Port-channel12, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel12, changed
state to up

SWx(config)#Int port-channel 12
SWx(config-if)#Swi trunk encap dot
SWx(config-if)#Swi mode trunk
```

To verify the configuration:

On SW1:

```
SW1#Show interface Po12 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        1,13,15,18,38,46,100,124

Port        Vlans in spanning tree forwarding state and not pruned
Po12        1,13,15,18,38,46,100,124
```


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

Task 2

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

CCIE Foundation

5.0

www.MicronicsTraining.com

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

3560

Switching

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

Lab 1
Basic 3560 Configuration I

Task 1

Shutdown all ports on all four Switchs except ports F0/19 and F0/20 on SW1 and SW2.

On All Switches:

```
SWx(config)#int range F0/1-24
SWx(config-if-range)#Shut
```

On SW1 and SW2:

```
Switch(config)#int range F0/19-20
Switch(config-if-range)#No Shut
```

To verify the configuration:

On Both Switches:

```
Switch#Show int 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		connected	1	a-full	a-100	10/100BaseTX

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Fa0/20                    connected    1          a-full  a-100 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
```

(The rest of the output is omitted)

Task 2

Configure Switch 1 to be in VTP domain called "CCIE"; this information should be propagated to Switch 2 via VTP messages. You can use any encapsulation or tagging to accomplish this task.

Before assigning a VTP domain name, there must be a trunk established between the two switches so the configurations can be propagated to the other switch.

On both switches:

```
Switch#Show interface trunk

Switch#
```

NOTE: The two 3560s switches are connected with 2 cross over ethernet cables, if one or both of these switches were 3550s the two ports would have negotiated an ISL trunk, actually, in the output of the "Show interface trunk" the connection will be displayed as "n-isl", because by default the ports are configured in dynamic desirable mode. With 3560 switches, the ports are not in dynamic desirable mode, a "Show interface F0/19 switchport" will reveal this information. By default, the ports in 3560 switches are configured in "Dynamic Auto" mode, and therefore, the port(s) must be configured statically to trunk or negotiate a trunk.

On Both switches:

```
Switch#Show cdp neighbors | B Device

Device ID          Local Intrfce       Holdtme    Capability     Platform   Port ID
Switch             Fas 0/20            126            S I      WS-C3560-2Fas 0/20
Switch             Fas 0/19            126            S I      WS-C3560-2Fas 0/19
```

Note the output of the above show command reveals the ports that connect the two switches to each other.

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

On Both switches:

```
Switch(config)#int range f0/19-20
Switch(config-if-range)#switchport trunk encapsulation isl
Switch(config-if-range)#switchport mode trunk
```

To verify the configuration:

On the first switches:

```
Switch#Show int trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           isl            trunking      1
Fa0/20      on           isl            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
```

Now that the trunk is established between the two switches, the next prerequisite for this task to be completed is the VTP domain name, the switches MUST be in the same VTP domain name, this means that you MUST use the "VTP domain" command in the VLAN database or in the global configuration mode:

On the first switch:

```
Switch(config)#VTP domain CCIE
```

By default, the 3560 switches are members of a domain called "NULL", therefore, after entering the above command, you should get the following console message unless the switch was a member of another domain:

Changing VTP domain name from NULL to CCIE

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

This task could also be accomplished within the "VLAN database" as follows:

```
Switch#Vlan database
Switch(vlan)#Vtp domain CCIE
Switch(vlan)#Exit
```

When any configuration is performed in the Vlan database, you must enter "exit" or the "apply" command for the changes to take effect.

Note the output of the following show command reveals that VTP messages propagated the VTP domain information to the second switch:

On the second switch:

```
Switch#Sh vtp status

VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 5
VTP Operating Mode              : Server
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
Configuration last modified by 0.0.0.0 at 3-1-93 04:31:53
Local updater ID is 0.0.0.0 (no valid interface found)
```

Task 3

The VTP domain configured in the previous step should be password protected using "Cisco" as the password.

On both switches:

```
Switch(config)#VTP password Cisco
```

You should get the following console message:

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

```
Setting device VLAN database password to Cisco
```

NOTE: If a domain name is not assigned to the switches and the default VTP domain name (NULL) is used, a password can not be assigned and you should see the following console error message:

```
%The VTP password cannot be set for NULL domain
```

The "VTP password" command can be entered in global configuration mode, privilege configuration mode or in the VLAN database mode.

The password command must be configured statically on both switches, because this change will NOT get propagated via VTP messages.

To verify the configuration:

On the First switch:

```
Switch#Show vtp password
```

This verifies the password, remember Spaces will not show

```
VTP Password: Cisco
```

On the Second switch:

```
Switch#Show VTP password

VTP Password: Cisco
```

The VTP password can be changed in three ways:

1. Privilege mode:

```
Switch#vtp password Cisco
```

2. Vlan Database:

```
Switch#Vlan database
Switch(vlan)#Vtp password Cisco
Switch#Exit
```

3. Global config mode:

```
Switch(config)#vtp password Cisco
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 4

The first Catalyst switch should be configured with a hostname of Cat-1 and the second Catalyst should have a hostname of Cat-2.

On the first Switch:

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

On the Second Switch:

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

Task 5

Cat-2 should NOT have the ability to create, delete or rename VLAN(s) or any VLAN information.

On Cat-2:

```
Cat-2(config)#Vtp mode client
```

This configuration can be performed in the vlan database or global config mode. The above command displays the command as it was entered in the global config mode. If you are asked to enter the command in the vlan database, you must first enter the "Vlan database" command in the privilege mode, then, enter "vtp client" and lastly the "apply" or the "exit" command is entered so the changes can take effect. Once the command is entered you should get the following console message:

```
Setting device to VTP CLIENT mode.
```

VTP Modes:

The switches can operate in three VTP modes and they are as follows:

- SERVER – The switch is able to delete, create, or rename VLAN information. A Catalyst 3560 switch that is configured in server mode participates in the VTP domain and propagates the VLAN information.

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

- CLIENT – In this mode the 3560 switches are able to receive and process the VTP messages, but they are not able to create, delete, or rename VLAN information. They can assign a port to a given VLAN that already exists. Catalyst 3560 in client mode participates in the VTP domain and propagates the VTP messages.

- Transparent – In this mode the switch is able to create, delete and modify the VLAN information but it will not propagate its VLAN information to other switches. Catalyst 3560 switches in this mode DO NOT participate in VTP domain. A Catalyst 3560 switch must be in this mode in order to create the extended-range VLANs (1006 – 4094), this configuration can only be performed in the global config mode.

Task 6

Create and configure the following VLAN assignments on the switches:

Router Interface	VLAN number	CAT Switches Port
R1 – F0/0	12	Cat-1 / F0/1
R2 – F0/0	12	Cat-1 / F0/2
R3 – F0/0	34	Cat-1 / F0/3
R4 – F0/0	34	Cat-1 / F0/4
R5 – F0/0	56	Cat-1 / F0/5
R6 – F0/0	56	Cat-1 / F0/6

On Cat-1:

```
Cat-1(config)#interface range F0/1 - 2
Cat-1(config-if)#switch mode access
Cat-1(config-if)#switch access vlan 12

Cat-1(config)#interface range F0/3 - 4
Cat-1(config-if)#switch mode access
Cat-1(config-if)#switch access vlan 34

Cat-1(config)#interface range F0/5 - 6
Cat-1(config-if)#switch mode access
Cat-1(config-if)#switch access vlan 56
```

Since both switches are in the same VTP domain and they are configured with the same password, the

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Vlan information should be propagated to the other switch (Cat-2).

On Cat-2:

```
Cat-2#Show vlan brie | 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

Cat-2#Show VTP Status

VTP Version                     : 2
Configuration Revision          : 3
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xB0 0x08 0xD6 0x80 0x12 0xF2 0xB4 0xC2
Configuration last modified by 0.0.0.0 at 3-1-93 01:12:17
```

On Cat-1:

```
Cat-1#Show VTP Status

VTP Version                     : 2
Configuration Revision          : 3
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Server
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xB0 0x08 0xD6 0x80 0x12 0xF2 0xB4 0xC2
Configuration last modified by 0.0.0.0 at 3-1-93 01:12:17
```


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

```
Local updater ID is 0.0.0.0 (no valid interface found)
```

Note, the VTP version is 1, Configuration revision is 3; the number of existing VLANs is 8 on both switches, (because they are synchronized), and the reason the VLAN information was propagated is because the VTP domain name and the password is identical on both switches and the switches are trunked.

Task 7

Configure Loopback 0 and Loopback 1 interfaces on Cat-1, use the IP addresses of 1.1.1.1 /8 and 11.1.1.1 /8 respectively and ensure that ONLY the IP address of Loopback 1 interface is used as the preferred source for the VTP IP updater address.

Note in the previous Task when the "show vtp status" command was entered on Cat-1, the last line of the output displayed "no valid interface found". Catalyst switches will use the IP address of the lowest physical interface number, if that interface does not have an IP address, then loopback 0 interface will be used as the source of all VTP messages, but this behavior can be changed by using the "VTP interface Loopback 1" global config command.

On Cat-1:

```
Cat-1(config)# Interface Loopback 0
Cat-1(config-if)# Ip address 1.1.1.1 255.0.0.0

Cat-1(config)# Interface Loopback 1
Cat-1(config-if)# Ip address 11.1.1.1 255.0.0.0

Cat-1#Show vtp status

VTP Version                     : 2
Configuration Revision          : 3
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Server
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xB0 0x08 0xD6 0x80 0x12 0xF2 0xB4 0xC2
Configuration last modified by 0.0.0.0 at 3-1-93 01:12:17
Local updater ID is 1.1.1.1 on interface Lo0 (first layer3 interface found)
```

NOTE: Loopback 0 is used as the source of all VTP messages. The following command changes the source to

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

its Loopback 1 interface:

```
Cat-1(config)#Vtp interface Loopback1 ONLY
```

The "ONLY" argument makes this interface mandatory. YOU MUST TYPE "LOOPBACK1" OR "LO1", OR ELSE IT WILL NOT WORK. IF "L1" IS ENTERED INSTEAD OF "LOOPBACK1", the IOS will take "L1", but it WILL NOT WORK.

To verify the configuration:

On Cat-1:

```
Cat-1#Show vtp status

VTP Version                     : 2
Configuration Revision          : 3
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Server
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xB0 0x08 0xD6 0x80 0x12 0xF2 0xB4 0xC2
Configuration last modified by 0.0.0.0 at 3-1-93 01:12:17
Local updater ID is 11.1.1.1 on interface Lo1 (preferred interface)
Preferred interface name is loopback1 (mandatory)
```

On Cat-2

```
Cat-2#Show vtp status

VTP Version                     : 2
Configuration Revision          : 18
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xB0 0x08 0xD6 0x80 0x12 0xF2 0xB4 0xC2
Configuration last modified by 0.0.0.0 at 3-1-93 01:12:17
```

Note this change has not been propagated, therefore, to force the propagation of this change, a VLAN is created, in this case "VLAN 80", so you can see that the change was made by the "Loopback 1" interface with an IP address of "11.1.1.1" on Cat-2. This VLAN should be deleted before proceeding to the next task.

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

On Cat-1:

```
Cat-1(config)#Vlan 80
Cat-1(config-vlan)#Exit
```

To verify the configuration:

On Cat-2:

```
Cat-2#Show vtp status

VTP Version                     : 2
Configuration Revision          : 4
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 9
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xC4 0x80 0x95 0xD5 0xC2 0x40 0x5C 0x2A
Configuration last modified by 11.1.1.1 at 3-1-93 01:24:34
```

Depending on the IOS, the above highlighted area may or may not show in the output.

To remove the VLAN:

On Cat-1:

```
Cat-1(config)#No vlan 80
```

To verify the configuration

On Cat-2:

```
Cat-2#Show vtp status

VTP Version                     : 2
Configuration Revision          : 5 → The version is incremented by one
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
VTP Traps Generation            : Disabled
MD5 digest                      : 0x97 0xC9 0x03 0x9B 0x50 0x06 0x2B 0xA2
Configuration last modified by 11.1.1.1 at 3-1-93 00:25:47

Cat-2#Show 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
12   VLAN0012                             active
34   VLAN0034                             active
56   VLAN0056                             active
```

NOTE: The VLAN "Vlan 80" is no longer there.

Task 8

Configure the switches such that flooded traffic is restricted to the trunk links that the traffic must use to reach the destination device.

To see the default setting:

On Cat-2:

```
Cat-2#Show vtp status

VTP Version                     : 2
Configuration Revision          : 5
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x72 0x36 0x88 0xF7 0xEB 0x4E 0x83 0x65
Configuration last modified by 0.0.0.0 at 3-1-93 00:40:51
```


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

VTP Pruning is disabled by default, enter the following command to enable VTP pruning:

On Cat-1:

```
Cat-1#VTP pruning
Pruning switched on
```

This command can be configured in privilege mode, Global config mode, and/or in the Vlan database. Once this feature is enabled, it will get propagated to the other switches within the VTP domain.

To verify the configuration on both switches:

On Cat-2:

```
Cat-2#Show vtp status

VTP Version                     : 2
Configuration Revision          : 6
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 8
VTP Operating Mode              : Client
VTP Domain Name                 : CCIE
VTP Pruning Mode                : Enabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0xDA 0xA4 0x1D 0xB6 0xD7 0x42 0x2A 0x0C
Configuration last modified by 0.0.0.0 at 3-1-93 00:45:22
```

NOTE: VTP messages propagate the change through the entire VTP domain.

Task 9

Configure Cat-1 and Cat-2 such that only the trunk ports (F0/19 and F0/20) and the ports that routers R1 to R6 are connected are in use, the rest of the ports should be configured in administratively down state.

On Both Switches:

```
Cat-x(config)#Int Range F0/1-6
Cat-x(config-if-range)#No Shut
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On Cat-1:

```
Cat-1#Show inter status | Exc disabled

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/1                        notconnect   12           auto   auto 10/100BaseTX
Fa0/2                        notconnect   12           auto   auto 10/100BaseTX
Fa0/3                        notconnect   34           auto   auto 10/100BaseTX
Fa0/4                        notconnect   34           auto   auto 10/100BaseTX
Fa0/5                        notconnect   56           auto   auto 10/100BaseTX
Fa0/6                        notconnect   56           auto   auto 10/100BaseTX
Fa0/19                       connected    trunk      a-full  a-100 10/100BaseTX
Fa0/20                       connected    trunk      a-full  a-100 10/100BaseTX
```

On Cat-2:

```
Cat-2#Show inter status | Exc disabled

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/1                        notconnect   1            auto   auto 10/100BaseTX
Fa0/2                        notconnect   1            auto   auto 10/100BaseTX
Fa0/3                        notconnect   1            auto   auto 10/100BaseTX
Fa0/4                        notconnect   1            auto   auto 10/100BaseTX
Fa0/5                        notconnect   1            auto   auto 10/100BaseTX
Fa0/6                        notconnect   1            auto   auto 10/100BaseTX
Fa0/19                       connected    trunk      a-full  a-100 10/100BaseTX
Fa0/20                       connected    trunk      a-full  a-100 10/100BaseTX
```

Task 10

Ensure that Cat-1 is the root bridge for the VLANs 12, 34 and Cat-2 is the root bridge for VLAN 56. Do NOT use the "**priority**" command to accomplish this task.

The following commands can be used to display the BID for a given switch:

- Show version | Inc MAC
- Show spanning-tree bridge

On Cat-1:

```
Cat-1#Show version | Inc MAC
```


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

```
Base ethernet MAC Address       : 00:12:7F:40:93:80
```

The following command reveals the base MAC address of the switch; The BID is a combination of priority and the base MAC address.

```
Cat-1#Show spanning-tree bridge

                                                    Hello  Max  Fwd
Vlan                    Bridge ID                    Time  Age  Dly  Protocol
---------------- ------------------------------ ------  ---  ---  --------
VLAN0001         32769 (32768,    1) 0012.7f40.9380    2    20   15  ieee
VLAN0012         32780 (32768,   12) 0012.7f40.9380    2    20   15  ieee
VLAN0034         32802 (32768,   34) 0012.7f40.9380    2    20   15  ieee
VLAN0056         32824 (32768,   56) 0012.7f40.9380    2    20   15  ieee
```

Note the priority starts with 32768, each VLAN that is created adds its VLAN number to the default priority value (If the base priority and the VLAN number is added within the parenthesis before the highlighted area, the sum will be the priority for that given VLAN, displayed to the left of the parenthisis); VLAN 12 adds 12 to the default priority value, therefore, the priority is 32780 and VLAN 34 adds 34 to the default priority value, therefore, the priority is 32802. Note that the MAC is the base MAC address and it remains the same, in this case (0012.7f40.9380).

NOTE: The MAC address of the switch that you are working on maybe different.

Enter the following command to reveal the BID and the root bridge for VLAN 12:

On Cat-1:

```
Cat-1#Show spanning-tree vlan 12
```

NOTE: The first section of the following output reveals information about the ROOT Bridge, obviously Cat-1 is NOT the root bridge:

```
VLAN0012
  Spanning tree enabled protocol ieee
  Root ID     Priority    32780
              Address     0012.7f40.9380
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

The second section of the following output reveals information about the LOCAL Bridge:

```
  Bridge ID  Priority    32780  (priority 32768 sys-id-ext 12)
             Address     0012.7f40.9380
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

The third section, reveals the state of the interface/s in this VLAN:

```
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/2            Desg FWD 19        128.4    P2p
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
```

Enter the following commands to configure Cat-1 to be the root bridge for VLANs 12 and 34:

On Cat-1:

```
Cat-1(config)#Spanning-tree vlan 12,34 root primary
```

The above command configures Cat-1 to be the root for VLANs 12 and 34; the "root" keyword is a macro that reduces the BID of the switch for a given VLAN(s) by a value of 8192 (The lower value is the preferred value). There are no spaces between the 12 and the comma and the 34.

```
Cat-1#Show spanning-tree vlan 12          Note 32768+ 12 - 8192 = 24588

VLAN0012
  Spanning tree enabled protocol ieee
  Root ID     Priority    24588
              Address     0012.7f40.9380
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24588  (priority 24576 sys-id-ext 12)
              Address     0012.7f40.9380
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Desg FWD 19        128.3    P2p
Fa0/2            Desg FWD 19        128.4    P2p
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
```

NOTE: Both interfaces are in FWD state.

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

On Cat-2:

```
Cat-2(config)##Spanning-tree vlan 56 root primary
```

To verify the configuration:

On Cat-2:

```
Cat-2#Show spanning vlan 56

VLAN0056
  Spanning tree enabled protocol ieee
  Root ID     Priority    24632
              Address     001d.e5d6.0000
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    24632  (priority 24576 sys-id-ext 56)
              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
```

Task 11

Cat-1 should be configured such that the ports that routers R1 to R6 are connected bypass listening and learning state. If any of these ports receive BPDU packets, they should transition into errdisable state. Use minimum number of commands to accomplish this task. This configuration should only be applied to the ports that the routers R1 - R6 are connected to as well as any future port that has this feature enabled. You should enable the F0/0 interfaces of R1 – R6.

On R1 – R6:

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


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-1:

```
Cat-1(config)#Spanning-tree portfast bpduguard default

Cat-1(config)#Interface range F0/1 - 6
Cat-1(config-if)#Spanning-tree portfast
```

You should see the following console messages:

```
%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 in 6 interfaces due to the range command
 but will only have effect when the interfaces are in a non-trunking
mode.
```

With the "spanning-tree portfast bpduguard default" command configured in global config mode, if any port that is portfast enabled receives BPDU packets, it will be shutdown in "err-disable" mode.

To verify the configuration:

On Cat-1:

```
Cat-1#Sh spanning-tree interface F0/1 portfast

VLAN0012            enabled
```

NOTE: If the output of the above show command states "no spanning tree info available for FastEthernet0/1", the F0/0 interface of R1 is may be in Shutdown mode.

To test the configuration:

On SW2:

```
Cat-2(config)#Spanning-tree portfast bpduguard default

Cat-2(config)#int f0/23
Cat-2(config-if)#swi mode acc
Cat-2(config-if)#spanning-tree portfast
Cat-2(config-if)#NO shut
```

Note if the F0/23 interface of Switch 3 is enabled, BPDUs will be generated and because of this configuration, F0/23 interface of Cat-2 will transition into "err-disable" mode, as follows:

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

On Cat-2 and Switch 3:

```
Switch(config)#int f0/23
Switch(config-if)#NO shut
```

You should see the following console messages:

```
%SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port FastEthernet0/23 with
BPDU Guard enabled. Disabling port.
%PM-4-ERR_DISABLE: bpduguard error detected on Fa0/23, putting Fa0/23 in
err-disable state
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/23, changed
state to down
%LINK-3-UPDOWN: Interface FastEthernet0/23, changed state to down
```

To verify that interface F0/23 is in err-disable mode:

On Cat-2:

```
Cat-2#Sh inter f0/23 status

Port      Name               Status       Vlan       Duplex  Speed Type
Fa0/23                       err-disabled 1            auto   auto 10/100BaseTX
```

To change the configuration back:

On Cat-2:

```
Cat-2(config)#NO spanning-tree portfast bpduguard default

Cat-2(config)#int f0/23
Cat-2(config-if)#Shut
Cat-2(config-if)#NO spanning-tree portfast
```

Task 12

You received a request from the IT department to monitor and analyze all the packets sent and received by the host connected to port F0/14 on Cat-1; you have connected the packet analyzer to port F0/15 on the same switch. Configure the switch to accommodate this request.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-1:

```
Cat-1(config)#monitor session 1 source interface F0/14 both
Cat-1(config)#monitor session 1 destination interface F0/15
```

Note the following:

- **There can only be up to 66 monitor sessions configured on a 3560 switch**
- **Their direction to monitor can be configured as Rx, Tx, or Both. Rx is for received traffic, Tx is for Transmitted traffic, and both is in both directions. Both is the default direction.**
- **To verify Enter the "Show monitor session 1" command.**

To verify the configuration:

On Cat-1:

```
Cat-1#Show monitor session 1

Session 1
---------
Type               : Local Session
Source Ports       :
    Both           : Fa0/14
Destination Ports  : Fa0/15
    Encapsulation  : Native
          Ingress  : Disabled
```

Task 13

You received another request from your IT department to keep track of all the MAC addresses that are learned by Cat-2 port F0/18. The switch must use the NMS located at 192.168.1.1 /24; this switch should send a community string of "Private" with the notification operation. You should use an IP address of 2.2.2.2 /8 to accomplish this task.

On Cat-2:

```
Cat-2(config)#Snmp-server host 192.168.1.1 traps Private
```

You should get the following console messages:

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

```
%IP_SNMP-3-SOCKET: can't open UDP socket
Unable to open socket on port 161
```

Since this switch is not configured with an IP address, it will fail to configure the Snmp server. Therefore, an IP address should be configured before entering the "snmp-server" command, the following command assigns an IP address of 2.2.2.2 /8 to a Loopback interface, understand that any Interface, and/or any IP address could have been used.

```
Cat-2(config)#Int Lo0
Cat-2(config-if)#IP addr 2.2.2.2 255.0.0.0
```

The following command identifies the NMS and sends a community string of Private with the notification operation.

```
Cat-2(config)#snmp-server host 192.168.1.1 traps Private
```

The following command configures the switch to send mac-address traps to the NMS:

```
Cat-2(config)#snmp-server enable traps mac-notification

Cat-2(config)#Inter f0/18
Cat-2(config-if)#snmp trap mac-notification added
```

The above command enables the SNMP trap on interface F0/18 and configures the switch to send MAC notification traps whenever a MAC-address is added. If the switch must be configured to report the MAC addresses that are removed and/or expired, then, "snmp trap mac-notification change removed" command must also be configured.

To verify the configuration:

On Cat-2:

```
Cat-2#Show mac-address-table notification inter f0/18

MAC Notification Feature is Disabled on the switch
Interface            MAC Added Trap MAC Removed Trap
---------            -------------- ----------------
FastEthernet0/18     Enabled        Disabled
```

Note the mac-notification is disabled, the following command will enable the mac-notification on the switch:

```
Cat-2(config)#mac address-table notification
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On Cat-2:

```
Cat-2#Show mac-address-table notification interface F0/18

MAC Notification Feature is Enabled on the switch

Interface            MAC Added Trap MAC Removed Trap
---------            -------------- ---------------

FastEthernet0/18     Enabled        Disabled
```

Task 14

Configure Cat-2's port F0/14 to limit the amount of bandwidth utilization for broadcast traffic to 50%.

On Cat-2:

```
Cat-2(config)#Interface F0/14
Cat-2(config-if)#Storm-control broadcast level 50.00
```

Storm-control can be used for Broadcast, Unicast and Multicast traffic, this command specifies traffic suppression level for a given type of traffic for a particular interface. The level can be a value from 0 to 100 and an optional fraction of a level can also be configured from 0 – 99. A threshold value of 100 percent means that no limit is placed for the specified type of traffic; a value of 0.0 means that the particular type of traffic is blocked all together.

To verify the configuration:

On Cat-2:

```
Cat-2#Show storm-control F0/14 broadcast

Interface  Filter State   Upper         Lower         Current
---------  ------------   -----------   -----------   ----------
Fa0/14     Link Down       50.00%        50.00%        0.00%
```

If you get "Link Down" under the "Filter State" column, the port might be down or a device is NOT

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

connected to this interface on the local switch.

Task 15

Mac addresses learnt dynamically by these two switches should not stay in the MAC address table if they are inactive for longer than 10 minutes.

By default, the MAC addresses that are inactive will expire within 300 seconds or 5 minutes; this task is asking for a 10 minute threshold, 10 minutes equates to 600 seconds; the following command sets the idle timer to 10 minutes:

On Both Switches:

The following command can be entered using the "Mac-address-table againg-time" or "Mac address-table aging-time 600".

```
Cat-x(config)#Mac address-table aging-time 600
```

To verify the configuration:

On Both Switches:

```
Cat-x#Sh mac address-table aging-time

Vlan    Aging Time
----    ----------
   1     600
  12     600
  34     600
  56     600
```

Task 16

For management purposes, assign an IP address of 10.1.1.11 /24 to Cat-1, with a default gateway of 10.1.1.100 /24. DO NOT enable IP routing to accomplish this task.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On Cat-1:

```
Cat-1(config)#Inter Vlan 1
Cat-1(config-if)#IP address 10.1.1.11 255.255.255.0
Cat-1(config-if)#NO shut

Cat-1(config)#IP default-gateway 10.1.1.100
```

To verify the configuration:

On Cat-1:

```
Cat-1#Sh ip interface vlan 1

Vlan1 is up, line protocol is up
  Internet address is 10.1.1.11/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
```

(The rest of the output is omitted)

```
Cat-1#Sh ip route

Default gateway is 10.1.1.100

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

Task 17

Configure R1 and R3 with the following IP addresses:

- R1 - F0/0 = 10.1.12.1 /24
- R3 - F0/0 = 10.1.34.3 /24

Configure Cat-1 to route between VLAN 12 and 34, use ping to verify the ommunication. The gateway for VLAN12 should be configured to be 10.1.12.100, and the gateway for VLAN 34 should be configured to be 10.1.34.100.

On R1:

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

```
R1(config)#Interface F0/0
R1(config-if)#IP address 10.1.12.1 255.255.255.0
R1(config-if)#No shut

R1(config)#IP route 0.0.0.0 0.0.0.0 10.1.12.100
```

On R3:

```
R3(config)#Interface F0/0
R3(config-if)#IP address 10.1.34.3 255.255.255.0
R3(config-if)#No shut

R3(config)#IP route 0.0.0.0 0.0.0.0 10.1.34.100
```

On Cat-1:

By default, IP routing is disabled on Catalyste 3560 switches, and it should be enabled using the following command:

```
Cat-1(config)#IP routing

Cat-1(config)#Interface Vlan 12
Cat-1(config-if)#Ip address 10.1.12.100 255.255.255.0

Cat-1(config)#Interface Vlan 34
Cat-1(config-if)#Ip address 10.1.34.100 255.255.255.0
```

Only one SVI can be associated with a VLAN. This is necessary when configuring InterVlan routing. When creating an SVI for a VLAN, the designated number must match the VLAN number. Remember that the local switch MUST be aware of the VLAN and have an interface in that given VLAN, or else the interface VLAN will NOT transition into UP/UP mode.

To verify the configuratiou:

On Cat-1:

```
Cat-1#Show IP interface brie | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  10.1.1.11       YES manual up                    up
Vlan12                 10.1.12.100     YES manual up                    up
Vlan34                 10.1.34.100     YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
Loopback1              11.1.1.1        YES manual up                    up
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Ping 10.1.34.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R3:

```
R3#Ping 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 18

Remove the SVI interfaces configured in the previous task. Configure InterVlan routing between VLANs 12 and 34. DO NOT use SVIs to accomplish this task. F0/1 interface of any router can be used to accomplish this task. Use the IP addressing from the previous task. Ensure to use an industry standard protocol/s to accomplish this task.

In the following configuration R5 is used to accomplish this task. Since R5's F0/0 is part of VLAN 56, R5's F0/1 is used to accomplish this task. This configuration is called "Router On a Stick".

On Cat-1:

```
Cat-1(config)#No Interface Vlan 12
Cat-1(config)#No Interface Vlan 34
```

R5's F0/1 interface is connected to Cat-2's F0/5 port. Since R5's F0/1 is used for routing purpose, the F0/5 interface of Cat-2 and the F0/1 interface of R5 must be configured as a trunk links.

On Cat-2:

```
Cat-2(config)#Interface F0/5
Cat-2(config-if)#Switchport trunk encap Dot1q
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Cat-2(config-if)#Switchport mode trunk
Cat-2(config-if)#No shut
```

On R5:

```
R5(config)#Interface F0/1
R5(config-if)#No Shut

R5(config)#Int f0/1.12
R5(config-if)#Encap dot1q 12
R5(config-if)#Ip address 10.1.12.100 255.255.255.0

R5(config)#Int f0/1.34
R5(config-if)#Encap dot1q 34
R5(config-if)#Ip address 10.1.34.100 255.255.255.0
```

To verify the configuration:

On R1 & R3:

```
R1#Clear arp
```

On R1:

```
R1#Ping 10.1.34.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.34.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R3:

```
R3#Ping 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, 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

Task 19

Optimize Cat-1 using the following policies:
Cat-1 should be configured such that its memory resources are optimized for routing.

Switch database management (SDM) are templates that can be configured to allocate memory resources in the switch for a specific feature depending on what the switch is used for in a given network.
A switch can be configured to use one of the following templates:

- **Access** – Used for QOS classification and Security.
- **Routing** – Used for routing
- **Vlan** – Disables routing and sets the switch to be a layer 2 switch.
- **Extended-match** – reformats routing memory space to allow 144-bit layer 3 TCAM support needed for WCCP and/or multiple VRF instances.

On Cat-1:

```
Cat-1(config)#Sdm prefer routing
```

You must reboot for these settings to take effect, DO NOT forget to save the configuration before reloading.

```
Cat-1#WR

Cat-1#Reload
```

To Verify the configuration after the reload:

On Cat-1:

```
Cat-1#Show sdm prefer

The current template is "desktop routing" template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1024 VLANs.
  number of unicast mac addresses:                  3K
  number of IPv4 IGMP groups + multicast routes:    1K
  number of IPv4 unicast routes:                    11K
    number of directly-connected IPv4 hosts:        3K
    number of indirect IPv4 routes:                 8K
  number of IPv4 policy based routing aces:         512
  number of IPv4/MAC qos aces:                      512
  number of IPv4/MAC security aces:                 1K
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To see the difference between Cat-1 (Using the default template) and Cat-2 (Using the routing):

On Cat-2:

```
Cat-2#Sh sdm prefer

The current template is "desktop default" template.
 The selected template optimizes the resources in
 the switch to support this level of features for
 8 routed interfaces and 1024 VLANs.

  number of unicast mac addresses:                  6K
  number of IPv4 IGMP groups + multicast routes:    1K
  number of IPv4 unicast routes:                    8K
    number of directly-connected IPv4 hosts:        6K

    number of indirect IPv4 routes:                 2K
  number of IPv4 policy based routing aces:         0
  number of IPv4/MAC qos aces:                      512
  number of IPv4/MAC security aces:                 1K
```

Note, the difference in memory allocation is revealed if the buffer allocation of Cat-2 is compared to Cat-1.

Task 20

Configure Cat-1's port F0/17 and F0/18 such that when client PCs connect to these ports, they automatically become a member of a given VLAN. Cat-1 should be configured to use 10.1.1.1 as the primary and 10.1.1.2 as the secondary VMPS server. Ensure that the local switch reconfirms the VLAN membership every half an hour and if the VMPS can not be contacted, the local switch will retry 5 times before considering the VMPS unavailable.

VMPS:

- The 3550 switch can't be setup as a VMPS server, but it can be configured as a VMPS client.
- The client communicates with the VMPS through Vlan Query Protocol (VQP).
- When a VMPS receives a VQP from the client, it searches its database for a MAC to VLAN mapping, and if the mapping is found, it conveys the VLAN information to the client and then the client assigns that given VLAN to the port that the client is connected to.

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

- The VMPS can operate in Secure mode, which means that if a MAC to VLAN mapping can not be found in its database, the VMPS will send a port-shutdown-message to the client and the client will shut down that given port, however, if the VMPS is not configured in a secure mode, it will send access-deny message, and the client will constantly monitor the port and will reject all traffic from that given port.

- The VMPS client periodically reconfirms the VLAN membership information received from the VMPS server. By default this is performed every 60 minutes, this interval can be changed using "VMPS reconfirm" global config command.

- If the VMPS client can't contact the VMPS server, it will retry to establish that communication three times and this value can be changed using vmps retry" command in the global config mode.

- The database is in form of an ASCII file saved on a TFTP server, which the VMPS server accesses.

On Cat-1:

Before configuring this task we should check some of the default parameters:

```
Cat-1#Show vmps

VQP Client Status:
--------------------

VMPS VQP Version:    1
Reconfirm Interval: 60 min
Server Retry Count: 3
VMPS domain server:


Reconfirmation status
---------------------
VMPS Action:          No Dynamic Port
```

VMPS VQP version is version 1, and the reconfirmation is at its default value of 60 minutes, and the retry value is set to 3. There are no VMPS servers.

```
Cat-1(config)#int range f0/17 - 18
Cat-1(config-if-range)#switchport mode access
Cat-1(config-if-range)#switchport access vlan dynamic
Cat-1(config-if-range)#NO shut
```

The above command sets ports F0/15 and F0/16 to VLAN dynamic, which means that they will acquire their VLAN information dynamically. The "NO Shut" command is required because these ports were shut

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

down earlier

```
Cat-1(config)#vmps reconfirm 30
Cat-1(config)#vmps retry 5
```

The above two commands configure the reconfirmation interval to 30 minutes and the retry counter to 5.

```
Cat-1(config)#vmps server 10.1.1.1 primary
Cat-1(config)#vmps server 10.1.1.2
```

These commands configure the primary and the secondary VMPS servers.

To verify the configuration:

On Cat-1:

```
Cat-1#Show vmps

VQP Client Status:
--------------------
VMPS VQP Version:   1
Reconfirm Interval: 30 min
Server Retry Count: 5
VMPS domain server: 10.1.1.2
                    10.1.1.1 (primary, current)

Reconfirmation status
---------------------
VMPS Action:          No Dynamic Port
```

Task 21

Configure trunking between Cat-1 and Cat-2 such that VLAN 12 does not get tagged when the traffic for this VLAN traverses the trunk.

Note the trunking encapsulation on the trunk links should have been DOT1Q; in the CCIE lab, when configuring a given section, the entire section should be read before configuring the individual tasks within that section.

When a trunk is configured with Dot1q, it can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic in the native VLAN ONLY. If a given VLAN should NOT be tagged as it

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

traverses the trunk link then, that VLAN should be set as the native VLAN.
When the native VLAN is changed, ensure that the change is configured on both switches or the trunk link will go down.

On Both Switches:

```
Cat-x(config)#Interface range F0/19-20
Cat-x(config-if-range)#Switchport trunk encap dot1q
```

To Verify the configuration:

On Cat-1:

```
Cat-1#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      1,12,34,56

Fa0/20      1.12,34,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      1
Fa0/20      1
```

To configure the native VLAN:

On Both Switches:

```
(config)#Interface range F0/19-20
(config-if-range)#Switchport trunk native VLAN 12
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show interface trunk
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Port        Mode         Encapsulation  Status        Native vlan
Fa0/19      on           802.1q         trunking      12
Fa0/20      on           802.1q         trunking      12

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

Port        Vlans allowed and active in management domain
Fa0/19      1,12,34,56
Fa0/20      1,12,34,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/19      34,56
Fa0/20      none
```

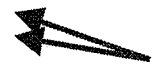

On Cat-2:

```
Cat-2#Show inter trunk

Port        Mode         Encapsulation  Status        Native vlan
Fa0/5       on           802.1q         trunking      1
Fa0/19      on           802.1q         trunking      12
Fa0/20      on           802.1q         trunking      12

Port        Vlans allowed on trunk
Fa0/5       1-4094
Fa0/19      1-4094
Fa0/20      1-4094

Port        Vlans allowed and active in management domain
Fa0/5       1,12,34,56
Fa0/19      1,12,34,56
Fa0/20      1,12,34,56

Port        Vlans in spanning tree forwarding state and not pruned
Fa0/5       1,12,34,56
Fa0/19      1,12,34,56
Fa0/20      none
```


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

Task 22

The IT department decided to stop monitoring port F0/14 from Task 14, you have received a new request to monitor port F0/14 on Cat-1 but the protocol analyzer is connected to port F0/18 on Cat-2. Configure the switches to accommodate this request.

On Cat-1:

```
Cat-1(config)#No monitor session 1

Cat-1(config)#Vlan 90
Cat-1(config-vlan)#Remote-span
Cat-1(config-vlan)#Exit
```

The creation of this VLAN can only be done in the global configuration mode, because this is the only mode that allows us to set the VLAN as remote-span. Ensure that this VLAN is propagated to Cat-2.

To verify the configuration:

On Cat-1:

```
Cat-1#Sh vlan brie | Inc 90

90   VLAN0090                         active
```

On Cat-2:

```
Cat-2#Sh vlan brie | Inc 90

90   VLAN0090                         active
```

On Cat-1:

```
Cat-1#Show vlan remote-span

Remote SPAN VLANs
------------------------------
90
```

On Cat-2:

```
Cat-2#Show vlan remote-span
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Remote SPAN VLANs
------------------------------
90
```

Note VLAN 90 should be displayed as remote-span on both switches.

On Cat-1:

```
Cat-1(config)#Monitor session 1 source interface F0/14
Cat-1(config)#Monitor session 1 destination remote vlan 90
```

To verify the configuration:

On Cat-1:

```
Cat-1#Show monitor session 1


Session 1
---------
Type               : Remote Source Session
Source Ports       :
    Both           : Fa0/14
Dest RSPAN VLAN    : 90
```

On Cat-2:

```
Cat-2(config)#Monitor session 1 source remote vlan 90
Cat-2(config)#Monitor session 1 destination interface F0/18
```

Port F0/18 is where the protocol analyzer is connected.

To verify the configuration:

On Cat-2:

```
Cat-2#Sh monitor session 1

Session 1
---------
Type               : Remote Destination Session
Source RSPAN VLAN  : 90
```


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

```
Destination Ports : Fa0/18
    Encapsulation : Native
          Ingress : Disabled
```

RSPAN extends SPAN by enabling remote monitoring of multiple switches across your network. The Traffic for RSPAN traverses over a user defined RSPAN VLAN (remote vlan), in this case VLAN 90. The SPAN traffic from port F0/14 is reflected to VLAN 90 (The RSPAN VLAN) and then forwarded over the trunk to port F0/18 an RSPAN destination.

Task 23

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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

LAB 2- Spanning-tree Protocol 802.1D

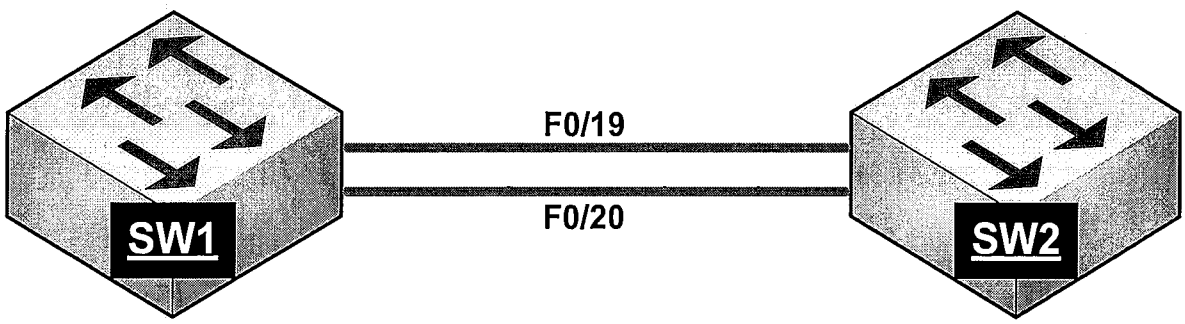

Task 1

Shutdown all ports on the four switches.

On All Switches:

```
Switch(config)#Int range f0/1-24
Switch(config-if-range)#Shut
```

To verify the configuration:

On All Switches:

```
Switch#Show interface status | Exc disabled|notconnect

Port      Name               Status       Vlan       Duplex  Speed Type
```

Task 2

Configure Dot1q trunking on the F0/19 and F0/20 interfaces of SW1 and SW2.

https://t.me/learningnets

On SW1 and SW2:

```
SW2(config)#Int range f0/19-20
SW2(config-if-range)#Switchport trunk encapsulation dot1q
SW2(config-if-range)#Switchport mode trunk
SW2(config-if-range)#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
Fa0/20      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      none
Fa0/20      none
```

Task 3

Which switch is the root bridge and why?

Before we start with the show commands, let's review the STP protocol:

When the switches come up, they will both think of themselves as the root bridge, and they will send BPDUs out every port advertising them as the root bridge. What does a BPDU look like?

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

2 Bytes	1 Byte	1 Byte	1 Byte	8 Bytes	4 Bytes	8 Bytes	2 Bytes	2 Bytes	2 Bytes	2 Bytes	2 Bytes
Protocol-ID	Version	Msg Type	Flags	Root ID	Root-Path-Cost	Bridge-ID	Port-ID	Msg Age	Max Age	Hello Time	Forward-delay

Let's explain the fields:

Protocol-ID	Indicates the type of the protocol, it's set to zero
Version	Identifies the version of the protocol, it's set to zero
Message Type	Indicates the type of message, it's set to zero
Flags	This field includes one of the following: • TC-bit, which signals a topology change • TCA-bit, which is set to ACK the receipt of a configuration Message with the TC-bit set
Root ID	The BID of the root bridge
Root Path Cost	Cumulative cost of the sending bridge to the root bridge
Bridge ID	Indicates the Priority and the BID of the sending bridge
Port ID	Indicates the port number through which the BPDU was sent
Message Age	The elapsed time since the root bridge sent the configuration message
Max-Age	Indicates when the current configuration message should be deleted
Hello Time	The time between the root bridge configuration messages
Forward-delay	indicates the legth of time that the bridge should wait before transitioning to a new state after a topology change

So initially, every switch will set the Root-ID and the Bridge-ID to the local BID's value.

Let's see the BID of each switch:

On SW1

```
SW1#Show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    32769
              Address     0012.7f40.9380
              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     0012.7f40.9380
              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        128.21   P2p
Fa0/20           Desg FWD 19        128.22   P2p
```

We can see that the BID which is a concatenation of Priority value and the MAC address in the Bridge-ID and the Root ID section of the above show command are identical, which means that this bridge MUST be the root bridge, and the area that is highlighted in green clearly states that the "This bridge is the root".

The receiving bridge compares the Root-id to its own Root-id, and the lower value wins and if the received Root-id is better (Lower) than the local Root-id, then, the local Root-id is replaced with the Root-id in the received BPDUs.

Since the MAC address is different on every switch, the priority is looked at first, and as a tie breaker the switch with a lowest MAC address becomes the Root bridge.

Let's look at SW2:

On SW2:

```
SW2#Show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    32769
              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    32769  (priority 32768 sys-id-ext 1)
              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 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
```

Another way of knowing which switch is the Root bridge is to use the following command:

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         32769 0012.7f40.9380        19    2   20  15   Fa0/19
```

NOTE: The last field (Root Port) indicates that the root bridge is found through F0/19 interface. Let's use CDP to find out the device that is connected to F0/19 interface:

```
SW2#Show cdp neighbor F0/19 | B Device ID

Device ID        Local Intrfce       Holdtme    Capability     Platform    Port ID
SW1              Fas 0/19              173          S I      WS-C3560-2Fas 0/19
```

Let's check SW1:

```
SW1#Show spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ----------
VLAN0001         32769 0012.7f40.9380         0    2   20  15
```

NOTE: The "Root Port" column is empty, which indicates that this switch is the Root bridge.

Task 4

Which port is the Root-Port?

Every Non-Root Bridge must select a Root Port. The Root Port is the closest port to the Root Bridge. The Root port calculation is based on the Root-Path-Cost, which is the cumulative cost of all links to the Root Bridge.

In this topology, SW2 is the None Root Bridge, so let's find out the Root Port:

On SW2:

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

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
```


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

```
Fa0/19           Root FWD 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
```

We can clearly see that the F0/19 of SW2 is the root port, but what if there is a tie?

Let's go through the golden rules that STP uses to break ties:

- A lower Root BID
- A lower Path cost to the Root Bridge
- A lower Sending BID
- A lower Sending Port-ID, which is the combination of "Priority.Port-id"

Since the Root Bridge is already known, let's go with the second rule and check the Path cost to the Root Bridge:

On SW2:

```
SW2#Sh spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ----------
VLAN0001         32769 0012.7f40.9380        19    2   20  15   Fa0/19
```

Let's shutdown the F0/19 interface and check the cost through F0/20 interface:

```
SW2(config)#Int F0/19
SW2(config-if)#Shut

SW2#Show spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------- --------- ----- --- ---  ----------
VLAN0001         32769 0012.7f40.9380        19    2   20  15   Fa0/20
```

Let's enable the F0/19 interface of SW2:

On SW2:

```
SW2(config)#Int F0/19
SW2(config-if)#No shut
```

In this case both F0/19 and F0/20 have the same cost.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

So since the cost to the Root Bridge is the same through both paths, let's check the next rule, which is the "Lower Sending BID", in this case it will be the same, since both interfaces are connected to the same Switch (SW1); therefore, let's look at the last rule, "The lowest sending Port-ID", to find out the

lowest sending port-id, we can use the "Show spanning-tree" command:

On SW2:

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

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Root FWD 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
```

We can see why the F0/19 interface is the Root port and the F0/20 interface is in "BLK" state, the "Prio.Nbr" column reveals the priority.Port-ID of the neighboring switch. You can see that the F0/19 interface and the F0/20 interface receive the same port-priority value from SW1, but the port-id is lower through the local F0/19 interface versus the F0/20 interface of SW2.

Task 5

Which port is the Designated-Port for the two segments?

There should be one designated port per segment, there are two segments connecting the two switches, since SW1 is the Root Bridge, and all the ports on the Root bridge will always be in designated state, ports F0/19 and F0/20 of SW1 is elected as the designated ports on the two segments; the designated ports are elected based on the lowest path cost.

let's verify:

On SW1:

```
SW1#Show spanning-tree root

                                        Root    Hello Max Fwd
Vlan                Root ID             Cost    Time  Age Dly  Root Port
---------------- ------------------ --------- ----- --- ---  ----------
VLAN0001         32769 0012.7f40.9380         0     2   20  15
```

NOTE: No matter which port is used on the root bridge (SW1), the cost is zero, and that is why all

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

interfaces on the Root bridge will always be in designated state because they will always be the closest interface to the root bridge.

Task 6

Which port is in the "BLK" state?

Once all the designated ports and the Root ports are determined, the rest of the port/s (Left over ports) will be in blocked state, let's verify:

On SW1:

```
SW1#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------

Number of blocked ports (segments) in the system : 0
```

Of course, there should NOT be any ports in blocking state on the root bridge. Let's verify the blocked port on SW2:

On SW2:

```
SW2#Show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------
VLAN0001             Fa0/20

Number of blocked ports (segments) in the system : 1
```

Let's verify that information:

On SW2:

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

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Root FWD 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 7

Configure SW2 such that its F0/20 interface transitions into "FWD" state and the F0/19 interface transitions into "BLK" state.

The "BLK" port is the port with the highest path cost, therefore, if the cost of the F0/20 interface is changed to be lower than the F0/19 interface, then the F0/20 interface will transition into "FWD" state and the F0/19 interafce will transition into "BLK" state. Let's test this:

On SW2:

```
SW2(config)#Int F0/20
SW2(config-if)#Spanning-tree cost 10
```

To verify the configuration:

On SW2:

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

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Root LIS 10        128.22   P2p

SW2#Show spannin | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Root LRN 19        128.22   P2p

SW2#Show spanning-tree | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Root FWD 10        128.22   P2p
```

We can see that the F0/20 goes through Listenening and learning state and transitions into "FWD" state, and the F0/19 transitions into "BLK" state.

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

Task 8

Remove the configuration commands from the previous task, and configure SW1 such that the F0/20 interface of SW2 transitions into "FWD" state and the F0/19 interface of SW2 transitions into "BLK" state.

On SW2:

```
SW2(config)#int f0/20
SW2(config-if)#No Spanning-tree cost 10
```

To verify the configuration:

On SW2:

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

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Root FWD 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p
```

To configure SW1

```
SW1(config)#Int F0/20
SW1(config-if)#Spanning-tree port-priority 0
```

To verify the configuration:

On SW1:

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

Interface        Role Sts Cost      Prio.Nbr Type

---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19          0.22   P2p
```

On SW2:

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

Interface        Role Sts Cost      Prio.Nbr Type
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Root FWD 19        128.22   P2p
```

As you can see, when it comes to port-pirority, it affects the neighboring switch.

Task 9

Configure SW2 to be the root bridge. You should use a macro to accomplish this task.

To accomplish this task using a MACRO, we can use, the "root Primary", let's test this MACRO:

On SW2:

```
SW2(config)#Spanning-tree  vlan 1 root primary
```

To verify the configuration:

On SW2:

```
SW2#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     001d.e5d6.0000
              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     001d.e5d6.0000
              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
```

NOTE: The default priority is 32768, and with every VLAN, the default value is incremented by the VLAN ID, in this case the ONLY VLAN in the Database is VLAN 1, therefore, 32768 + 1 = 32769.

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

Using the "Spanning-tree root primary" Macro, the total priority is reduced by 8192, so:

32769 – 8192 = 24577, and we know that the switch with the lowest priority will become the root bridge.

Task 10

Remove the command from the previous task, and configure SW2 to be the root bridge. You should NOT use a macro to accomplish this task.

On SW2:

```
SW2(config)#No spanning-tree vlan 1 root pri
```

To verify the configuration:

On SW1:

```
SW1#Show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    32769
              Address     0012.7f40.9380
              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     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          0.22   P2p
```

On SW2:

```
SW2(config)#Spanning-tree vlan 1 priority 0
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW2:

```
SW2#Show spanning-tree vlan 1

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    1
              Address     001d.e5d6.0000
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    1      (priority 0 sys-id-ext 1)
              Address     001d.e5d6.0000
              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
```

Task 11

Remove the command from the previous task, and configure two VLANs 100 and 200. SW1 should be configured such that on SW2 the traffic for VLAN 100 takes the F0/19 interface, whereas, the traffic for VLAN 200 takes the F0/20 interface. Configure this task with redundancy in mind.

On SW2:

```
SW2(config)#No Spanning-tree vlan 1 priority 0
```

On SW1:

```
SW1(config)#int f0/20
SW1(config-if)#No spanning-tree port-priority 0

SW1(config)#vtp domain tst
Changing VTP domain name from NULL to tst

SW1(config)#VLAN 100,200
```


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

```
SW1(config-vlan)#exit
```

To verify the configuration:

On SW2:

```
SW2#Show vlan brie | 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
100  VLAN0100                            active
200  VLAN0200                            active
```

We can see that the configured VLANs (100 and 200) are propagated to SW2 via VTP messages. Let's configure the load sharing part of this task:

```
SW1(config)#Int F0/19
SW1(config-if)# Spanning-tree vlan 100 port-priority 16

SW1(config-if)#int f0/20
SW1(config-if)#Spanning-tree vlan 200 port-priority 16
```

To verify the configuration:

On SW2:

The output of the following show commands reveal that on SW2 the traffic for VLAN 100 uses the F0/19 interface, whereas, the traffic for VLAN 200 uses the F0/20 interface.

```
SW2#Show spanning-tree vlan 100 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Root FWD 19        128.21   P2p
Fa0/20           Altn BLK 19        128.22   P2p

SW2#Show spanning-tree vlan 200 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
---------------- ---- --- --------- -------- -----
Fa0/19           Altn BLK 19        128.21   P2p
Fa0/20           Root FWD 19        128.22   P2p
```

Let's verify these values on SW1

On SW1:

```
SW1#Show spanning-tree vlan 100 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        16.21    P2p
Fa0/20           Desg FWD 19        128.22   P2p

SW1#Show spanning-tree vlan 200 | B Interface

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- -----
Fa0/19           Desg FWD 19        128.21   P2p
Fa0/20           Desg FWD 19        16.22    P2p
```

Task 12

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

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

CCIE Foundation

5.0

www.MicronicsTraining.com

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

DMVPN

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1 - DMVPN – Phase #1 with Static Mapping

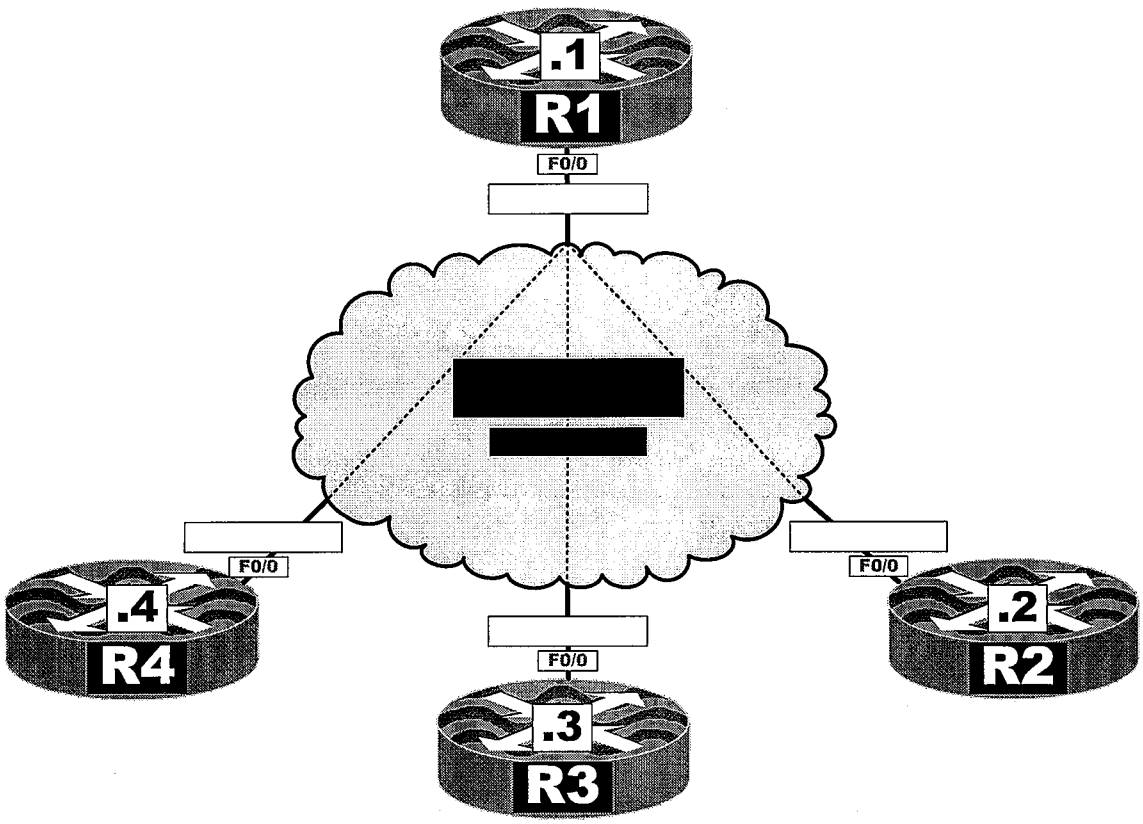

Task 1

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
```

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, 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 = 1/2/4 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.

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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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.

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.

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

- 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
   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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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
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
```


https://t.me/learningnets

```
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

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.4, 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
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.

https://t.me/learningnets

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

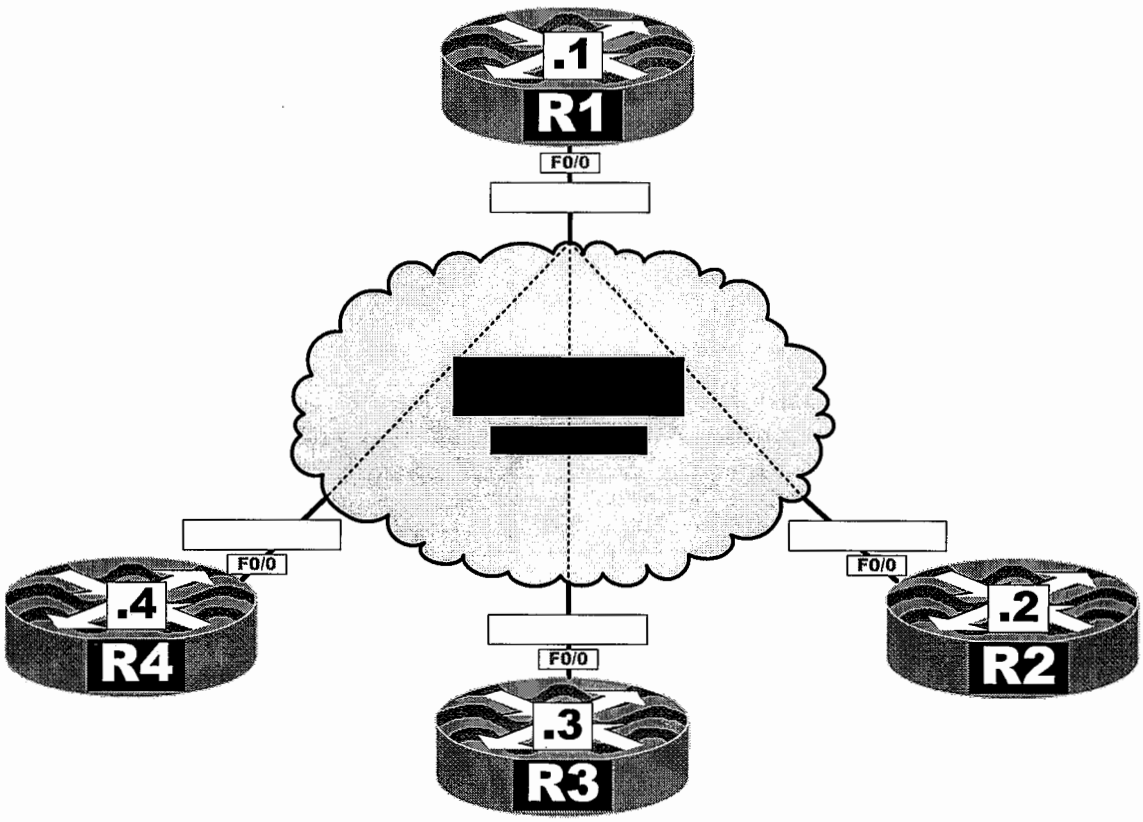

Task 1

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 .

On SW1:

© https://t.me/learningnets

```
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
```

On R3:

```
R3(config)#Int f0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
!!!!!
```


© [illegible] Rights Reserved
https://t.me/learningnets

```
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 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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

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
```


©[illegible] All rights reserved
https://t.me/learningnets

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
   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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   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

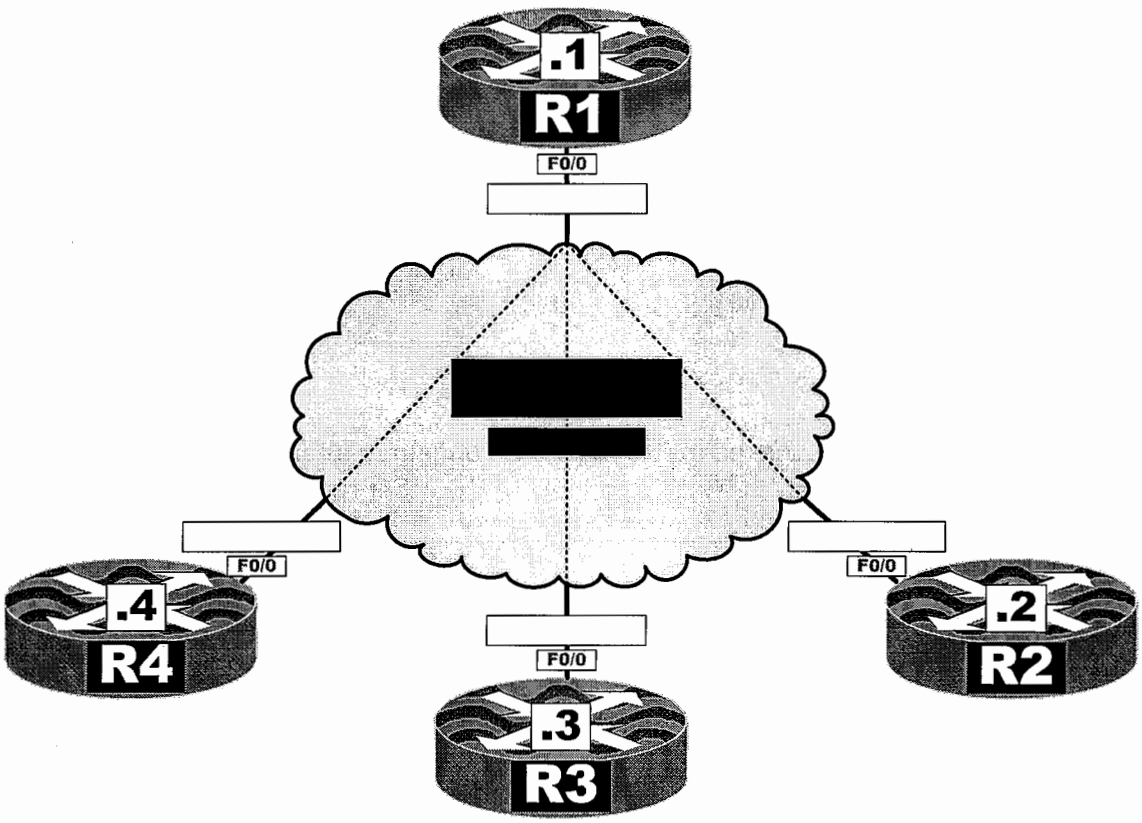

Task 1

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 .

On SW1:

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

```
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
```

On R3:

```
R3(config)#Int f0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, timeout is 2 seconds:
!!!!!
```


© 2010 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

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
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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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:

```
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
```


© 2011 Narbik Kocharians All rights reserved
https://t.me/learningnets

```
   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
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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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
```

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:
!!!!!
```


©https://t.me/learningnets

```
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

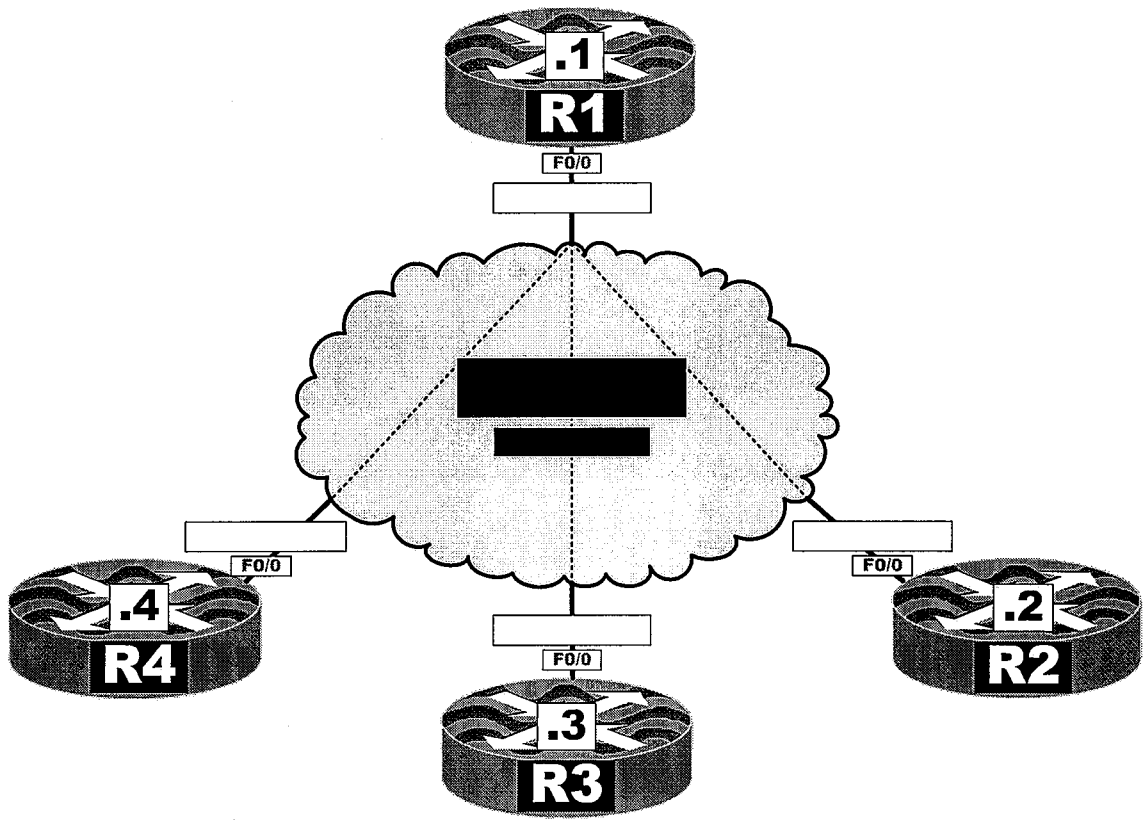

Task 1

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.

On SW1:

©https://t.me/learningnets

```
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
```

On R3:

```
R3(config)#Int f0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.1.1, 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 = 1/2/4 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
```

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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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:

```
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
```


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

```
 (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
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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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/learninghets

Lab 5 – Running Routing Protocols on DMVPN Phase #1

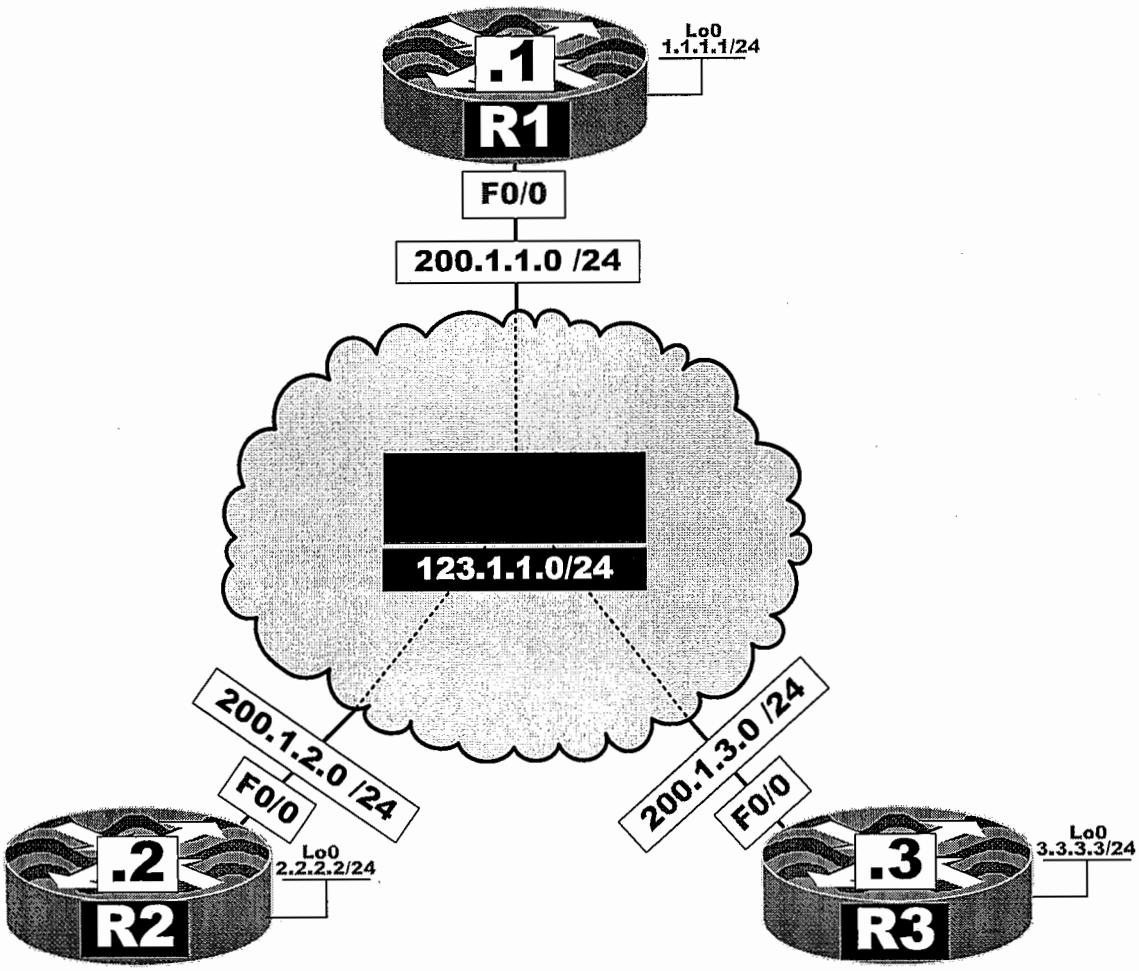

Task 1

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 in this

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

R2(config)#ip route 0.0.0.0 0.0.0.0 200.1.2.10
```


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

On R3:

```
R3(config)#Int f0/0
R3(config-if)#ip addr 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:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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
```


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

Question:
Why is R2 able to Ping R3 when R2 or R3 have a single mapping pointing to the hub (R1)?

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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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
```

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
```


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

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
```

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  2 123.1.1.3 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)
  1 123.1.1.1 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:

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

On R1:

```
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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
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
```


© 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
```

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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 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:

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

```
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
```

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#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, 00:12:57, 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
```


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

```
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
```

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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 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
   Network          Next Hop            Metric LocPrf Weight Path
```


https://t.me/learningnets

```
*> 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

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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
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

https://t.me/learningnets

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:

```
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:
```


© 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

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.

© [illegible] reserved
https://t.me/learningnets

Lab 6 – Running Routing Protocols on DMVPN Phase #2 and Phase #3

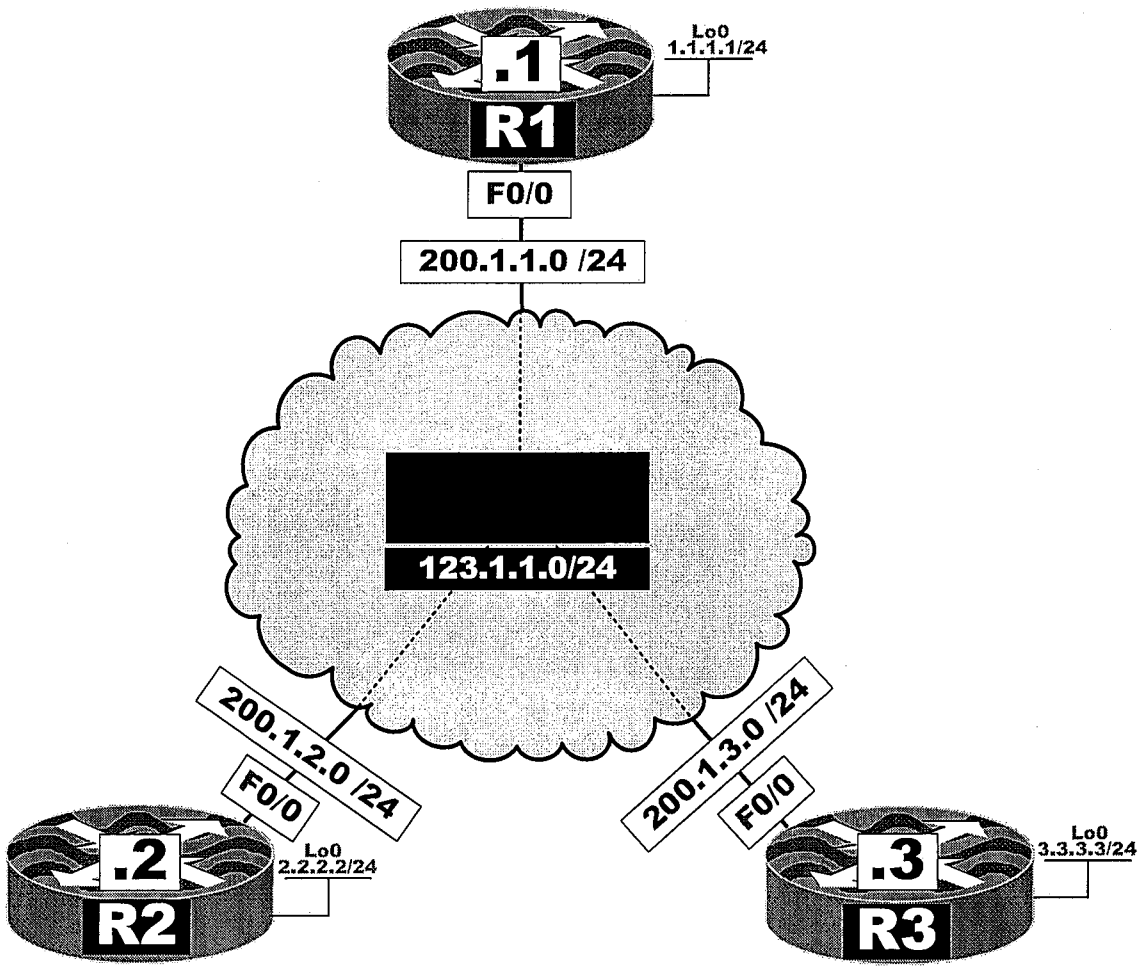

Task 1

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 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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

R2(config)#ip route 0.0.0.0 0.0.0.0 200.1.2.10
```

On R3:

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

```
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
```


https://t.me/learningnets

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 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:

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
```

On R2:

https://t.me/learningnets

```
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
```

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?

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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)
  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.
```


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

```
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

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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
```

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
```


© 2013 Narbik Kocharians All rights reserved
https://t.me/learningnets

```
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:

```
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
```


© 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 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.
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
```


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

```
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 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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

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
```


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

```
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

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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:

```
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
```


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

```
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
```

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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

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.
```


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

```
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. 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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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:

On R1:

```
R1(config)#No router bgp 100

R1(config)#Router eigrp 100
R1(config-router)#No au
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:

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

```
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)#No au
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)#No au
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.

To test the configuration:

On R2:

```
R2#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)
  1 123.1.1.1 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 question 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
   NBMA address: 200.1.3.3
```


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

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

CCIE Foundation

5.0

www.MicronicsTraining.com

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

RIPv2

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

Lab 1
Configuring RIPv2

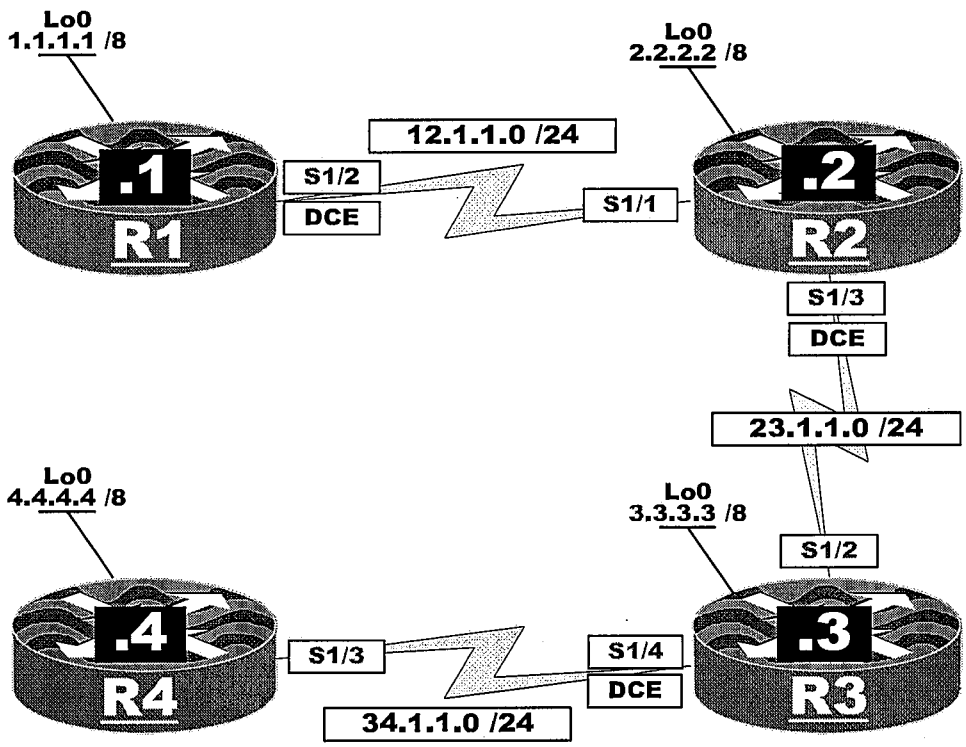

Task 1

Configure the routers based on the above topology. DO NOT configure any routing protocol.

To build the logical topology should be very easy, let's start with R1 and R2:

On R1:

```
R1(config)#int s1/2
```

By looking at the above diagram we can see that R1 has the DCE end of the cable, therefore, the "clock rate" must be configured on the S1/2 interface of R1. Whenever we see a serial connection, we should ask ourselves which device has the DCE end of the serial cable and configure clocking on the interface of that device.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#clock rate 64000
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```

On R2:

```
R2(config)#int s1/1
R2(config-if)#ip addr 12.1.1.2 255.255.255.0
R2(config-if)#No shu

R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0

R2(config-if)#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
```

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)#Int Lo0
R3(config-if)#Ip addr 3.3.3.3 255.0.0.0

R3(config-if)#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
```

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)#Int lo0
R4(config-if)#Ip addr 4.4.4.4 255.0.0.0
```

To verify the configuration:

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

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R1#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              12.1.1.1        YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
```

On R2:

```
R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R2#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/1              12.1.1.2        YES manual up                    up
Serial1/3              23.1.1.2        YES manual up                    up
Loopback0              2.2.2.2         YES manual up                    up
```

On R3:

```
R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R3#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              23.1.1.3        YES manual up                    up
Serial1/4              34.1.1.3        YES manual up                    up
Loopback0              3.3.3.3         YES manual up                    up
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R4:

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

Interface                  IP-Address      OK? Method Status                Protocol
Serial1/3                  34.1.1.4        YES manual up                    up
Loopback0                  4.4.4.4         YES manual up                    up
```

Task 2

Configure RIPv2 on all routers and advertise their directly connected networks. When configuring RIP, ensure that auto summarization is disabled.

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 1.0.0.0
```

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)#network 23.0.0.0
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 23.0.0.0
R3(config-router)#network 34.0.0.0
R3(config-router)#network 3.0.0.0
```

On R4:

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

```
R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#network 34.0.0.0
R4(config-router)#network 4.0.0.0
```

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 [120/1] via 12.1.1.2, 00:00:03, Serial1/2
R     3.0.0.0/8 [120/2] via 12.1.1.2, 00:00:03, Serial1/2
R     4.0.0.0/8 [120/3] via 12.1.1.2, 00:00:03, Serial1/2
R     23.0.0.0/8 [120/1] via 12.1.1.2, 00:00:03, Serial1/2
R     34.0.0.0/8 [120/2] via 12.1.1.2, 00:00:03, Serial1/2
```

On R4:

```
R4#Show ip route rip | B Gate
Gateway of last resort is not set

R     1.0.0.0/8 [120/3] via 34.1.1.3, 00:00:22, Serial1/3
R     2.0.0.0/8 [120/2] via 34.1.1.3, 00:00:22, Serial1/3
R     3.0.0.0/8 [120/1] via 34.1.1.3, 00:00:22, Serial1/3
R     12.0.0.0/8 [120/2] via 34.1.1.3, 00:00:22, Serial1/3
R     23.0.0.0/8 [120/1] via 34.1.1.3, 00:00:22, Serial1/3

R4#Show ip protocol | B Gateway

    Gateway         Distance      Last Update
    34.1.1.3             120      00:00:14
  Distance: (default is 120)
```

NOTE: Even though RIP does not have a neighbor table, you can always see the gateways (routers) that the local router is exchanging routes with, in this case 34.1.1.3 (R3).

On R3:

```
R3#Show ip protocol | B Gateway

    Gateway         Distance      Last Update
    23.1.1.2             120      00:00:28
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
    34.1.1.4                120       00:00:26
  Distance: (default is 120)
```

The output of the above show command identifies R2 and R4 as neighbors.

Task 3

Configure these routers such that update interval, invalidation timer, Holddown timer and flush timers are double the default value.

To find out the default parameters:

On R1:

```
R1#Show ip protocols | Inc seconds

  Sending updates every 30 seconds, next due in 10 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
```

On R1:

```
R1(config)#Router rip
R1(config-router)#Timers Basic ?
<0-4294967295>  Interval between updates

R1(config-router)#Timers basic 60 ?
  <1-4294967295>  Invalid

R1(config-router)#Timers basic 60 360 ?
  <0-4294967295>  Holddown

R1(config-router)#Timers basic 60 360 360 ?
  <1-4294967295>  Flush

R1(config-router)#Timers basic 60 360 360 480
```

On All Routers:

```
Rx(config)#Router rip
Rx(config-router)#Timers basic 60 360 360 480
```


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

To verify the configuration:

On R1:

```
R1#Show ip protocol | Inc seconds

  Sending updates every 60 seconds, next due in 41 seconds
  Invalid after 360 seconds, hold down 360, flushed after 480
```

Task 4

These routers should suppress a flash update if the regular update is due in 10 seconds or less.

On All Router:

```
Rx(config)#router rip
Rx(config-router)#flash-update-threshold 10
```

The above configuration configures all four routers to suppress a regularly scheduled flash update, if the regular update is due in 10 seconds or less.

RIP enabled routers send their entire routing table every 30 seconds, this is done because RIP does not have a keepalive mechanism. RIP routing protocol also performs a flash update if there is a topology change, in which case only the changes are sent to all directly connected neighbors.
Let's assume that we just sent an update and 20 seconds later there is a topology change, RIP will trigger a flash update to convey the changes and 10 seconds later it will send a regular update that will include the changes, the flash update in this case is redundant and it can be suppressed using the "Flash-update-threshold" command, this command suppresses the triggered flash update if it's 10 seconds before a regular update.

To verify the configuration:

On All Routers:

```
R4#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 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 2 – RIPv2 Authentication

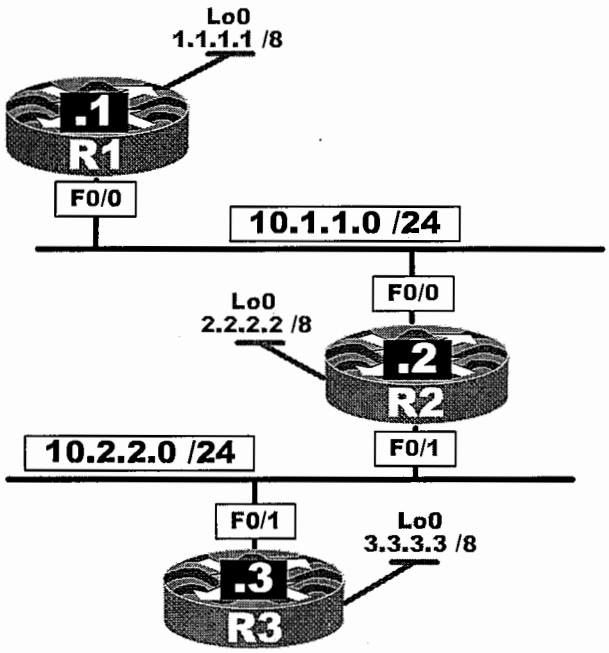

Task 1

Configure routers and the switches based on the above diagram. DO NOT configure RIPv2.

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)#No shut
```

On SW2:

```
SW2(config)#Int Range F0/2-3
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 F0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#Ip addr 10.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int Lo0
R1(config-if)#Ip addr 1.1.1.1 255.0.0.0
```

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

R2(config)#Int F0/1
R2(config-if)#IP addr 10.2.2.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int Lo0
R2(config-if)#Ip addr 2.2.2.2 255.255.255.0
```

On R3:

```
R3(config)#Int F0/1
R3(config-if)#Ip addr 10.2.2.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Int Lo0
R3(config-if)#Ip addr 3.3.3.3 255.0.0.0
```

To verify 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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R1#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.1        YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
```


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

On R2:

```
R2#Ping 10.2.2.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.2.2.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R2#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        10.1.1.2        YES manual up                    up
FastEthernet0/1        10.2.2.2        YES manual up                    up
Loopback0              2.2.2.2         YES manual up                    up
```

Task 2

Configure RIPv2 on R1 and R2 and advertise their directly connected networks in this routing protocol. You should disable auto summarization when configuring RIP.

On R1:

NOTE: The "Network 0.0.0.0" statement advertises all directly connected interfaces in this routing domain.

```
R1(config-if)#router rip
R1(config-router)#ver 2
R1(config-router)#No au
R1(config-router)#netw 0.0.0.0
```

On R2:

```
R2(config-if)#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
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

R     2.0.0.0/8 [120/1] via 10.1.1.2, 00:00:03, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R        10.2.2.0/24 [120/1] via 10.1.1.2, 00:00:03, FastEthernet0/0
```

On R2:

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

R     1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:05, FastEthernet0/0
```

Task 3

Configure clear text RIPv2 authentication between R1 and R2. Use "**cisco**" as the password for this authentication.

In RIPv2 there are two types of authentication, clear text and MD5. When configuring clear text authentication, the key numbers don't have to match, meaning that R1 can use one key number and R2 can use a totally different one, but the passwords MUST MATCH.

RIPv2 uses key-chain based authentication, meaning that a keychain is configured, a key is identified, and under that key a password is configured.

On Both Routers:

```
Rx(config)#Key chain TST
Rx(config-keychain)#key 1
Rx(config-keychain-key)#key-string cisco

Rx(config-keychain-key)#int f0/0
Rx(config-if)#ip rip authentication key-chain TST
```

To verify the configuration:

On R1:

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

```
R1#Sh ip proto | Inc TST|Key

    Interface              Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0        2     2                    TST

R1#Show ip route rip | B Gate
Gateway of last resort is not set

R     2.0.0.0/8 [120/1] via 10.1.1.2, 00:00:11, FastEthernet0/0
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R        10.2.2.0/24 [120/1] via 10.1.1.2, 00:00:11, FastEthernet0/0
```

Task 4

Configure RIPv2 on R3 and advertise its directly connected interfaces in this routing protocol. Configure RIPv2 MD5 authentication between R2 and R3, these routers should use "**cisco23**" as the password for this authentication.

On R3:

```
R3(config-subif)#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
```

To verify the configuration:

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

R     1.0.0.0/8 [120/2] via 10.2.2.2, 00:00:03, FastEthernet0/1
R     2.0.0.0/8 [120/1] via 10.2.2.2, 00:00:03, FastEthernet0/1
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
R        10.1.1.0/24 [120/1] via 10.2.2.2, 00:00:03, FastEthernet0/1
```

On R2 and R3:

```
Rx(config)#key chain TST23
Rx(config-keychain)#key 1
Rx(config-keychain-key)#key-string cisco23
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Rx(config)#int f0/1
Rx(config-if)#ip rip authentication key-chain TST23
Rx(config-if)#ip rip authentication mode md5
```

To verify the configuration:

On R2:

```
R2#Show ip protocol | s Default

  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    FastEthernet0/0       2     2                   TST
    FastEthernet0/1       2     2                   TST23
    Loopback0             2     2

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

R     1.0.0.0/8 [120/1] via 10.1.1.1, 00:00:10, FastEthernet0/0
R     3.0.0.0/8 [120/1] via 10.2.2.3, 00:00:04, FastEthernet0/1
```

When configuring MD5 authentication, the passwords and the key numbers MUST match on both routers, or else the routers will not exchange routing updates.

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 3
Configuring RIPv2 Updates

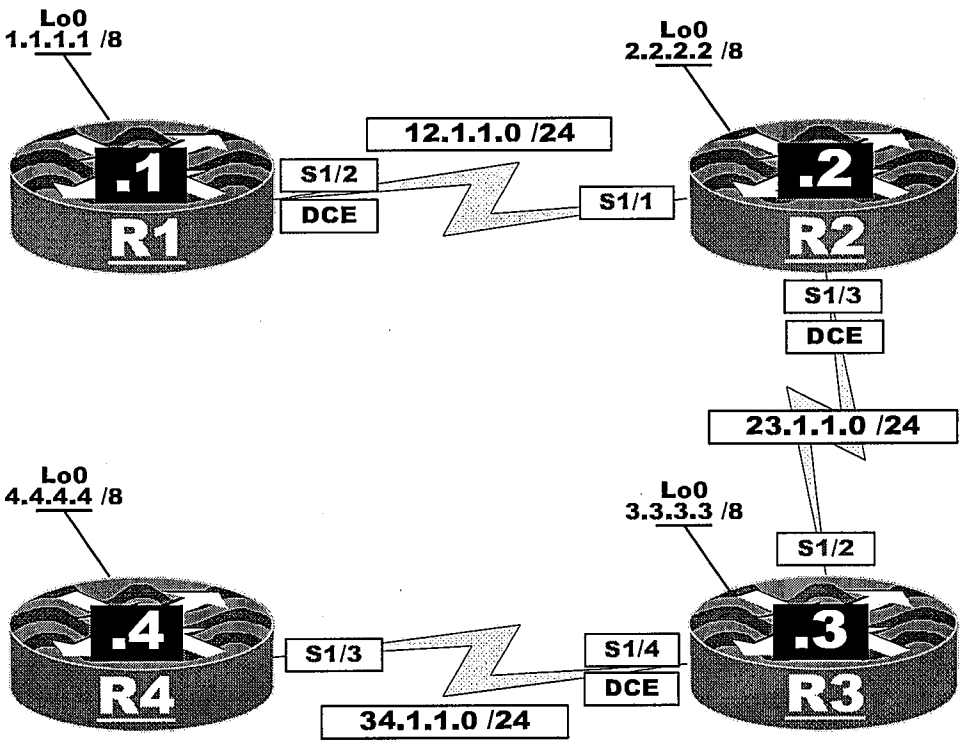

Task 1

Configure the routers based on the above topology. DO NOT configure any routing protocol.

To build the logical topology should be very easy, let's start with R1 and R2:

On R1:

```
R1(config)#int s1/2
R1(config-if)#clock rate 64000
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#int s1/1
R2(config-if)#ip addr 12.1.1.2 255.255.255.0
R2(config-if)#No shu

R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0

R2(config-if)#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
```

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)#Int Lo0
R3(config-if)#Ip addr 3.3.3.3 255.0.0.0

R3(config-if)#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
```

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)#Int lo0
R4(config-if)#Ip addr 4.4.4.4 255.0.0.0
```

To verify the configuration:

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
```


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

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R1#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              12.1.1.1        YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
```

On R2:

```
R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R2#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/1              12.1.1.2        YES manual up                    up
Serial1/3              23.1.1.2        YES manual up                    up
Loopback0              2.2.2.2         YES manual up                    up
```

On R3:

```
R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R3#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              23.1.1.3        YES manual up                    up
Serial1/4              34.1.1.3        YES manual up                    up
Loopback0              3.3.3.3         YES manual up                    up
```

On R4:

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

Interface              IP-Address      OK? Method Status                Protocol
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Serial1/3              34.1.1.4        YES manual up                    up
Loopback0              4.4.4.4         YES manual up                    up
```

Task 2

R1 and R2 should exchange their updates using a Multicast, whereas, R2 and R3 should exchange their updates via Unicast. These routers should advertise their Loopback0 interface in this routing domain.

By default, RIPv2 sends its updates to a reserved Multicast destination address of 224.0.0.9, this behavior can be changed such that the updates are sent to the neighbor's IP address (Unicast) or to the broadcast (255.255.255.255) destination address, in this task we only need to be concerned with R2 and R3's configuration, since R1 and R2 will send their updates to 224.0.0.9 (Multicast) address which is the default behavior.

On R1:

```
R1(config)#Router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#Netw 12.0.0.0
R1(config-router)#Netw 1.0.0.0
```

On R2:

```
R2(config)#router rip
R2(config-router)#passive-interface S1/3
R2(config-router)#Netw 12.0.0.0
R2(config-router)#Netw 23.0.0.0
R2(config-router)#Netw 2.0.0.0
```

The passive-interface command stops RIP from sending updates out of the specified interface. This router will still process the updates that it receives through its 51/3 interface.

Another way to handle this task is to configure "passive-interface default". This command stops sending RIP updates out of all interfaces, and then we can configure "neighbor" command to send updates ONLY to the specified neighbor/s. If the neighbor command is used without the passive-interface command, RIPv2 will send two updates, one to the Multicast and another to a Unicast destination.

```
R2(config-router)#neighbor 23.1.1.3
```


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

On R3:

```
R3(config)#router rip
R3(config-router)#passive-interface S1/2
R3(config-router)#neighbor 23.1.1.2
R3(config-router)#Netw 3.0.0.0
R3(config-router)#Netw 23.0.0.0
```

To verify the configuration:

On R2:

```
R2#Debug ip rip

RIP: sending v2 update to 224.0.0.9 via Serial1/1 (12.1.1.2)
RIP: build update entries
        2.0.0.0/8 via 0.0.0.0, metric 1, tag 0

RIP: sending v2 update to 23.1.1.3 via Serial1/3 (23.1.1.2)
RIP: build update entries
        1.0.0.0/8 via 0.0.0.0, metric 2, tag 0
        2.0.0.0/8 via 0.0.0.0, metric 1, tag 0
        12.1.1.0/24 via 0.0.0.0, metric 1, tag 0
```

NOTE: R2 and R3 are using Unicast to exchange updates, whereas, R1 and R2 use the reserved Multicast address of 224.0.0.9 to exchange updates.

To stop the debug:

On R2:

```
R2#Undebug all
All possible debugging has been turned off
```

Task 3

Configure R3 and R4 to exchange RIPv2 updates via Broadcast.

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config)#Router rip
R3(config-router)#Netw 34.0.0.0

R3(config)#int s1/4
R3(config-if)#ip rip v2-broadcast
```

On R4:

```
R4(config)#router rip
R4(config-router)#No au
R4(config-router)#ver 2
R4(config-router)#Network 34.0.0.0
R4(config-router)#Netw 4.0.0.0

R4(config)#int s1/3
R4(config-if)#ip rip v2-broadcast
```

The "IP rip v2-broadcast" is used to send version 2 updates to a destination address of 255.255.255.255; this is very handy when and if the hosts/routers do NOT listen to Multicast addresses.

To verify the configuration:

On R4:

```
R4#Debug ip rip
RIP protocol debugging is on

RIP: sending v2 update to 255.255.255.255 via Serial1/3 (34.1.1.4)
RIP: build update entries
        4.0.0.0/8 via 0.0.0.0, metric 1, tag 0
```

To stop the debug:

On R3:

```
R2#Undebug all
All possible debugging has been turned off
```

Do not erase the configuration, proceed to the next lab.

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

Lab 4
Injection of Default routes in RIPv2

Task 1

R1 should NOT be configured with any routing protocol. Network 12.0.0.0 should not be advertised by R2. Configure a static route for network 1.0.0.0 /8 on R2 with a next hop IP address of 12.1.1.1, and configure a static default route on R1 pointing to R2 (12.1.1.2).

On R1:

```
R1(config)#No router rip

R1(config)#ip route 0.0.0.0 0.0.0.0 12.1.1.2
```

On R2:

```
R2(config)#router rip
R2(config-router)#No Netw 12.0.0.0

R2(config)#ip route 1.0.0.0 255.0.0.0 12.1.1.1
```

Task 2

In order to provide full reachability between the routers, R2 should be configured to inject a default route into RIP routing domain. You should configure this task in four different ways.

Default route injection can be configured in four different ways, as follows:

1. **Default-information originate**
2. **Configuring a Static default route, and then redistributing the static route into RIP routing domain.**
3. **IP default-network command.**
4. **Network 0.0.0.0 with a static default route to Null0**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

First method:

On R2:

```
R2(config)#router rip
R2(config-router)#default-information originate
```

To verify the configuration:

On R4:

```
R4#Show ip route rip | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:12, Serial1/3
```

To test the configuration:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 104/106/108 ms
```

The second method:

Remove the previous command from R2 and "Clear ip route *" to implement the changes.

On R2:

```
R2(config)#Router rip
R2(config-router)#No default-information originate

R2#Clear ip route *
```

To verify the configuration

On R3 or R4:

```
Rx#Show ip route rip | Inc 0.0.0.0/0
Rx#
```

The output of the above show command verifies that the default route is no longer injected into the RIP routing domain.

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

Remove the static route for network 1.0.0.0 from R2.

On R2:

```
R2(config)#No ip route 1.0.0.0 255.0.0.0 10.1.12.1
```

Configure a static default route pointing to 12.1.1.1 and then redistribute the static route into RIP routing domain. Since the task specifically asks for a default route to be injected, redistributing the previous static route (ip route 1.0.0.0 255.0.0.0) into the RIP routing domain will NOT accomplish the task based on the requirements.

```
R2(config)#ip route 0.0.0.0 0.0.0.0 10.1.12.1

R2(config)#router rip
R2(config-router)#redistribute static
```

To verify the configuration:

On R4:

```
R4#Show ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:14, Serial1/3
```

To test the configuration:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 104/106/108 ms
```

The third method:

Remove the static default route that was configured in the previous step:

On R2:

```
R2(config)#No ip route 0.0.0.0 0.0.0.0 10.1.12.1
```

Remove the "redistribute static" that was configured in the previous step:

```
R2(config)#router rip
R2(config-router)#No redistribute static
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#Clear ip route *
```

To verify the configuration:

On R3 and R4:

```
Rx#Show ip route rip | inc 0.0.0.0/0
Rx#
```

Configure ip default-network command to inject a default route into RIP's routing domain.

On R2:

```
R2(config)#ip default-network 2.0.0.0
```

The above command injects a default route into RIP's routing domain. This tells the neighboring router/s that if they are not aware of a given route/s, they should send the traffic to network 2.0.0.0 (R2).

NOTE: The specified network address MUST be advertised in RIP, in this case network 2.0.0.0 was advertised in one of the previous task (Task 2), if the network specified in the "Ip default-network" command is NOT advertised, the local router will NOT advertise it as a candidate default, and the last point to remember is that the specified Network in this command must be a major network (Class A, B or C) and NOT a subnet of a major network or else the command will NOT work.

To verify and test the configuration:

On R4:

```
R4#Show ip route rip | inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:14, Serial1/3

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

We are getting unreachable messages and the ping is not successful. The actual echo messages are getting to R2, but R2 does not have a route to network 1.0.0.0 /8, therefore, to rectify this problem we have to configure a static route on R2 for network 1.0.0.0 /8 that is pointing to 12.1.1.1:

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

On R2:

```
R2(config)#Ip route 1.0.0.0 255.0.0.0  12.1.1.1
```

To test the configuration

On R4:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/100/140 ms
```

The forth method:

A static route to NULL0 is configured on R2 and then Network 0.0.0.0 is advertised in RIP. Before configuring this task the following command must be removed, this command was configured in the previous step:

On R2:

```
R2(config)#No ip default-network 2.0.0.0
```

To configure the forth method:

```
R2(config)#router rip
R2(config-router)#Network 0.0.0.0

R2(config)#ip route 0.0.0.0 0.0.0.0 null0
```

To verify the configuration:

On R4:

You may have to "Clear ip route *" on R2 to propagate the default route'

```
R4#Show ip route | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:02, Serial1/3

R4#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 = 56/95/148 ms
```

Task 3

Remove the following command/s:

- The last solution from the previous task
- The static route on R2 pointing to network 1.0.0.0 /8
- The static default route from R1 that is pointing to 12.1.1.2

On R2:

```
R2(config)#No Ip route 0.0.0.0 0.0.0.0 Null0
R2(config)#No ip route 1.0.0.0 255.0.0.0 10.1.12.1

R2(config)#Router rip
R2(config-router)#No Network 0.0.0.0
```

On R1:

```
R1(config)#No ip route 0.0.0.0 0.0.0.0 12.1.1.2
```

Task 4

Configure EIGRP AS 100 on R1 and R2; advertise the following networks in this routing domain:

- On R1 - Advertise its directly connected networks
- On R2 - Advertise the link to R1.

On R1:

```
R1(config)#Router eigrp 100
R1(config-router)#Netw 1.0.0.0
```


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

```
R1(config-router)#Netw 12.1.1.1 0.0.0.0
```

On R2:

```
R2(config-router)#Router eigrp 100
R2(config-router)#Netw 12.1.1.2 0.0.0.0
```

Note Eigrp, RIPv2 and BGP are classless routing protocols but by default they have a classful nature, meaning that they summarize networks at their major network boundary. This behavior can be disabled by configuring the "No Auto-summary" command. In IOS version 15, the auto-summary is disabled by default.

To verify the configuration:

```
R2#Show ip route eigrp | B Gate

D      1.0.0.0/8 [90/20640000] via 12.1.1.1, 00:03:09, Serial1/1
```

Task 5

Do not redistribute between the routing protocols and instead configure R2 to inject a default route into RIP and Eigrp's routing domain. Choose any method to accomplish this task. R3 and R4 must use the default route to reach network 1.0.0.0 /8.

On R2:

The following is one of many ways to inject a route into Eigrp's routing domain, these methods will be configured and tested in the Eigrp section:

```
R2(config)#int S1/1
R2(config-if)#ip summary-address eigrp 100 0.0.0.0 0.0.0.0

R2(config)#Router rip
R2(config-router)#Default-information originate

R2#Clear ip route *
```

To verify the configuration:

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4#Show ip route | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:12, Serial1/3
```

To test the configuration:

On R4:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 88/96/116 ms
```

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is 12.1.1.2 to network 0.0.0.0

D*    0.0.0.0/0 [90/21024000] via 12.1.1.2, 00:00:15, Serial1/2
```

To test the configuration:

On R1:

```
R1#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 = 84/100/128 ms
```

Task 6

Since the users that are connected to R3 and R4 will ONLY use the default route from the previous step to reach network 1.0.0.0 /8, configure R2 such that it injects the default route only if network 1.0.0.0 is in R2's routing table.

On R2:

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

```
R2(config)#access-list 1 permit 1.0.0.0 0.255.255.255

R2(config)#route-map TEST permit 10
R2(config-route-map)#match ip addr 1

R2(config)#router rip
R2(config-router)#default-information originate route-map TEST

R2#Clear ip route *
```

This configuration is called conditional default route injection. In this configuration a default route is injected ONLY if the condition of the route-map called "TEST" is true. The condition of the route map can ONLY be true if the network specified in the access-list can be matched in the routing table.

To test the configuration:

On R2:

```
R2#Show ip route | Inc 1.0.0.0/8

D       1.0.0.0/8 [90/20640000] via 12.1.1.1, 00:00:36, Serial1/1
```

NOTE: Now that R2 can see Network 1.0.0.0 in its routing table, the default route is injected, the following show command verifies the default route:

On R4:

```
R4#Show ip route | Inc 0.0.0.0/0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:24, Serial1/3
```

To test this task further, Loopback 0 is shutdown and the default route is once again verified on R4:

On R1:

```
R1(config)#int lo0
R1(config-if)#Shut
```

Note the above command shuts down Lo0 interface, and the following command verifies that R2 no longer has network 1.0.0.0 /8 in its routing table:

On R2:

```
R2#Show ip route | Inc 1.0.0.0/8
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2#
```

Since R2 does NOT have network 1.0.0.0 in its routing table the default route is removed:

On R4:

```
R4#Show ip route | Inc 0.0.0.0/0
R4#
```

To see the default route in R3 and R4's routing table, Lo0 interface is enabled:

On R1:

```
R1(config-if)#int lo0
R1(config-if)#No shut
```

Network 1.0.0.0 is up, and the following show command verifies that R2 can see network 1.0.0.0/8 in its routing table:

On R2:

```
R2#Show ip route | Inc 1.0.0.0/8

D      1.0.0.0/8 [90/20640000] via 12.1.1.1, 00:00:55, Serial1/1
```

R3 and R4 have the default route back:

On R3:

```
R3#Sh ip route rip | Inc 0.0.0.0/0

R*     0.0.0.0/0 [120/1] via 23.1.1.2, 00:00:11, Serial1/2
```

On R4:

```
R4#Sh ip route rip | Inc 0.0.0.0/0

R*     0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:02, Serial1/3
```

Do not erase the configuration, proceed to the next lab.

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

Lab 5 Filtering routes

Task 1

Configure the following Loopback interfaces on R2, this may override Lo0's IP Address:

Lo0 – 2.2.0.2 /24	Lo5 – 2.2.5.2 /24
Lo1 – 2.2.1.2 /24	Lo6 – 2.2.6.2 /24
Lo2 – 2.2.2.2 /24	Lo7 - 2.2.7.2 /24
Lo3 – 2.2.3.2 /24	Lo8 – 2.2.8.2 /24
Lo4 – 2.2.4.2 /24	Lo9 – 2.2.9.2 /24

On R2:

```
R2(config)#int lo0
R2(config-if)#ip addr 2.2.0.2 255.255.255.0

R2(config)#int lo1
R2(config-if)#ip addr 2.2.1.2 255.255.255.0

R2(config)#int lo2
R2(config-if)#ip addr 2.2.2.2 255.255.255.0

R2(config)#int lo3
R2(config-if)#ip addr 2.2.3.2 255.255.255.0

R2(config)#int lo4
R2(config-if)#ip addr 2.2.4.2 255.255.255.0

R2(config)#int lo5
R2(config-if)#ip addr 2.2.5.2 255.255.255.0

R2(config)#int lo6
R2(config-if)#ip addr 2.2.6.2 255.255.255.0

R2(config)#int lo7
R2(config-if)#ip addr 2.2.7.2 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#int lo8
R2(config-if)#ip addr 2.2.8.2 255.255.255.0

R2(config)#int lo9
R2(config-if)#ip addr 2.2.9.2 255.255.255.0
```

There is no need to re-advertise these loopback interfaces, since they are subsets of network 2.0.0.0, and since network 2.0.0.0 is advertised in RIPv2, these networks will also be advertised.

To verify the configuration:

On R4:

```
R4#Show ip route rip | B Gate
Gateway of last resort is 34.1.1.3 to network 0.0.0.0

R*    0.0.0.0/0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
      2.0.0.0/24 is subnetted, 10 subnets
R        2.2.0.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.1.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.2.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.3.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.4.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.5.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.6.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.7.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.8.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R        2.2.9.0 [120/2] via 34.1.1.3, 00:00:07, Serial1/3
R     3.0.0.0/8 [120/1] via 34.1.1.3, 00:00:07, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 34.1.1.3, 00:00:07, Serial1/3
```

Task 2

Remove the following configurations from R1 and R2:

- Eigrp 100 should be removed from both R1 and R2
- Remove the "default-information originate route-map TST" from R2
- Remove the route-map "TEST" from R2
- Remove the access-list 1 from R2

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

On R2:

```
R2(config)#No router eigrp 100
R2(config)#No access-list 1
R2(config)#No route-map TEST

R2(config)#Router rip
R2(config-router)#No Default-information originate
```

On R1:

```
R1(config)#No router eigrp 100
```

Task 3

Configure RIPv2 on R1 and advertise its directly connected networks. R2 should be configured to run RIPv2 on its S1/1 interface.

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 1.0.0.0
```

On R2:

```
R2(config)#Router rip
R2(config-router)#Netw 12.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route rip | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 10 subnets
R        2.2.0.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.1.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R        2.2.2.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.3.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.4.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.5.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.6.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.7.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.8.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R        2.2.9.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
R     3.0.0.0/8 [120/2] via 12.1.1.2, 00:00:19, Serial1/2
R     4.0.0.0/8 [120/3] via 12.1.1.2, 00:00:19, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 12.1.1.2, 00:00:19, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
R        34.1.1.0 [120/2] via 12.1.1.2, 00:00:19, Serial1/2
```

Task 4

Configure R1 such that it filters the even numbered subnets (2.2.0.0, 2.2.2.0, 2.2.4.0, 2.2.6.0, 2.2.8.0) of the networks added in the previous task. Use minimum number of lines in the access-list to accomplish this task

The following access-list filters the even numbered subnets of network 2.2.0.0 from R1's routing table:

On R1:

```
R1(config)#access-list 1 deny 2.2.0.0 0.0.254.0
R1(config)#access-list 1 permit any

R1(config)#router rip
R1(config-router)#distribute-list 1 in S1/2

R1#Clear ip route *
```

Explanation of the access-list:

Let's say the following access-list is configured:

Access-list 1 permit 2.0.0.0 0.255.255.255

The above access-list is permitting any network that matches the result of the anding of network 2.0.0.0 and its real mask.

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

In order for us to do the anding, we must get the regular mask; to get the regular mask, we can subtract 0.255.255.255 from 255.255.255.255, which results to 255.0.0.0; therefore, the anding is performed based on the following; remember anding is nothing but multiplication, we are multiplying the bits:

00000010.00000000.00000000.00000000 → Network (In this case 2.0.0.0)
11111111.11111111.11111111.00000000 → Regular Mask (In this case 255.255.255.0)

The result of this anding is 2.0.0.0, therefore, if the result of the AND matches network 2.0.0.0, the network is permitted, if NOT, it's denied.

In this task the access-list is configured as following:

Access-list 1 deny 2.2.0.0 0.0.254.0

Therefore, in order for us to do the anding, we have to get the regular mask for the inverse mask of 0.0.254.0, which is; 0.0.254.0 minus 255.255.255.255 which equals to 255.255.1.255

00000010.00000010.00000000.00000000 → (Network)
11111111.11111111.00000001.11111111 → (Regular mask)

The result of this AND is 2.2.0.0, therefore, if it matches 2.2.0.0, it should be denied. If the network is replaced with other subnets of network 2.2.0.0, we will see that the EVEN numbered third octet networks are all denied. For example, let's look at subnet 2.2.8.0:

00000010.00000010.00001000.00000000 → (Network)
11111111.11111111.00000001.11111111 → (Regular mask)

The result of the above AND is 2.2.0.0, therefore, it is denied. Let's look at a odd numbered third octete subnet, let's test subnet 2.2.1.0:

00000010.00000010.00000001.00000000 → (Network)
11111111.11111111.00000001.11111111 → (Regular mask)

The result of the above AND is 2.2.1.0, and this subnet does not match the access-list, therefore, it is not denied.

Finally, we should permit the rest of the networks so networks (34.1.1.0 /24, 23.1.1.0 /24, 4.0.0.0 /8, 3.0.0.0 /8, and the ODD numbered third octet of 2.2.0.0 network) are still advertised.

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1#Clear ip route *

R1#Show ip route rip | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 5 subnets
R        2.2.1.0 [120/1] via 12.1.1.2, 00:00:20, Serial1/2
R        2.2.3.0 [120/1] via 12.1.1.2, 00:00:20, Serial1/2
R        2.2.5.0 [120/1] via 12.1.1.2, 00:00:20, Serial1/2
R        2.2.7.0 [120/1] via 12.1.1.2, 00:00:20, Serial1/2
R        2.2.9.0 [120/1] via 12.1.1.2, 00:00:20, Serial1/2
R     3.0.0.0/8 [120/2] via 12.1.1.2, 00:00:20, Serial1/2
R     4.0.0.0/8 [120/3] via 12.1.1.2, 00:00:20, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 12.1.1.2, 00:00:20, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
R        34.1.1.0 [120/2] via 12.1.1.2, 00:00:20, Serial1/2
```

Task 5

Configure R3 such that it filters the odd 3rd octet subnets of network 2.2.0.0. Use minimum number of lines in the access-list to accomplish this task

Let's verify the existing routing table on R3:

On R3:

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

R     1.0.0.0/8 [120/2] via 23.1.1.2, 00:00:20, Serial1/2
      2.0.0.0/24 is subnetted, 10 subnets
R        2.2.0.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.1.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.2.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.3.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.4.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.5.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.6.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
```


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

```
R        2.2.7.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.8.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R        2.2.9.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
R     4.0.0.0/8 [120/1] via 34.1.1.4, 00:00:27, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 23.1.1.2, 00:00:20, Serial1/2
```

On R3:

```
R3(config)#access-list 1 deny 2.2.1.0  0.0.254.255
R3(config)#access-list 1 permit any

R3(config)#router rip
R3(config-router)#distribute-list 1 in S1/2

R3#Clear ip route *
```

To Verify the configuration:

```
R3#Sh ip route rip | Inc 2.2.
Gateway of last resort is not set

R     1.0.0.0/8 [120/2] via 23.1.1.2, 00:00:04, Serial1/2
      2.0.0.0/24 is subnetted, 5 subnets
R        2.2.0.0 [120/1] via 23.1.1.2, 00:00:04, Serial1/2
R        2.2.2.0 [120/1] via 23.1.1.2, 00:00:04, Serial1/2
R        2.2.4.0 [120/1] via 23.1.1.2, 00:00:04, Serial1/2
R        2.2.6.0 [120/1] via 23.1.1.2, 00:00:04, Serial1/2
R        2.2.8.0 [120/1] via 23.1.1.2, 00:00:04, Serial1/2
R     4.0.0.0/8 [120/1] via 34.1.1.4, 00:00:11, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
R        12.1.1.0 [120/1] via 23.1.1.2, 00:00:04, Serial1/2
```

Task 6

R2 should be configured such that R4 filters network 1.0.0.0 /8.

On R2:

```
R2(config)#access-list 1 permit 1.0.0.0 0.255.255.255
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#router rip
R2(config-router)#offset-list 1 out 13 S1/3

R2#Clear ip route *
```

To Verify the configuration;

On R3:

```
R3#Show ip route rip | Inc 1.0.0.0/8

R     1.0.0.0/8 [120/15] via 23.1.1.2, 00:00:17, Serial1/2
```

NOTE: The output of the above command reveals that from R3's perspective network 1.0.0.0/8 has a hop count of 15; since in RIP a hop count of 16 is unreachable, R4 will reject the route.

On R4:

```
R4#Show ip route | Inc 1.0.0.0/8
R4#
```

Things to remember about offset-lists:

1. **It only works with RIP and Eigrp routing protocols.**
2. **In RIP, it's used to add to the existing hop count.**
3. **If all networks are to be affected, then "offset-list 0" can be used.**

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

CCIE Foundation

5.0

www.MicronicsTraining.com

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

EIGRP

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1
Configuring EIGRP and adjusting timers

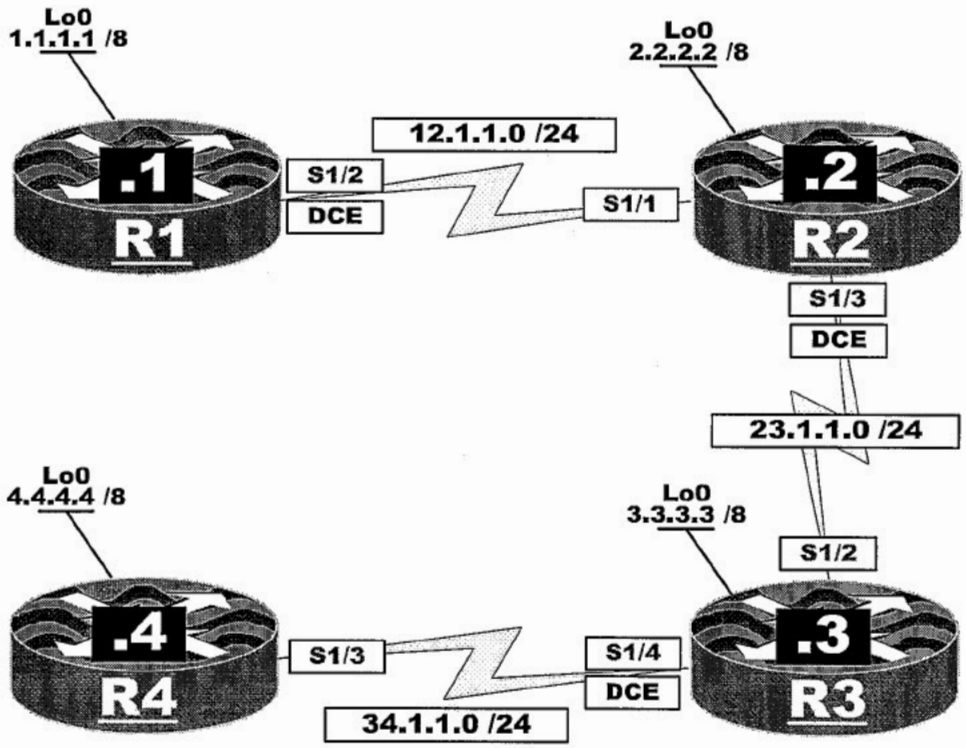

Task 1

Configure the routers based on the above diagram. DO NOT configure Eigrp.

To build the logical topology should be very easy, let's start with R1 and R2:

On R1:

```
R1(config)#int s1/2
R1(config-if)#clock rate 64000
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```


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

On R2:

```
R2(config)#int s1/1
R2(config-if)#ip addr 12.1.1.2 255.255.255.0
R2(config-if)#No shu

R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0

R2(config-if)#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
```

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)#Int Lo0
R3(config-if)#Ip addr 3.3.3.3 255.0.0.0

R3(config-if)#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
```

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)#Int lo0
R4(config-if)#Ip addr 4.4.4.4 255.0.0.0
```

To verify the configuration:

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R1#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              12.1.1.1        YES manual up                    up
Loopback0              1.1.1.1         YES manual up                    up
```

On R2:

```
R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R2#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/1              12.1.1.2        YES manual up                    up
Serial1/3              23.1.1.2        YES manual up                    up
Loopback0              2.2.2.2         YES manual up                    up
```

On R3:

```
R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R3#Show ip int br | Exc unass

Interface              IP-Address      OK? Method Status                Protocol
Serial1/2              23.1.1.3        YES manual up                    up
Serial1/4              34.1.1.3        YES manual up                    up
Loopback0              3.3.3.3         YES manual up                    up
```

On R4:

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


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

```
Interface              IP-Address      OK? Method Status                Protocol
Serial1/3              34.1.1.4        YES manual up                    up
Loopback0              4.4.4.4         YES manual up                    up
```

Task 2

Configure EIGRP AS 100 on all routers and advertise their directly connected networks.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#netw 12.1.1.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

NOTE: Since R1 is the edge router, it could have been configured using the "Network 0.0.0.0 0.0.0.0" statement instead of the two network statements. With "Network 0.0.0.0" statement, you are running Eigrp on all existing and future interface/s in the specified AS. Some engineers configure the above two statements using the following:

```
R1(config)#router eigrp 100
R1(config-router)#netw 12.1.1.0 0.0.0.255
R1(config-router)#netw 1.0.0.0 0.255.255.255
```

In the above example, you are running Eigrp AS 100 on all interfaces in Networks 1.0.0.0/8 and 12.1.1.0/24 network. But the best way to configure Eigrp is to be very specific, like the configuration that was performed on R1 for this task.

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#netw 12.1.1.2 0.0.0.0
R2(config-router)#netw 2.2.2.2 0.0.0.0
R2(config-router)#netw 23.1.1.2 0.0.0.0
```

On R3:

```
R3(config-if)#router eigrp 100
R3(config-router)#netw 23.1.1.3 0.0.0.0
R3(config-router)#netw 34.1.1.3 0.0.0.0
R3(config-router)#netw 3.3.3.3 0.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R4:

```
R4(config)#router eigrp 100
R4(config-router)#netw 34.1.1.4 0.0.0.0
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, 00:01:48, Serial1/2
D      3.0.0.0/8 [90/21152000] via 12.1.1.2, 00:01:00, Serial1/2
D      4.0.0.0/8 [90/21664000] via 12.1.1.2, 00:00:13, Serial1/2
       23.0.0.0/24 is subnetted, 1 subnets
D         23.1.1.0 [90/21024000] via 12.1.1.2, 00:01:45, Serial1/2
       34.0.0.0/24 is subnetted, 1 subnets
D         34.1.1.0 [90/21536000] via 12.1.1.2, 00:00:53, Serial1/2
```

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D      1.0.0.0/8 [90/21664000] via 34.1.1.3, 00:00:37, Serial1/3
D      2.0.0.0/8 [90/21152000] via 34.1.1.3, 00:00:37, Serial1/3
D      3.0.0.0/8 [90/20640000] via 34.1.1.3, 00:00:37, Serial1/3
       12.0.0.0/24 is subnetted, 1 subnets
D         12.1.1.0 [90/21536000] via 34.1.1.3, 00:00:37, Serial1/3
       23.0.0.0/24 is subnetted, 1 subnets
D         23.1.1.0 [90/21024000] via 34.1.1.3, 00:00:37, Serial1/3
```

Task 3

These routers should send their hellos every 30 seconds and if they don't receive a hello from their neighboring router for 120 seconds, the neighbor adjacency should be terminated.

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

On R1:

```
R1(config)#Int S1/2
R1(config-if)#ip hello-interval eigrp 100 30
R1(config-if)#ip hold-time eigrp 100 120
```

On R2:

```
R2(config)#Int S1/1
R2(config-if)#ip hello-interval eigrp 100 30
R2(config-if)#ip hold-time eigrp 100 120

R2(config)#Int S1/3
R2(config-if)#ip hello-interval eigrp 100 30
R2(config-if)#ip hold-time eigrp 100 120
```

On R3:

```
R3(config)#Int S1/4
R3(config-if)#ip hello-interval eigrp 100 30
R3(config-if)#ip hold-time eigrp 100 120

R3(config)#Int S1/2
R3(config-if)#ip hello-interval eigrp 100 30
R3(config-if)#ip hold-time eigrp 100 120
```

On R4:

```
R4(config)#Int S1/3
R4(config-if)#ip hello-interval eigrp 100 30
R4(config-if)#ip hold-time eigrp 100 120
```

To verify the configuration:

To see the hello interval for a given interface:

On R4:

```
R4#Show ip eigrp interfaces detail | S Se1/3

Se1/3               1        0/0        23       5/190        286
0
  Hello-interval is 30, Hold-time is 120
  Split-horizon is enabled
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 6/7
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Topology-ids on interface - 0
  Authentication mode is not set
```

The output of the above show command reveals the "Hello interval", "Hold-time", the state of "Eigrp's Spplit-horizon", and if authentication is enabled.

```
R4#Show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)

H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   34.1.1.3                Se1/3            99 00:10:44   23   1140  0  18
```

The first column "H":
This column specifies the order in which the neighbor/s were found, this number starts with "0".

The second column "Address":
This is the IP address of the neighboring router.

The third column "Interface":
The interface through which the local router receives hello packets from the peer.

The forth column "Holdtime":
This column specifies the length in time that the local router will wait to hear from the peer before declaring it down.

The fifth column "Uptime":
This column specifies how long the local router has known the neighbor.

The sixth column "SRTT":
Smooth round-trip time. The amount of time required for an Eigrp packet to be sent to this neighbor and for the local router to receive an ACK of that packet.

The seventh column "RTO":
Retransmission timeout. This column specifies the amount of time Eigrp waits before resending a packet from the retransmission queue to a neighbor.

The eighth column "Q Count":
The number of Eigrp packets in the Queue waiting to be sent.

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

The ninth "Seq Num":
This column specifies the sequence number of the last packet that was received from the neighbor.

Task 4

Ensure that these routers use bandwidth only and NOT bandwidth + delay to calculate their composite metric for a given network.

This task is accomplished by setting the "K" values; the command to configure the "K" values is: "Metric weight tos k1 k2 k3 k4 k5"

This command is entered in router Eigrp config mode, where "tos" value is always set to 0, and the default K values which are multipliers are: K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 0

The K values effect the result of the following formula:
[(K1 * Bandwidth) + (K2 * bandwidth / 256 – load) + (K3 * delay)] * 256

If K5 does not equal to zero, then an additional operation is performed as follows:
[K5 / K4 + reliability]

Therefore, if only the "bandwidth" is what should be considered, then, we must set the value of "K3" to zero. But remember that this change must occur on all routers or else the routers won't establish neighbor adjacency.

Enter the following command to display the default values for the "K" multipliers:

On R1:

```
R1#Show ip protocol | Inc K

  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
```

Before changing the "K" multipliers, the routing table is checked for network 2.0.0.0/8:

On R1:

```
R1#Show ip route eigrp | Inc 2.0.0.0/8

D     2.0.0.0/8 [90/20640000] via 12.1.1.2, 00:16:01, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The composite metric is 20640000.

On All Routers:

```
Rx(config)#Router eigrp 100
Rx(config-router)#metric weight 0 1 0 0 0 0
```

Once the K values are changes, the routers will reestablish their neighbor adjacency based on the new values. You should see the following console message on R1:

```
%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 12.1.1.2 (Serial1/2) is up:
new adjacency
```

To Verify the configuration:

On R1:

```
R1#Show ip protocol | Inc EIGRP metric

  EIGRP metric weight K1=1, K2=0, K3=0, K4=0, K5=0

R1#Show ip route eigrp | Inc 2.0.0.0/8

D      2.0.0.0/8 [90/20000000] via 12.1.1.2, 00:00:56, Serial1/2
```

Note before the change the composite metric for network 2.0.0.0/8 was 20640000 and after the change its 20000000.

Task 5

DO NOT ERASE THE CONFIGURATION, CONTINUE TO THE NEXT LAB.

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

Lab 2 - EIGRP Metric

Task 1

Remove Eigrp and configure RIPv2 and disable auto-summary on R1 and advertise its directly connected networks in RIP routing domain.

On R1:

```
R1(config)#No router eigrp 100

R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#netw 12.0.0.0
R1(config-router)#netw 1.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip protocol | Inc Routing Protocol

Routing Protocol is "rip"
```

Task 2

Reconfigure R2 to advertise its S1/1 and Loopback 0 interfaces in RIPv2; configure a loopback interface (Lo1) on this router and assign an IP address of 22.2.2.2 /8; advertise S1/3 and lo1 interface of this router in EIGRP 100.

On R2

```
R2(config-subif)#router rip
R2(config-router)#ver 2
R2(config-router)# No au
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#netw 12.0.0.0
R2(config-router)#netw 2.0.0.0
R2(config-router)#passive-interface F0/0

R2(config)#Int lo1
R2(config-if)#ip addr 22.2.2.2 255.0.0.0

R2(config-router)#router eigrp 100
R2(config-router)#No netw 2.2.2.2 0.0.0.0
R2(config-router)#No netw 12.1.1.2 0.0.0.0
R2(config-router)#netw 22.2.2.2 0.0.0.0
```

To verify the configuration:

On R2:

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

R     1.0.0.0/8 [120/1] via 12.1.1.1, 00:00:16, Serial1/1

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

D     3.0.0.0/8 [90/20000000] via 23.1.1.3, 00:08:00, Serial1/3
D     4.0.0.0/8 [90/20000000] via 23.1.1.3, 00:08:00, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
D        34.1.1.0 [90/20000000] via 23.1.1.3, 00:08:00, Serial1/3
```

On R4:

```
R4#Show ip route eigrp | B Gate
Gateway of last resort is not set

D     3.0.0.0/8 [90/20000000] via 34.1.1.3, 00:08:38, Serial1/3
D     22.0.0.0/8 [90/20000000] via 34.1.1.3, 00:02:14, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/20000000] via 34.1.1.3, 00:08:38, Serial1/3
```

Task 3

Perform a mutual redistribution between RIPv2 and EIGRP such that R1 sees the EIGRP routes with the correct hop count. For example, R1 should see network 3.0.0.0 two hops

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

away, and network 4.0.0.0 three hops away. Future networks should be redistributed with a hop count of 1.

To resolve this task, an access-list is configured to permit every network in the EIGRP routing domain, a route-map is configured to reference a specific network and assign a hop count to it using the "Set metric" command.

On R2:

```
R2(config)#access-list 3 permit 3.0.0.0 0.255.255.255
R2(config)#access-list 4 permit 4.0.0.0 0.255.255.255
R2(config)#access-list 34 permit  10.1.34.0 0.0.0.255
R2(config)#access-list 23 permit  10.1.23.0 0.0.0.255
R2(config)#access-list 22 permit 22.0.0.0 0.255.255.255

R2(config)#route-map TEST permit 10
R2(config-route-map)#match ip addr 3
R2(config-route-map)#set metric 2

R2(config)#route-map TEST permit 20
R2(config-route-map)#match ip addr 4
R2(config-route-map)#set metric 3

R2(config)#route-map TEST permit 30
R2(config-route-map)#match ip addr 23
R2(config-route-map)#set metric 1

R2(config)#route-map TEST permit 40
R2(config-route-map)#match ip addr 34
R2(config-route-map)#set metric 2

R2(config)#route-map TEST permit 50
R2(config-route-map)#match ip addr 22
R2(config-route-map)#set metric 1

R2(config)#route-map TEST permit 60
R2(config-route-map)#set metric 1
```

NOTE: The "route-map TEST permit 60" must be configured for the future networks that maybe redistributed into RIP, these networks will be assigned a hop count of 1 as per requirement of this task.

```
R2(config-router)#router rip
R2(config-router)#redistribute eigrp 100 route-map TEST

R2(config)#router eigrp 100
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#redistribute rip metric 1500 20000 255 1 1500
```

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 [120/1] via 12.1.1.2, 00:00:15, Serial1/2
R     3.0.0.0/8 [120/2] via 12.1.1.2, 00:00:15, Serial1/2
R     4.0.0.0/8 [120/3] via 12.1.1.2, 00:00:15, Serial1/2
R     22.0.0.0/8 [120/1] via 12.1.1.2, 00:00:15, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
R        23.1.1.0 [120/1] via 12.1.1.2, 00:00:15, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
R        34.1.1.0 [120/2] via 12.1.1.2, 00:00:15, Serial1/2
```

Task 4

DO NOT ERASE THE CONFIGURATION, CONTINUE TO THE NEXT LAB.

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

Lab 3 - EIGRP Summarization

Task 1

Remove RIPv2 and configure Eigrp 100 on R1; configure the following loopback interfaces on R1 and advertise them in this routing domain:

Interface	IP Address
Lo0	1.1.0.1 /24
Lo1	1.1.1.1 /24
Lo2	1.1.2.1 /24
Lo3	1.1.3.1 /24

On R1:

```
R1(config)#int lo0
R1(config-if)#ip addr 1.1.0.1 255.255.255.0

R1(config-if)#int lo1
R1(config-if)#ip addr 1.1.1.1 255.255.255.0

R1(config-if)#int lo2
R1(config-if)#ip addr 1.1.2.1 255.255.255.0

R1(config-if)#int lo3
R1(config-if)#ip addr 1.1.3.1 255.255.255.0

R1(config)#No router rip

R1(config)#router eigrp 100
R1(config-router)#netw 12.1.1.1 0.0.0.0
R1(config-router)#netw 1.0.0.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 2

Remove the following from R2:

- Remove RIP
- Remove access-list 3, 4, 22, 23 and 34
- Remove the route-map "TEST"
- RIP should NOT be redistributed into Eigrp
- Advertise R2's S1/1 in Eigrp AS 100 routing domain
- Remove the "Metric weight" command from all routers
- Remove Lo1 interface and configure a Lo0 interface with an IP address of 2.2.2.2 /8; advertise Lo0 interface in Eigrp AS 100 routing domain

On R2

```
R2(config)#No router rip

R2(config)#No access-list 3
R2(config)#No access-list 4
R2(config)#No access-list 22
R2(config)#No access-list 23
R2(config)#No access-list 34

R2(config)#No route-map TEST

R2(config)#router eigrp 100
R2(config-router)#No redistribute rip

R2(config)#No int lo1

R2(config)#Int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0

R2(config)#Router eigrp 100
R2(config-router)#Netw 12.1.1.2 0.0.0.0
R2(config-router)#Netw 2.2.2.2 0.0.0.0
```

On All Routers:

```
Rx(config)#Router eigrp 100
Rx(config-router)#No metric weight
```


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

Task 3

Configure the following Loopback interfaces on R2 and advertise them in Eigrp 100. This configuration will override the IP address of R2's Lo0 interface.

Interface	IP Address
Lo0	2.2.4.2 /24
Lo1	2.2.5.2 /24
Lo2	2.2.6.2 /24
Lo3	2.2.7.2 /24

On R2:

```
R2(config)#int lo0
R2(config-if)#ip addr 2.2.4.2 255.255.255.0

R2(config-if)#int lo1
R2(config-if)#ip addr 2.2.5.2 255.255.255.0


R2(config-if)#int lo2
R2(config-if)#ip addr 2.2.6.2 255.255.255.0

R2(config-if)#int lo3
R2(config-if)#ip addr 2.2.7.2 255.255.255.0

R2(config)#router eigrp 100
R2(config-router)#netw 2.0.0.0
```

Task 4

Configure the following Loopback interfaces on R3 and advertise them in Eigrp 100. This configuration will override the IP address of R3's Lo0 interface.

Interface	IP Address
Lo0	3.3.8.3 /24
Lo1	3.3.9.3 /24
Lo2	3.3.10.3 /24
Lo3	3.3.11.3 /24

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3(config)#Int Lo0
R3(config-if)#ip addr 3.3.8.3 255.255.255.0

R3(config)#Int Lo1
R3(config-if)#ip addr 3.3.9.3 255.255.255.0

R3(config)#Int Lo2
R3(config-if)#ip addr 3.3.10.3 255.255.255.0

R3(config)#Int Lo3
R3(config-if)#ip addr 3.3.11.3 255.255.255.0

R3(config)#router eigrp 100
R3(config-router)#Network 3.0.0.0
```

Task 5

Configure the following Loopback interfaces on R4 and advertise them in Eigrp 100. This configuration will override the IP address of R4's Lo0 interface.

Interface	IP Address
Lo0	4.4.12.4 /24
Lo1	4.4.13.4 /24
Lo2	4.4.14.4 /24
Lo3	4.4.15.4 /24

On R4:

```
R4(config)#Int Lo0
R4(config-if)#Ip address 4.4.12.4 255.255.255.0

R4(config)#Int Lo1
R4(config-if)#Ip address 4.4.13.4 255.255.255.0

R4(config)#Int Lo2
R4(config-if)#Ip address 4.4.14.4 255.255.255.0

R4(config)#Int Lo3
```


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

```
R4(config-if)#Ip address 4.4.15.4 255.255.255.0

R4(config)#router eigrp 100
R4(config-router)#Network 4.0.0.0
```

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 [90/21664000] via 34.1.1.3, 00:08:22, Serial1/3
D        1.1.1.0 [90/21664000] via 34.1.1.3, 00:08:22, Serial1/3
D        1.1.2.0 [90/21664000] via 34.1.1.3, 00:08:22, Serial1/3
D        1.1.3.0 [90/21664000] via 34.1.1.3, 00:08:22, Serial1/3
      2.0.0.0/24 is subnetted, 4 subnets
D        2.2.4.0 [90/21152000] via 34.1.1.3, 00:05:17, Serial1/3
D        2.2.5.0 [90/21152000] via 34.1.1.3, 00:05:17, Serial1/3
D        2.2.6.0 [90/21152000] via 34.1.1.3, 00:05:17, Serial1/3
D        2.2.7.0 [90/21152000] via 34.1.1.3, 00:05:17, Serial1/3
      3.0.0.0/24 is subnetted, 4 subnets
D        3.3.8.0 [90/20640000] via 34.1.1.3, 00:03:44, Serial1/3
D        3.3.9.0 [90/20640000] via 34.1.1.3, 00:03:44, Serial1/3
D        3.3.10.0 [90/20640000] via 34.1.1.3, 00:03:44, Serial1/3
D        3.3.11.0 [90/20640000] via 34.1.1.3, 00:03:44, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/21536000] via 34.1.1.3, 00:08:22, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/21024000] via 34.1.1.3, 00:08:22, Serial1/3
```

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        2.2.4.0 [90/20640000] via 12.1.1.2, 00:06:20, Serial1/2
D        2.2.5.0 [90/20640000] via 12.1.1.2, 00:06:20, Serial1/2
D        2.2.6.0 [90/20640000] via 12.1.1.2, 00:06:20, Serial1/2
D        2.2.7.0 [90/20640000] via 12.1.1.2, 00:06:20, Serial1/2
      3.0.0.0/24 is subnetted, 4 subnets
D        3.3.8.0 [90/21152000] via 12.1.1.2, 00:04:47, Serial1/2
D        3.3.9.0 [90/21152000] via 12.1.1.2, 00:04:47, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
D        3.3.10.0 [90/21152000] via 12.1.1.2, 00:04:47, Serial1/2
D        3.3.11.0 [90/21152000] via 12.1.1.2, 00:04:47, Serial1/2
      4.0.0.0/24 is subnetted, 4 subnets
D        4.4.12.0 [90/21664000] via 12.1.1.2, 00:03:37, Serial1/2
D        4.4.13.0 [90/21664000] via 12.1.1.2, 00:03:37, Serial1/2
D        4.4.14.0 [90/21664000] via 12.1.1.2, 00:03:37, Serial1/2
D        4.4.15.0 [90/21664000] via 12.1.1.2, 00:03:37, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/21024000] via 12.1.1.2, 00:09:47, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
D        34.1.1.0 [90/21536000] via 12.1.1.2, 00:09:25, Serial1/2
```

Task 6

Configure R2 such that it advertises a summary route for its Loopback interfaces to R3.

Unlike OSPF, in EIGRP, summarization can be performed on any router; in Eigrp, summarization must be configured in the interface configuration mode which is very specific, meaning that only the neighbor/s through that interface will receive the summary route.

On R2:

```
R2(config)#int S1/3
R2(config-if)#ip summary-address eigrp 100 2.2.4.0 255.255.252.0
```

To verify the configuration:

On R2:

```
R2#Show ip route eigrp | Inc 2.2.4.0

D       2.2.4.0/22 is a summary, 00:01:17, Null0
```

◄ **Note, when summarization is configured, Eigrp will inject a discard route for loop avoidance**

On R3:

```
R3#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 [90/21152000] via 23.1.1.2, 00:11:43, Serial1/2
```


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

```
D        1.1.1.0 [90/21152000] via 23.1.1.2, 00:11:43, Serial1/2
D        1.1.2.0 [90/21152000] via 23.1.1.2, 00:11:43, Serial1/2
D        1.1.3.0 [90/21152000] via 23.1.1.2, 00:11:43, Serial1/2
     2.0.0.0/22 is subnetted, 1 subnets
D        2.2.4.0 [90/20640000] via 23.1.1.2, 00:00:53, Serial1/2
     4.0.0.0/24 is subnetted, 4 subnets
D        4.4.12.0 [90/20640000] via 34.1.1.4, 00:05:56, Serial1/4
D        4.4.13.0 [90/20640000] via 34.1.1.4, 00:05:56, Serial1/4
D        4.4.14.0 [90/20640000] via 34.1.1.4, 00:05:56, Serial1/4
D        4.4.15.0 [90/20640000] via 34.1.1.4, 00:05:56, Serial1/4
     12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/21024000] via 23.1.1.2, 00:11:43, Serial1/2
```

Task 7

R2 should summarize its loopback interfaces and advertise a single summary route plus all the specific routes of the summary to R1.

The leak-map was introduced in IOS 12.3(14)T and it must be configured under the physical interface and NOT a sub-interface.

Configuring leak-map option allows us to advertise a component route (One or more specific networks of the summary route) that would otherwise be suppressed by manual summarization.

There are three rules to remember:

- **If the leak-map is configured to reference a route-map that does not exist, only the summary route is advertised and the more specific routes are suppressed.**
- **If the leak-map is configured to reference a route-map and the route-map is referencing an access-list that does NOT exist, then the summary route plus all the specific routes are advertised.**
- **If the leak-map is configured to reference a route-map and the route-map matches on an access-list, all the permitted networks by the access-list will be advertised along with the summary route.**

On R2:

```
R2(config)#int S1/1
R2(config-if)#ip summary-address eigrp 100 2.2.4.0 255.255.252.0 leak-map TEST21
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config)#route-map TEST21 permit 10
```

Note the leak-map is referencing a route-map and the route-map is NOT referencing an access-list, therefore, the summary route plus all the specific routes are advertised. If the route-map references an access-list that does NOT exist, the same will result.

To Verify the configuration:

On R1:

```
R1#Show ip route eigrp | B Gate
Gateway of last resort is not set

      2.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
D        2.2.4.0/22 [90/20640000] via 12.1.1.2, 00:00:48, Serial1/2
D        2.2.4.0/24 [90/20640000] via 12.1.1.2, 00:00:28, Serial1/2
D        2.2.5.0/24 [90/20640000] via 12.1.1.2, 00:00:28, Serial1/2
D        2.2.6.0/24 [90/20640000] via 12.1.1.2, 00:00:28, Serial1/2
D        2.2.7.0/24 [90/20640000] via 12.1.1.2, 00:00:28, Serial1/2
      3.0.0.0/24 is subnetted, 4 subnets
D        3.3.8.0 [90/21152000] via 12.1.1.2, 00:11:39, Serial1/2
D        3.3.9.0 [90/21152000] via 12.1.1.2, 00:11:39, Serial1/2
D        3.3.10.0 [90/21152000] via 12.1.1.2, 00:11:39, Serial1/2
D        3.3.11.0 [90/21152000] via 12.1.1.2, 00:11:39, Serial1/2
      4.0.0.0/24 is subnetted, 4 subnets
D        4.4.12.0 [90/21664000] via 12.1.1.2, 00:10:29, Serial1/2
D        4.4.13.0 [90/21664000] via 12.1.1.2, 00:10:29, Serial1/2
D        4.4.14.0 [90/21664000] via 12.1.1.2, 00:10:29, Serial1/2
D        4.4.15.0 [90/21664000] via 12.1.1.2, 00:10:29, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/21024000] via 12.1.1.2, 00:16:39, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
D        34.1.1.0 [90/21536000] via 12.1.1.2, 00:16:17, Serial1/2
```

Task 8

Configure R4 to summarize its loopback interfaces and advertise a summary route plus network 4.4.13.0 /24 to R3.

On R4:

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

```
R4(config)#Int S1/3
R4(config-if)#ip summary-addres eigrp 100 4.4.12.0 255.255.252.0 leak-map TEST43

R4(config)#route-map TEST43 permit 10
R4(config-route-map)#match ip addr 13

R4(config)#access-list 13 permit 4.4.13.0 0.0.0.255
```

To verify the configuration:

On R4:

```
R4#Show ip route eigrp | Inc Null0

D       4.4.12.0/22 is a summary, 00:01:49, Null0
```

The output of the above show command reveals the Null0 discard route that is created whenever summarization is configured.

On R3:

```
R3#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 [90/21152000] via 23.1.1.2, 00:19:34, Serial1/2
D        1.1.1.0 [90/21152000] via 23.1.1.2, 00:19:34, Serial1/2
D        1.1.2.0 [90/21152000] via 23.1.1.2, 00:19:34, Serial1/2
D        1.1.3.0 [90/21152000] via 23.1.1.2, 00:19:34, Serial1/2
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.4.0 [90/20640000] via 23.1.1.2, 00:08:44, Serial1/2
      4.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
D        4.4.12.0/22 [90/20640000] via 34.1.1.4, 00:01:21, Serial1/4
D        4.4.13.0/24 [90/20640000] via 34.1.1.4, 00:00:57, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/21024000] via 23.1.1.2, 00:19:34, Serial1/2
```

Task 9

Configure R3 to summarize its loopback interfaces and advertise a single summary plus 3.3.10.0 /24 and 3.3.11.0 /24 to R4.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R3:

```
R3(config)#int S1/4
R3(config-if)#ip summary-addr eigrp 100 3.3.8.0 255.255.252.0 leak-map TEST34

R3(config)#route-map TEST34 permit 10
R3(config-route-map)#match ip addr 34

R3(config)#access-list 34 permit 3.3.10.0 0.0.0.255
R3(config)#access-list 34 permit 3.3.11.0 0.0.0.255
```

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 [90/21664000] via 34.1.1.3, 00:22:02, Serial1/3
D        1.1.1.0 [90/21664000] via 34.1.1.3, 00:22:02, Serial1/3
D        1.1.2.0 [90/21664000] via 34.1.1.3, 00:22:02, Serial1/3
D        1.1.3.0 [90/21664000] via 34.1.1.3, 00:22:02, Serial1/3
      2.0.0.0/22 is subnetted, 1 subnets
D        2.2.4.0 [90/21152000] via 34.1.1.3, 00:11:11, Serial1/3
      3.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D        3.3.8.0/22 [90/20640000] via 34.1.1.3, 00:00:47, Serial1/3
D        3.3.10.0/24 [90/20640000] via 34.1.1.3, 00:00:24, Serial1/3
D        3.3.11.0/24 [90/20640000] via 34.1.1.3, 00:00:24, Serial1/3
      4.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
D        4.4.12.0/22 is a summary, 00:03:48, Null0
      12.0.0.0/24 is subnetted, 1 subnets
D        12.1.1.0 [90/21536000] via 34.1.1.3, 00:22:02, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
D        23.1.1.0 [90/21024000] via 34.1.1.3, 00:22:02, Serial1/3
```

Task 10

Configure R1 to inject a default route into Eigrp 100 routing domain. Do NOT use any global or router configuration command as part of the solution to accomplish this task.

© 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 0.0.0.0 0.0.0.0
```

To Verify the configuration:

To see the discard route:

On R1:

```
R1#Show ip route eigrp | Inc Null0

D*    0.0.0.0/0 is a summary, 00:03:19, Null0
```

On R2:

```
R2#Sh ip route eigrp | Inc Serial1/1

D*     0.0.0.0/0 [90/20640000] via 12.1.1.1, 00:00:37, Serial1/1
```

NOTE: All the specific routes are suppressed and ONLY the summary route is advertised.

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 - EIGRP Authentication & Advanced Configuration

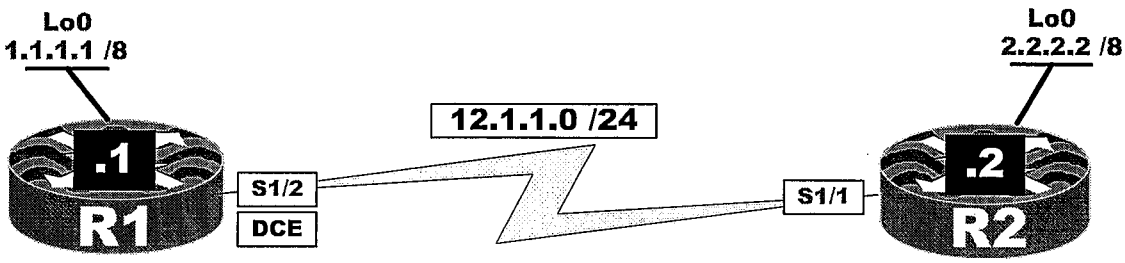

Task 1

Configure the routers based on the above diagram.

On R1:

```
R1(config)#int s1/2
R1(config-if)#clock rate 64000
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```

On R2:

```
R2(config)#int s1/1
R2(config-if)#ip addr 12.1.1.2 255.255.255.0
R2(config-if)#No shu

R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0
```

To verify the configuration:

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

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms
```

Task 2

Configure EIGRP in AS100 on both routers and advertise their directly connected networks.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#netw 12.1.1.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

On R2:

```
R2(config)#router eigrp 100
R2(config-router)#netw 2.2.2.2 0.0.0.0
R2(config-router)#netw 12.1.1.2 0.0.0.0
```

To verify the configuration:

On R1:

```
R1#Show ip route Eigrp | B Gate

D     2.0.0.0/8 [90/20640000] via 12.1.1.2, 00:00:20, Serial1/2
```

Task 3

Enable Eigrp authentication between R1 and R2; use "**cisco**" as the password.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

EIGRP uses key chain base authentication just like RIPv2.

On Both Routers:

```
Rx(config)#key chain TEST
Rx(config-keychain)#key 1
Rx(config-keychain-key)#key-string cisco
```

On R1:

```
R1(config)#int S0/0.12
R1(config-subif)#ip authentication key-chain eigrp 100 TEST
R1(config-subif)#ip authentication mode eigrp 100 md5
```

On R2:

```
R2(config)#int S0/0.21
R2(config-subif)#ip authentication key-chain eigrp 100 TEST
R2(config-subif)#ip authentication mode eigrp 100 md5
```

To verify the configuration:

On R1:

```
R1#Show ip route eigrp

D    2.0.0.0/8 [90/2297856] via 10.1.12.2, 00:00:27, Serial0/0.12

R1#Show ip eigrp interfaces detail S1/2 | Inc Authenti

  Authentication mode is md5,  key-chain is "TEST"
```

Task 4

Configure R2 such that it advertises EIGRP routes with greater than 110 hops as unreachable.

On R2:

```
R2(config)#router eigrp 100
```


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

```
R2(config-router)#metric maximum-hops 110
```

The above command will advertise the routes with a hop count higher than 110 as unreachable, the default setting is 100 hops, and it can be increased to 255.

Task 5

Configure a static route on R2 for network 3.0.0.0 /8 reachable by R2 and redistribute the static route into EIGRP routing process. Ensure that R1 can see this network in its routing table.

On R2:

```
R2(config)#ip route 3.0.0.0 255.0.0.0 Null0
```

NOTE: When redistributing static, connected, or routes from another EIGRP routing domain you do not have to specify the metric.

```
R2(config)#router eigrp 100
R2(config-router)#redistribute static
```

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, 00:04:30, Serial1/2
D EX   3.0.0.0/8 [170/20512000] via 12.1.1.2, 00:01:27, Serial1/2
```

NOTE: The routes that are redistributed into Eigrp routing domain show up as "EX" meaning that they are external, external to eigrp routing domain. In Eigrp routing protocols the routes can be either internal or external, the internal routes have an administrative distance of 90, whereas, the external routes have an administrative distance of 170.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 6

On R1, change the administrative distance of all existing and future external networks to 110.

Before we change the administrative distance of the external routes, we should check the default:

On R1:

```
R1#Show ip route eigrp | Inc D EX

D EX  3.0.0.0/8 [170/20512000] via 12.1.1.2, 00:05:02, Serial1/2
```

Note, by default the administrative distance of Eigrp external routes is 170. The following command changes the administrative distance of Eigrp external routes to 110.

Remember, since administrative distance can NOT be advertised to another router, this will ONLY affect the router on which it's configured on.

On R1:

```
R1(config)#router eigrp 100
R1(config-router)#distance eigrp 90 110
```

The first value after the "distance eigrp" command specifies the administrative distance of all Eigrp internal routes and the second value specifies the administrative distance of all Eigrp external routes.

To Verify the configuration:

On R1:

```
R1#Show ip route eigrp | Inc D EX

D EX  3.0.0.0/8 [110/20512000] via 12.1.1.2, 00:00:47, Serial1/2
```

Task 7

Ensure that R1 never uses more than 45% of the bandwidth for EIGRP updates; you should use an EIGRP specific command to accomplish this task.

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

On R1:

```
R1(config-router)#int S1/2
R1(config-subif)#ip bandwidth-percent eigrp 100 45
```

By default, Eigrp can use up to %50 of the bandwidth on a given interface, the above command can be used to change the default, in this case, it's changed to 45 percent.

Task 8

You were asked to limit Eigrp's Bandwidth utilization on R2 to 386 Kbits/sec, but you can not remember the EIGRP specific command, use another method to accomplish the same task.

On R2:

```
R2(config-router)#int S1/1
R2(config-subif)#bandwidth 772
```

By default, Eigrp can utilize upto 50% of the bandwidth, if you multiply the desired value by two (In this case 772) and set the bandwidth of the interface to that value, Eigrp will use half of that value which is the desired value.

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 - EIGRP Stub

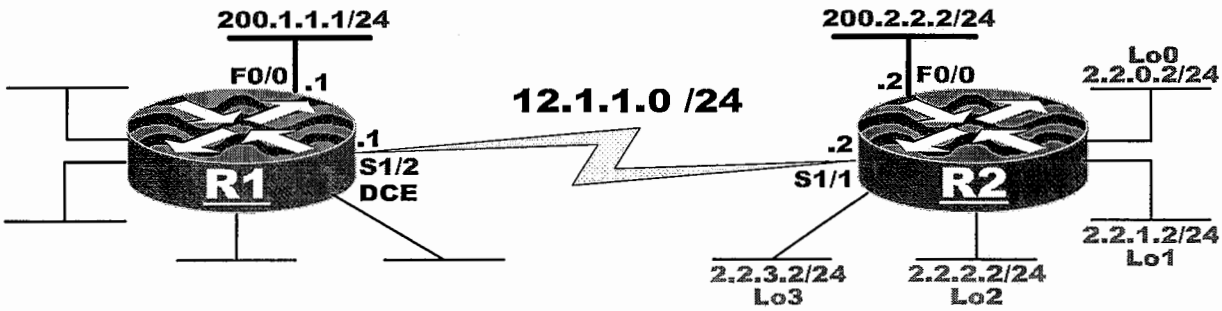

Task 1

Configure the routers based on the above topology. DO NOT configure any routing protocols.

On R1:

```
R1(config)#int s1/2
R1(config-if)#clock rate 64000
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 1.1.0.1 255.255.255.0

R1(config)#Int lo1
R1(config-if)#ip addr 1.1.1.1 255.255.255.0

R1(config)#Int lo2
R1(config-if)#ip addr 1.1.1.1 255.255.255.0

R1(config)#Int lo3
R1(config-if)#ip addr 1.1.1.1 255.255.255.0

R1(config)#Int F0/0
R1(config-if)#ip addr 200.1.1.1 255.255.255.0
R1(config-if)#No shut
```


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

On R2:

```
R1(config)#int s1/1
R1(config-if)#ip addr 12.1.1.2 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 2.2.0.2 255.255.255.0

R1(config)#Int lo1
R1(config-if)#ip addr 2.2.1.2 255.255.255.0

R1(config)#Int lo2
R1(config-if)#ip addr 2.2.2.2 255.255.255.0

R1(config)#Int lo3
R1(config-if)#ip addr 2.2.3.2 255.255.255.0

R1(config)#Int F0/0
R1(config-if)#ip addr 200.2.2.2 255.255.255.0
R1(config-if)#No shut
```

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

SW1(config)#int F0/2
SW1(config-if)#Swi mode acc
SW1(config-if)#Swi acc vlan 200
SW1(config-if)#No shut
```

Task 2

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 method to accomplish this task. **DO NOT run EIGRP on the F0/0 interfaces of these two routers.**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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:

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:

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

Let's configure R2 with EIGRP named configuration:

```
R2(config)#Router eigrp tst
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
```


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

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 3

Configure R1 and R2 to summarize their Loopback interfaces in EIGRP.

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 4

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**

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
```


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

```
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
```

Task 5

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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:

```
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
```


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

Task 6

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
  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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 7

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
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
```


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

```
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 8

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:

```
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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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 9

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
```

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 10

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.

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

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
```

Task 11

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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
```

Task 12

Erase the startup config and reload the routers before proceeding to the next lab.

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

CCIE Foundation

5.0

www.MicronicsTraining.com

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

OSPF

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Lab 1 – Advertising Networks

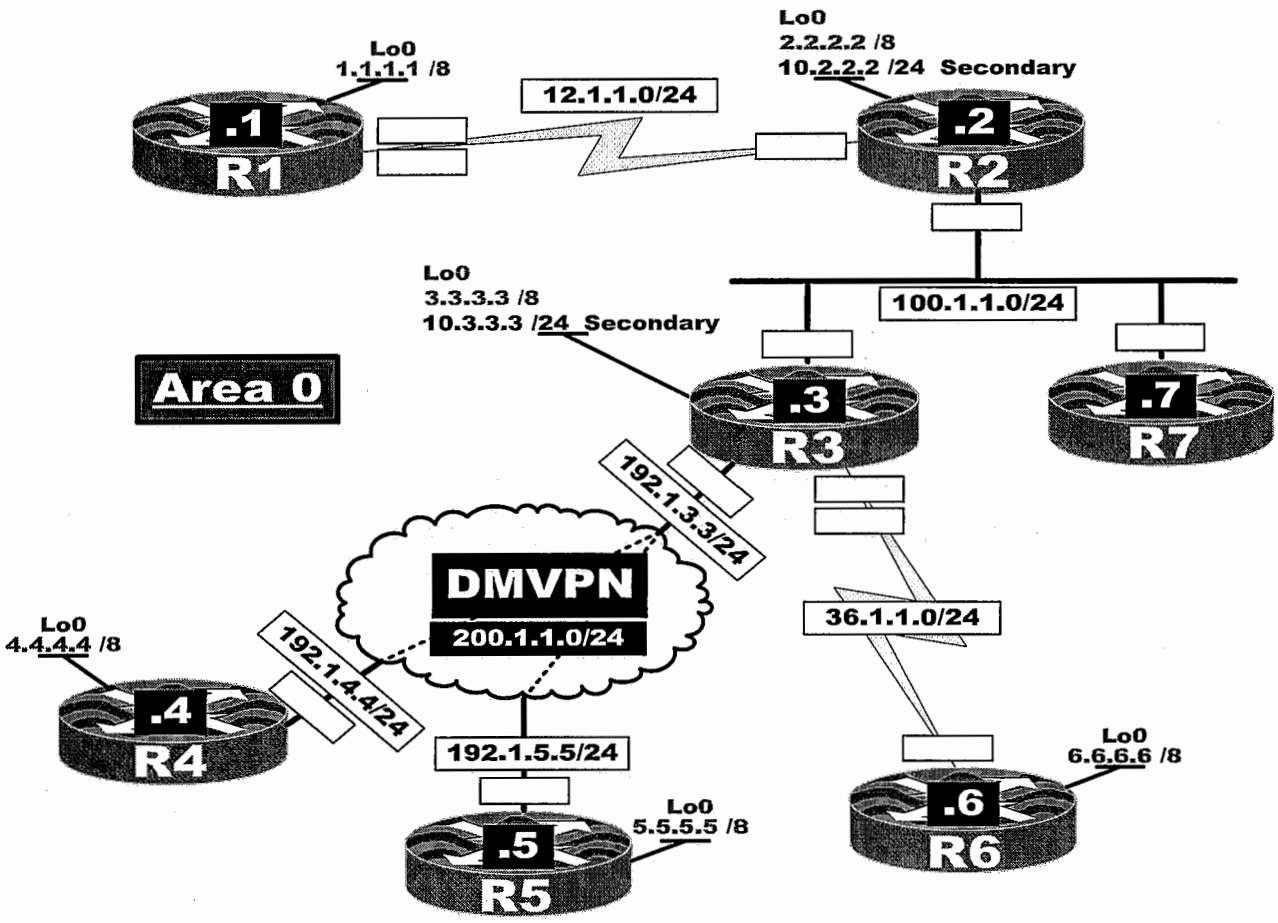

Task 1

Configure the routers and the switches based on the above diagram; DO NOT configure OSPF. When configuring the DMVPN, R3 should be the NHS and R4 and R5 should be the NHCs or the spokes; the hub router (R3) should not be configured with any static mappings, the spokes should be configured in a point-to-point manner.

Let's start with configuration of SW1; we can see that routers R2, R3 and R7 should be in the same VLAN, since the VLAN ID is not specified we will use VLAN 100.

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

On SW1:

```
SW1(config)#Int range F0/2-3,F0/7
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc v 100
SW1(config-if-range)#No shu
```

We can also see that R3, R4 and R5 are using their F0/1 interface for the DMVPN network, this means that we need to configure the appropriate ports on SW2 for this connection, since the lab has not specified any IP address for SW2, we will use the same subnet with a ".10" as the host portion. Let's configure SW2 for this connection:

On SW2:

```
SW2(config)#Int range f0/3-5
SW2(config-if-range)#No swi
SW2(config-if-range)#No shut

SW2(config)#Int f0/3
SW2(config-if)#ip addr 192.1.3.10 255.255.255.0

SW2(config)#Int f0/4
SW2(config-if)#ip addr 192.1.4.10 255.255.255.0

SW2(config)#int f0/5
SW2(config-if)#ip addr 192.1.5.10 255.255.255.0

SW2(config)#ip routing
```

Now that we are done with the switching aspect of this lab, let's start the configuration of the routers starting with R1's connection to R2:

On R1:

```
R1(config)#Int S1/2
R1(config-if)#IP address 12.1.1.1 255.255.255.0
R1(config-if)#clock rate 64000
R1(config-if)#No shut

R1(config)#Int Lo0
R1(config-if)#IP address 1.1.1.1 255.0.0.0
```

On R2:

```
R2(config)#Int S1/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-subif)#IP addres 12.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int Lo0
R2(config-if)#IP address 2.2.2.2 255.0.0.0
R2(config-if)#IP address 10.2.2.2 255.255.255.0 secondary

R2(config)#Int F0/0
R2(config-if)#IP address 123.1.1.2 255.255.255.0
R2(config-if)#No shut
```

To verify and test the configuration:

On R1:

```
R1#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

Let's configure R2's F0/0, R3's F0/0 and R7's G0/0 connection:

On R3:

```
R3(config)#Int F0/0
R3(config-if)#IP address 100.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Int Lo0
R3(config-if)#IP address 3.3.3.3 255.0.0.0
R3(config-if)#IP address 10.3.3.3 255.255.255.0 Secondary
```

On R7:

```
R7(config)#Int G0/0
R7(config-if)#IP address 100.1.1.7 255.255.255.0
R7(config-if)#No shu
```

On R2:

```
R2(config)#Int F0/0
R2(config-subif)#IP address 100.1.1.2 255.255.255.0
R2(config-if)#No shut
```


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

To verify and test the configuration:

On R2:

```
R2#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 = 1/1/1 ms

R2#Ping 100.1.1.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.1.1.7, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/4 ms
```

Let's configure the serial connection between R3 and R6:

On R3:

```
R3(config)#int s1/6
R3(config-if)#clock rate 64000
R3(config-if)#ip addr 36.1.1.3 255.255.255.0
R3(config-if)#No shut
```

On R6:

```
R6(config)#int s1/3
R6(config-if)#ip addr 36.1.1.6 255.255.255.0
R6(config-if)#No shut

R6(config)#Int lo0
R6(config-if)#ip addr 6.6.6.6 255.0.0.0
```

To verify and test the configuration:

On R6:

```
R6#Ping 36.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 36.1.1.3, 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 = 28/28/32 ms
```

The only other connection left to configure are the routers connecting to the DMVPN network, we have already configured SW2 for this connection, so let's start with R3's configuration:

On R3:

The following IP address is configured based on the diagram:

```
R3(config)#Int f0/1
R3(config-if)#ip addr 192.1.3.3 255.255.255.0
R3(config-if)#No shut
```

We need to have reachability to the endpoints of our DMVPN network, these are 192.1.4.4/24 and 192.1.5.5/24; to accomplish this we can simply configure a static route for each endpoint and point the static route to the IP address of the switchport of SW2 to which the local router is connected to:

```
R3(config)#IP route 192.1.4.4 255.255.255.255 192.1.3.10
R3(config)#IP route 192.1.5.5 255.255.255.255 192.1.3.10
```

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 192.1.3.3 255.255.255.255 192.1.4.10
R4(config)#ip route 192.1.5.5 255.255.255.255 192.1.4.10

R4(config)#int Lo0
R4(config-if)#ip addr 4.4.4.4 255.0.0.0
```

On R5:

```
R7(config)#Int F0/1
R7(config-if)#IP address 192.1.5.5 255.255.255.0
R7(config-if)#No shut

R5(config)#ip route 192.1.3.3 255.255.255.255 192.1.5.10
R5(config)#ip route 192.1.4.4 255.255.255.255 192.1.5.10

R5(config)#int Lo0
R5(config-if)#ip addr 5.5.5.5 255.0.0.0
```


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

To verify and test the configuration:

On R3:

```
R3#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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

R3#Ping 192.1.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R4:

```
R4#Ping 192.1.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.1.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Now that these three routers have reachability to each other, we can configure the DMVPN portion of this diagram:

On R3:

```
R3(config)#int tunnel 1
R3(config-if)#IP addr 200.1.1.3 255.255.255.0
R3(config-if)#tunnel source F0/1
R3(config-if)#tunnel mode gre multipoint
R3(config-if)#ip nhrp network-id 333
```

On R4:

```
R4(config)#int tunnel 1
R4(config-if)#ip addr 200.1.1.4 255.255.255.0
R4(config-if)#tunnel source f0/1
R4(config-if)#tunnel destination 192.1.3.3
R4(config-if)#ip nhrp network-id 444
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config-if)#ip nhrp map 200.1.1.3 192.1.3.3
R4(config-if)#ip nhrp nhs 200.1.1.3
```

On R5:

```
R5(config)#int tunnel 1
R5(config-if)#ip addr 200.1.1.5 255.255.255.0
R5(config-if)#tunnel source f0/1
R5(config-if)#tunnel destination 192.1.3.3
R5(config-if)#ip nhrp network-id 444
R5(config-if)#ip nhrp map 200.1.1.3 192.1.3.3
R5(config-if)#ip nhrp nhs 200.1.1.3
```

To verify and test the configuration:

On R3:

```
R3#Show ip nhrp

200.1.1.4/32 via 200.1.1.4
   Tunnel1 created 00:01:22, expire 01:59:18
   Type: dynamic, Flags: unique registered
   NBMA address: 192.1.4.4
200.1.1.5/32 via 200.1.1.5
   Tunnel1 created 00:02:27, expire 01:57:32
   Type: dynamic, Flags: unique registered
   NBMA address: 192.1.5.5

R3#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

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 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R4:

```
R4#Ping 200.1.1.5
```


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

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms
```

Task 2

Configure OSPF area 0 on R2, R3's F0/0 and R7's G0/0 interfaces. The router-id of these routers should be based on the following:
R2 – 0.0.0.2
R3 – 0.0.0.3
R7 – 0.0.0.7

Before we run OSPF on any layer two protocol we must know the capabilities of that protocol from OSPF's perspective; the following states the important aspects of an Ethernet segment:

- **OSPF's network type is "Broadcast".**
- **Timers are set to 10/40, meaning that the OSPF Hellos are exchanged every 10 seconds, and the Dead interval is set to 40 seconds.**
- **There is a DR/BDR election: The router with the highest OSPF interface priority is elected as the DR, the default priority is set to 1, and the range is 0 – 255. A priority of 0 means that the local router will NOT participate in the DR/BDR election.**
- **Next-hop is the IP address of the router that originated the route.**
- **The routers use Multicast addresses of 224.0.0.5 and 224.0.0.6.**

Let's configure and verify all the above items:

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Network 100.1.1.2 0.0.0.0 area 0
```

In the above configuration the OSPF router-id is set to "0.0.0.2"; in OSPF the router id uniquely identifies the router within the entire routing domain and MUST be unique within the entire OSPF routing domain. OSPF router id is a 32 bit dotted decimal value, IT IS NOT AN IP ADDRESS. However, since IP addresses are also 32 bit dotted decimal values, an IP address can also be used as the OSPF router id. If the router-id is not configured, numerically highest IP address of any Loopback interface will be chosen as the router id, if one does not exist, then the highest IP address configured on the local router will be chosen as the OSPF

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

router id.

In OSPF the "Network" statement can be configured in different ways, the above network statement could have been configured in the following ways:

Network 0.0.0.0 0.0.0.0 area 0
This means that the existing and future interface/s that have an IP address will run in area 0. Rememeber that if an interface is redistributed into OSPF routing protocol, the redistributed interface will show up in the routing table as an intra-area route and NOT external, because intra-area is preferred over inter-area which is preferred over external routes.

Network 100.0.0.0 0.255.255.255 area 0
The above network statement means that any subnet within the major network (100.0.0.0/8) should run in area 0.

Network 100.1.1.0 0.0.0.255
The above network statement means that any host within network 100.1.1.0/24 will run in area 0.

Network 100.1.1.1 0.0.0.0 area 0
This is by far the best way to run OSPF on a given interface; it is the most specific manner in which OSPF can be configured on a given interface.

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Network 100.1.1.3 0.0.0.0 area 0
```

You should see the following console message stating that the local router has established an OSPF adjacency with 0.0.0.2 (R2) and they are in full state:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

On R7:

```
R7(config)#Router ospf 1
R7(config-router)#router-id 0.0.0.7
R7(config-router)#Network 100.1.1.7 0.0.0.0 area 0
```

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on GigabitEthernet0/0 from LOADING
to FULL, Loading Done
```


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

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on GigabitEthernet0/0 from LOADING
to FULL, Loading Done
```

To verify the configuration:

On R2:

```
R2#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.3           1   FULL/BDR        00:00:33    100.1.1.3       FastEthernet0/0
0.0.0.7           1   FULL/DROTHER    00:00:36    100.1.1.7       FastEthernet0/0
```

We can see that the local router has established two adjacencies, one with 0.0.0.3 (R3) and the second one with 0.0.0.7 (R7). It looks like R3 is the BDR, whereas, R7 is neighter the DR nor the BDR.

```
R2#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Fa0/0        1     0               100.1.1.2/24       1     DR    2/2
```

The output of the above show command reveals the following:

- OSPF Process is 1
- Local router's F0/0 interface is configured in area 0
- The IP address in this area is 100.1.1.2/24
- The OSPF cost of this interface is 1
- The local router is the DR for this segment
- The local router has two neighbors

Let's verify the timers:

```
R2#Show ip ospf interface F0/0 Inc Network|Timer|ID

  Internet Address 100.1.1.2/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 0.0.0.2, Network Type BROADCAST, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
  Designated Router (ID) 0.0.0.2, Interface address 100.1.1.2
  Backup Designated router (ID) 0.0.0.3, Interface address 100.1.1.3
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

You can see that the router id is 0.0.0.2, the network type is "Broadcast", and the cost of this interface is 1, and the state is DR. One of the most important points of this output is the Hello and Dead intervals.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Show ip ospf route

            OSPF Router with ID (0.0.0.2) (Process ID 1)

                Base Topology (MTID 0)

    Area BACKBONE(0)

    Intra-area Route List
*   100.1.1.0/24, Intra, cost 1, area 0, Connected
      via 100.1.1.2, FastEthernet0/0
```

Task 3

Configure OSPF Area 0 on the Serial1/2 interface of R1 and the S1/1 interface of R2. R2 should also run OSPF area 0 on its Loopback 0 interface (Primary and the Secondary IP addresses). You should NOT use a Network statement to accomplish this task. The loopback interface/s of these two routers should be advertised with their correct mask.

On R1:

```
R1(config)#int S1/2
R1(config-if)#ip ospf 1 area 0
```

OSPF can be configured in two different modes: Router configuration mode or interface configuration mode. This task states that a "Network" command can NOT be used; therefore, the interface configuration mode is used to run OSPF area 0 on the S1/2 interface of R1. When OSPF is configured directly on the interface, the IOS will automatically start the OSPF process for you, let's verify:

```
R1#Show run | S router ospf

router ospf 1
```

Let's set the router id of the router:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
```


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

```
R1(config)#int Lo0
R1(config-if)#ip ospf network point-to-point
R1(config-if)#ip ospf 1 area 0
```

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 S1/1
R2(config-if)#ip ospf 1 area 0
```

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/  -        00:00:35    12.1.1.2        Serial1/2
```

Task 4

Configure the point-to-point frame-relay connection between R3 and R6 in OSPF area 0. R6 and R3 should also run OSPF area 0 on their Loopback 0 interfaces, R3 should run OSPF area 0 on the Primary and the Secondary IP addresses of its loopback0 interface. Ensure that the loopback interfaces are advertised with their correct mask.

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Network 36.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 3.3.3.3 0.0.0.0 area 0
R3(config-router)#Network 10.3.3.3 0.0.0.0 area 0

R3(config)#Int Lo0
R3(config-if)#IP ospf network point-to-point
```

On R6:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R6(config)#Router ospf 1
R6(config-router)#router-id 0.0.0.6
R6(config-router)#Network 6.6.6.6 0.0.0.0 area 0
R6(config-router)#Network 36.1.1.6 0.0.0.0 area 0

R6(config)#Int Lo0
R6(config-if)#IP ospf network point-to-point
```

To verify the configuration:

On R6:

```
R6#Show ip route ospf | B Gate
Gateway of last resort is not set

O     1.0.0.0/8 [110/1564] via 36.1.1.3, 00:00:07, Serial1/3
O     2.0.0.0/8 [110/783] via 36.1.1.3, 00:00:07, Serial1/3
O     3.0.0.0/8 [110/782] via 36.1.1.3, 00:00:07, Serial1/3
      10.0.0.0/24 is subnetted, 2 subnets
O        10.2.2.0 [110/783] via 36.1.1.3, 00:00:07, Serial1/3
O        10.3.3.0 [110/782] via 36.1.1.3, 00:00:07, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/1563] via 36.1.1.3, 00:00:07, Serial1/3
      100.0.0.0/24 is subnetted, 1 subnets
O        100.1.1.0 [110/782] via 36.1.1.3, 00:00:07, Serial1/3
```

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 100.1.1.2, 00:04:47, FastEthernet0/0
O     2.0.0.0/8 [110/2] via 100.1.1.2, 00:05:07, FastEthernet0/0
O     6.0.0.0/8 [110/782] via 36.1.1.6, 00:01:05, Serial1/6
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O        10.2.2.0/24 [110/2] via 100.1.1.2, 00:05:07, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/782] via 100.1.1.2, 00:04:57, FastEthernet0/0
```

Task 5

Configure OSPF area 0 on network 200.1.1.0 /24; you should also run OSPF area 0 on the loopback 0 interfaces of R4 and R5. These loopback interfaces should be advertised with

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

their correct mask. The tunnel interfaces in DMVPN network should be configured with OSPF network type of Non-broadcast.
Without changing the network types ensure that these routers form an OSPF adjacency and exchange routes.

On R3, R4 and R5:

```
Rx(config)#int tunnel 1
Rx(config-if)#IP ospf network non-broadcast
```

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 200.1.1.4 0.0.0.0 area 0

R4(config)#Int Lo0
R4(config-if)#IP Ospf network point-to-point
```

On R5:

```
R5(config)#Router ospf 1
R5(config-router)#router-id 0.0.0.5
R5(config-router)#Network 5.5.5.5 0.0.0.0 area 0
R5(config-router)#Network 200.1.1.5 0.0.0.0 area 0

R5(config)#Int Lo0
R5(config-if)#IP Ospf network point-to-point
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Network 200.1.1.3 0.0.0.0 area 0
```

To verify the configuration:

On R3:

```
R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.6           0   FULL/  -        00:00:36    36.1.1.6        Serial1/6
0.0.0.2           1   FULL/DR         00:00:37    100.1.1.2       FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
0.0.0.7           1   FULL/DROTHER    00:00:31    100.1.1.7       FastEthernet0/0
```

R3 has NOT established an adjacency with R4 or R5; let's verify:

Since the network type is "non-broadcast" and we can not change the network type, then the only other choice is to use unicast using the "neighbor commands on R3. We also need to configure the tunnel interface of the spoke routers (R4 and R5) to OSPF priority of zero.

On R4 and R5:

```
R5(config)#int tunnel 1
R5(config-if)#ip ospf priority 0
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Neighbor 200.1.1.4
R3(config-router)#Neighbor 200.1.1.5
```

You should see the following console messages stating that neighbor adjacencies were established with R4 and R5.

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Tunnel1 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.5 on Tunnel1 from LOADING to FULL,
Loading Done
```

On R6:

```
R6#Show ip route ospf | B Gate
Gateway of last resort is not set

O     1.0.0.0/8 [110/1564] via 36.1.1.3, 00:03:53, Serial1/3
O     2.0.0.0/8 [110/783] via 36.1.1.3, 00:03:53, Serial1/3
O     3.0.0.0/8 [110/782] via 36.1.1.3, 00:03:53, Serial1/3
O     4.0.0.0/8 [110/1782] via 36.1.1.3, 00:01:28, Serial1/3
O     5.0.0.0/8 [110/1782] via 36.1.1.3, 00:01:28, Serial1/3
      10.0.0.0/24 is subnetted, 2 subnets
O        10.2.2.0 [110/783] via 36.1.1.3, 00:03:53, Serial1/3
O        10.3.3.0 [110/782] via 36.1.1.3, 00:03:53, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/1563] via 36.1.1.3, 00:03:53, Serial1/3
      100.0.0.0/24 is subnetted, 1 subnets
O        100.1.1.0 [110/782] via 36.1.1.3, 00:03:53, Serial1/3
```


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

```
O      200.1.1.0/24 [110/1781] via 36.1.1.3, 00:03:53, Serial1/3
```

Task 6

R2 should be configured such that the secondary IP address of its loopback 0 interface is NOT advertised in OSPF. DO NOT use any global configuration mode command as part of the solution to accomplish this task.

When configuring OSPF in the interface configuration mode, the "Secondaries none" keyword can be used to not advertise the secondary IP address/es.

To see the secondary IP address of R2's Lo0 interface:

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, 00:30:19, Serial1/2
O     3.0.0.0/8 [110/783] via 12.1.1.2, 00:08:52, Serial1/2
O     4.0.0.0/8 [110/1783] via 12.1.1.2, 00:06:27, Serial1/2
O     5.0.0.0/8 [110/1783] via 12.1.1.2, 00:06:27, Serial1/2
O     6.0.0.0/8 [110/1564] via 12.1.1.2, 00:08:52, Serial1/2
      10.0.0.0/24 is subnetted, 2 subnets
O        10.2.2.0 [110/782] via 12.1.1.2, 00:30:19, Serial1/2
O        10.3.3.0 [110/783] via 12.1.1.2, 00:08:52, Serial1/2
      36.0.0.0/24 is subnetted, 1 subnets
O        36.1.1.0 [110/1563] via 12.1.1.2, 00:08:52, Serial1/2
      100.0.0.0/24 is subnetted, 1 subnets
O        100.1.1.0 [110/782] via 12.1.1.2, 00:30:19, Serial1/2
O     200.1.1.0/24 [110/1782] via 12.1.1.2, 00:08:52, Serial1/2
```

On R2:

```
R2(config)#Int Lo0
R2(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, 00:31:56, Serial1/2
O     3.0.0.0/8 [110/783] via 12.1.1.2, 00:10:29, Serial1/2
O     4.0.0.0/8 [110/1783] via 12.1.1.2, 00:08:04, Serial1/2
O     5.0.0.0/8 [110/1783] via 12.1.1.2, 00:08:04, Serial1/2
O     6.0.0.0/8 [110/1564] via 12.1.1.2, 00:10:29, Serial1/2
      10.0.0.0/24 is subnetted, 1 subnets
O        10.3.3.0 [110/783] via 12.1.1.2, 00:10:29, Serial1/2
      36.0.0.0/24 is subnetted, 1 subnets
O        36.1.1.0 [110/1563] via 12.1.1.2, 00:10:29, Serial1/2
      100.0.0.0/24 is subnetted, 1 subnets
O        100.1.1.0 [110/782] via 12.1.1.2, 00:31:56, Serial1/2
O     200.1.1.0/24 [110/1782] via 12.1.1.2, 00:10:29, Serial1/2
```

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 2 – OSPF Broadcast Networks

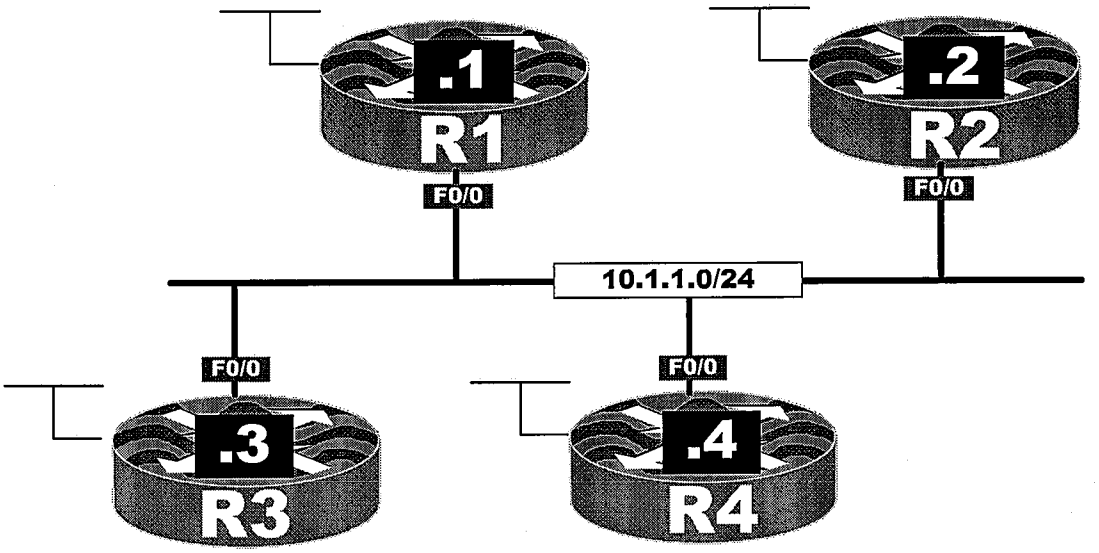

Task 1:

Configure the routers and the switches based on the above diagram. DO NOT configure OSPF.

On SW1:

```
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)#No shu
```

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config)#int F0/0
R1(config-if)#ip addr 10.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```

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

R2(config)#Int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0
```

On R3:

```
R3(config)#int F0/0
R3(config-if)#ip addr 10.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Int lo0
R3(config-if)#ip addr 3.3.3.3 255.0.0.0
```

On R4:

```
R4(config)#int F0/0
R4(config-if)#ip addr 10.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int lo0
R4(config-if)#ip addr 4.4.4.4 255.0.0.0
```

To verify 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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```


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

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

Task 2

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:

```
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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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


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
```


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

```
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 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

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 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 3

Lab Setup:

To copy and paste the initial configurations, go to "http://www.micronics.nl"→ "OSPF" →"Lab-2-Task-2".

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

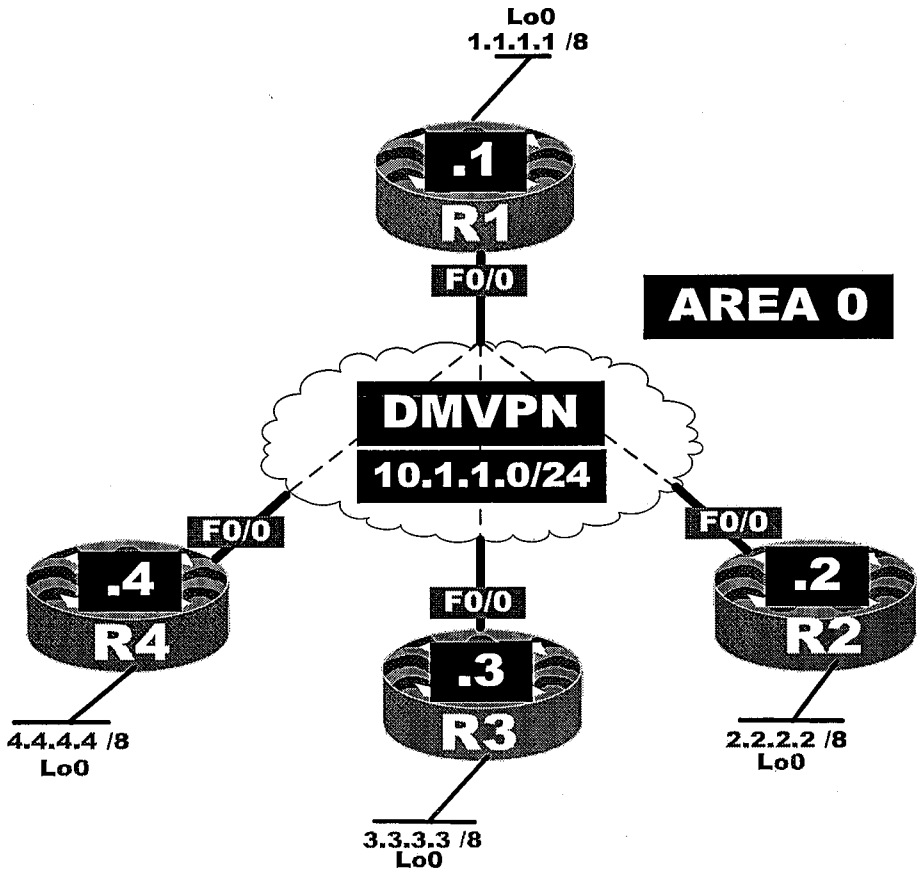

Configure OSPF on the tunnel and Loopback 0 interfaces of all routers based on the following policy:

- 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:

© 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 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
```

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#
```


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

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
 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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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
```

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
```


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

```
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:

```
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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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.
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:
```


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

```
.....
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
```

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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
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:
```


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

```
   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

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 3 – OSPF Non-Broadcast Networks

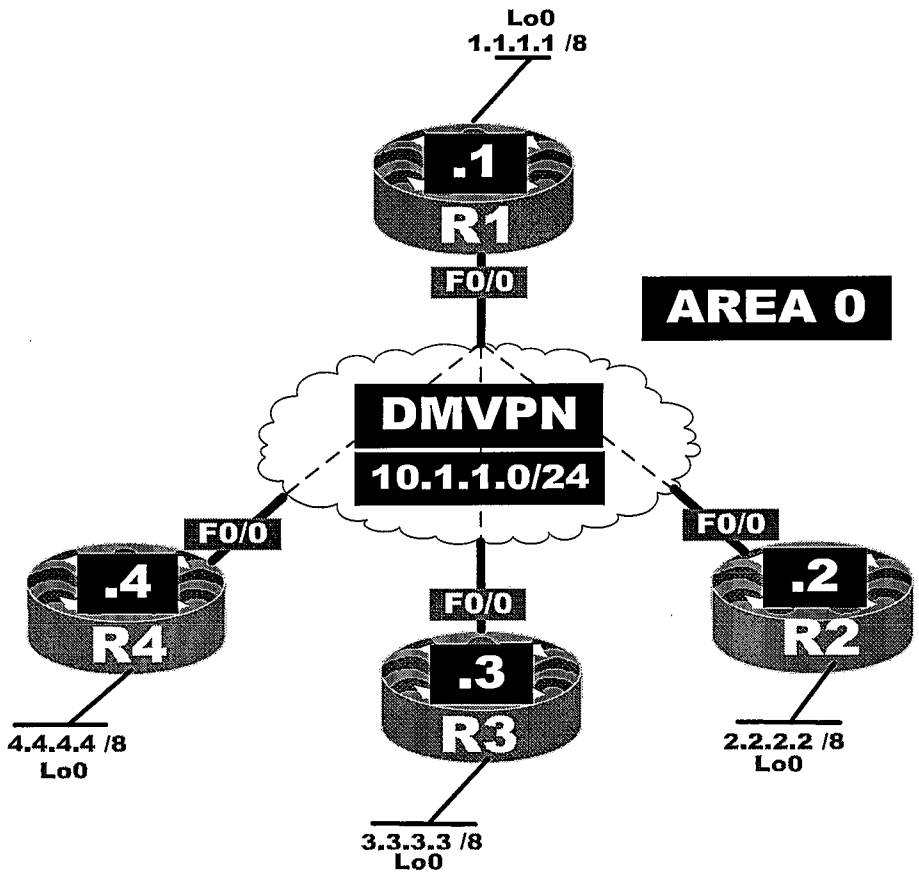

Lab Setup:

To copy and paste the initial configurations, go to "http://www.micronics.nl"→ "OSPF" →"Lab-3".

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

      1.0.0.0/32 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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
   Type: static, Flags:
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
   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
10.1.1.4/32 via 10.1.1.4
   Tunnel1234 created 00:00:21, never expire
```


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

```
   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

      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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      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:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

R4#Ping 2.2.2.2
```


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

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 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 4 – OSPF Point-to-Point Networks

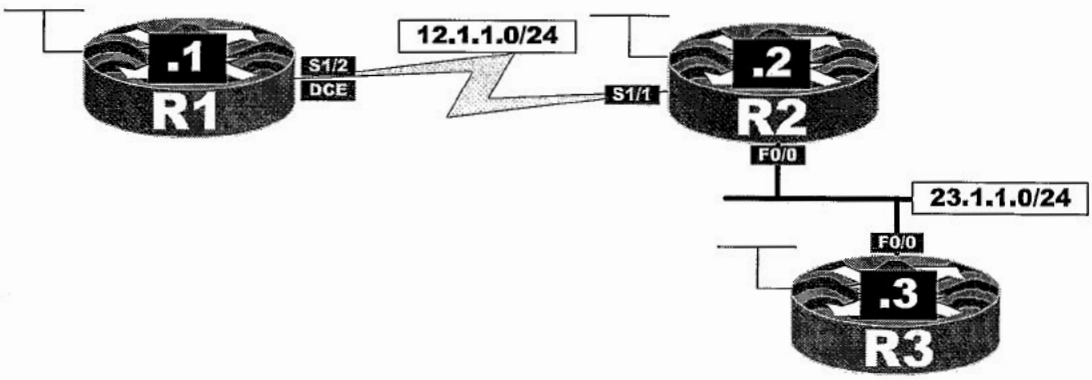

Lab Setup:

To copy and paste the initial configurations, go to "http://www.micronics.nl"→ "OSPF" →"Lab-4".

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
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
```


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

```
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

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

  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
```


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

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 5 – OSPF "Point-to-Multipoint" & "Point-to-Multipoint Non-Broadcast" Networks

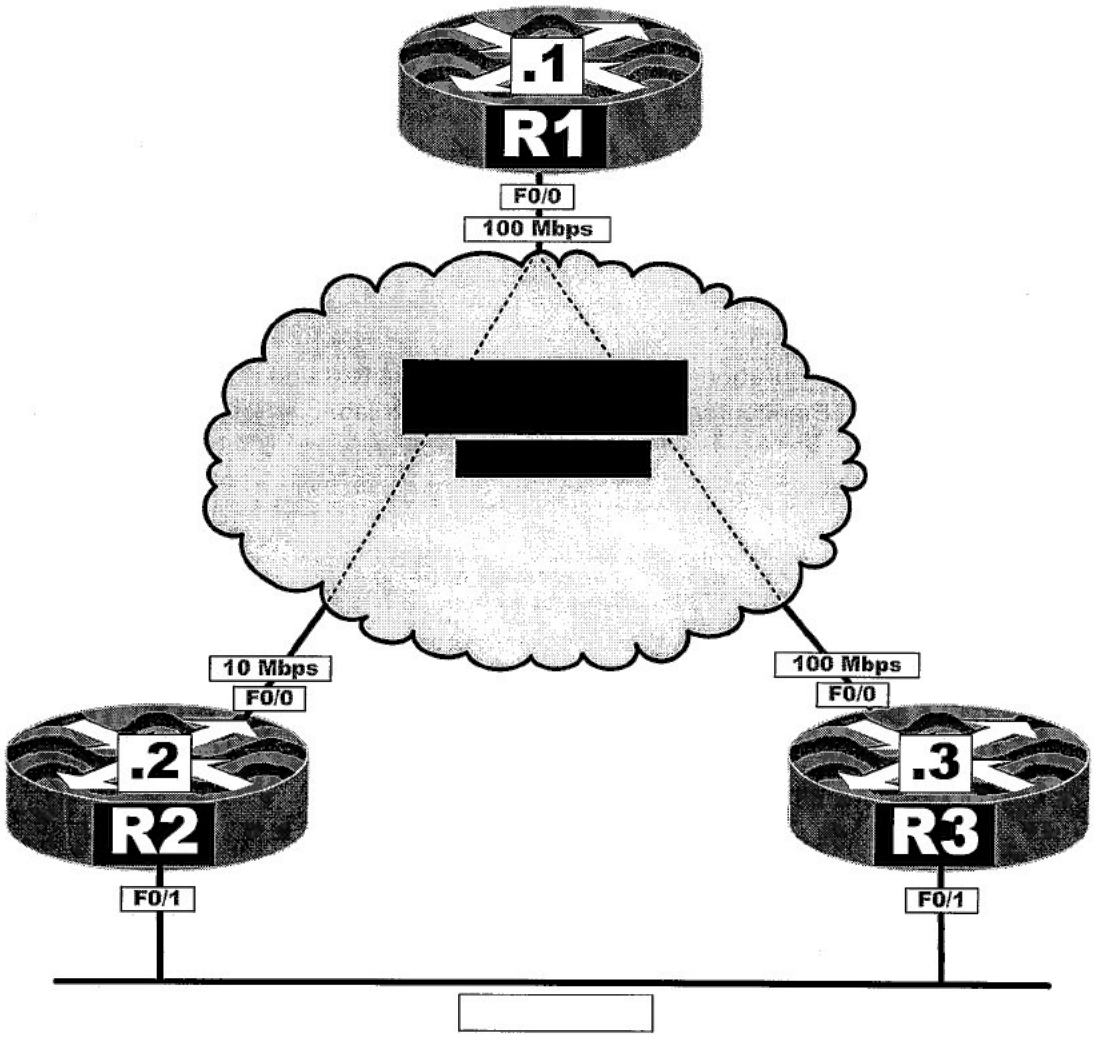

Lab Setup:

To copy and paste the initial configurations, go to "http://www.micronics.nl"→ "OSPF" →"Lab-5".

© 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
 ip ospf network point-to-multipoint
```


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

```
 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:

1. **The routers use 224.0.0.5**

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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.

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

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

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
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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

- **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

Lab 6 - OSPF Cost

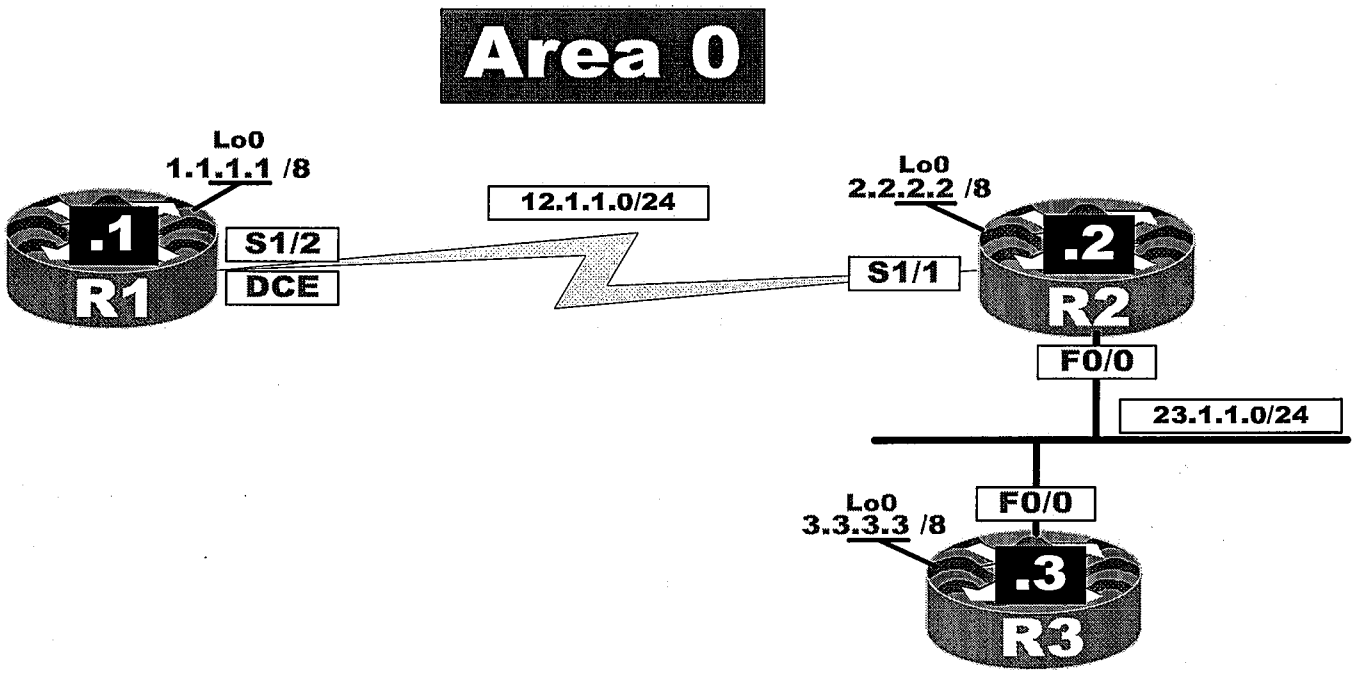

Task 1

Configure the routers and the switches based on the above diagram. DO NOT configure OSPF.

On SW1:

```
SW1(config)#Int range F0/2-3
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc vlan 23
SW1(config-if-range)#No shu
```

On R1:

```
R1(config)#int S1/2
R1(config-if)#clock rate 64000
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1(config-if)#No shu

R1(config-if)#int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0
```

On R2:

```
R2(config)#int s1/1
R2(config-if)#ip addr 12.1.1.2 255.255.255.0
R2(config-if)#No shu

R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.0.0.0

R2(config)#int f0/0
R2(config-if)#ip addr 23.1.1.2 255.255.255.0
R2(config-if)#No shu
```

On R3:

```
R3(config-if)#int lo0
R3(config-if)#ip addr 3.3.3.3 255.0.0.0

R3(config)#int f0/0
R3(config-if)#ip addr 10.1.23.3 255.255.255.0
R3(config-if)#No shu
```

To verify the configuration:

On R2:

```
R2#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.23.3, 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 all three routers in OSPF area 0 and advertise their directly connected networks in this area. Ensure that all routers have NLRI to every advertised network. The loopback interfaces are advertised with their correct mask. Configure router ids of 0.0.0.1, 0.0.0.2, and 0.0.0.3 for R1, R2 and R3 respectively.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Netw 0.0.0.0 0.0.0.0 are 0

R1(config-router)#int lo0
R1(config-if)#ip ospf net point-to-point
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#netw 0.0.0.0 0.0.0.0 are 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 0.0.0.0 0.0.0.0 are 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, 00:00:33, Serial1/2
O     3.0.0.0/8 [110/783] via 12.1.1.2, 00:00:15, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O        23.1.1.0 [110/782] via 12.1.1.2, 00:00:56, Serial1/2
```

Task 3

Configure R1 such that it advertises a cost of 20 for its loopback 0 interface.

You should check the cost of network 1.0.0.0 /8 which is advertised to R2 by R1.

On R2:

```
R2#Show ip route 1.0.0.0 | Inc metric

  Known via "ospf 1", distance 110, metric 782, type intra area
      Route metric is 782, traffic share count is 1
```

NOTE: From R2's perspective the cost of R1's loopback 0 interface is 782; this is the result of adding the cost of the serial interface plus the cost of the loopback 0 interface. Let's see the cost of R2's serial 1/1 interface and R1's Loopback 0 interface:

On R2:

```
R2#Show ip ospf int s1/1 | Inc Cost

  Process ID 1, Router ID 0.0.0.2, Network Type POINT_TO_POINT, Cost: 781
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

Why is the cost of this serial interface 781? Let's see:

```
R2#Show int s1/1 | Inc BW

  MTU 1500 bytes, BW 128 Kbit/sec, DLY 20000 usec,
```

The cost in OSPF is equal to 100,000,000 (The reference)/128 (The software Bandwidth of the interface:

100,000,000/128000 = 781.5
The fractions are dropped, so the cost is 781. Now, let's see the cost of the R1's Loopback 0 interface.

On R1:

```
R1#Show ip ospf int lo0 | Inc Cost
```


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

```
  Process ID 1, Router ID 0.0.0.1, Network Type POINT_TO_POINT, Cost: 1
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

Let's calculate the cost of this loopback interface:

```
R1#Show int lo0 | Inc BW

  MTU 1514 bytes, BW 8000000 Kbit/sec, DLY 5000 usec,
```

100,000,000/8,000,000 = 1
Remember that we can not use fractions. Therefore, 1 + 781 = 782.

Enter the following to change the cost of the lo0 interface on R1:

On R1

```
R1(config-if)#int lo0
R1(config-if)#ip ospf cost 20
```

To verify the configuration:

On R2:

```
R2#Show ip route 1.0.0.0 | Inc metric

  Known via "ospf 1", distance 110, metric 801, type intra area
      Route metric is 801, traffic share count is 1
```

NOTE: From R2's perspective the cost is 801, which is the sum of 781 (The cost of the serial interface) plus 20.

Task 4

In the future a gigabit interface will be installed on one of the routers in this routing domain. Ensure that the routers in this routing domain can differentiate the cost of a gigabit versus Fast Ethernet interface, and ensure the cost of the other interfaces are adjusted proportionally.

On All Routers:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Rx(config-if)#router ospf 1
Rx(config-router)#auto-cost reference-bandwidth 1000

Rx#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

By default, OSPF calculates the cost of an interface by dividing the reference (100,000,000 bps) by the bandwidth (Software Bandwidth) of that interface.

The default value is NOT recommended because sub-optimal routing can occur when your network has interface/s with a bandwidth greater than 100 Mbps, because OSPF cannot differentiate between 100 Mbps interface and an interface with a bandwidth that is greater than 100 Mbps.

The "IP OSPF COST" command enables you to change the OSPF cost of a given interface, but a better way to accomplish this is to change the default reference value. This can be accomplished by modifying the "Auto-cost reference-bandwidth"command. If you are planning to use this command, remember that this command should be configured on every router in the OSPF routing domain.

To verify the configuration:

On R1:

```
R1#Show ip ospf inter Lo0 | Inc Cost

  Process ID 1, Router ID 0.0.0.1, Network Type POINT_TO_POINT, Cost: 20
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name

R2#Show ip ospf int br

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               2.2.2.2/8          1     P2P   0/0
Se1/1        1     0               12.1.1.2/24        7812  P2P   1/1
Fa0/0        1     0               23.1.1.2/24        10    BDR   1/1
```

Task 5

Remove the "**Auto-cost reference-bandwidth**" command configured in the previous task.

On All Routers:

```
Rx(config-if)#router ospf 1
```


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

```
Rx(config-router)#NO auto-cost reference-bandwidth 1000

Rx#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

Task 6

Configure the routers such that the Fast Ethernet interface of these routers will have a cost of 70. The other interfaces should have their cost calculated proportionally.

The equation used by OSPF is as follows:
Reference/Bandwidth = Cost

Rearranging the formula, we get the following:
Reference = Cost * Bandwidth; which means 70 * 100,000,000 = 7,000,000,000

Since the configured reference value is in units of Mbps (1,000,000), and by default it's set to 100 which means 100,000,000 bps, now we have to divide the result by 1000,000 to get the actual reference value, which is 7000.

Therefore, the cost reference should be changed to 7000

On All Routers:

```
Rx(config-router)#router ospf 1
Rx(config-router)#auto-cost reference-bandwidth 7000
```

To verify the configuration:

On R3:

```
R3#Show ip ospf int F0/0 | Inc Cost

  Process ID 1, Router ID 0.0.0.3, Network Type BROADCAST, Cost: 70

R3#Show ip ospf int br

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               3.3.3.3/8          1     P2P   0/0
Fa0/0        1     0               23.1.1.3/24        70    DR    1/1
```


© 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 7 – OSPF Authentication

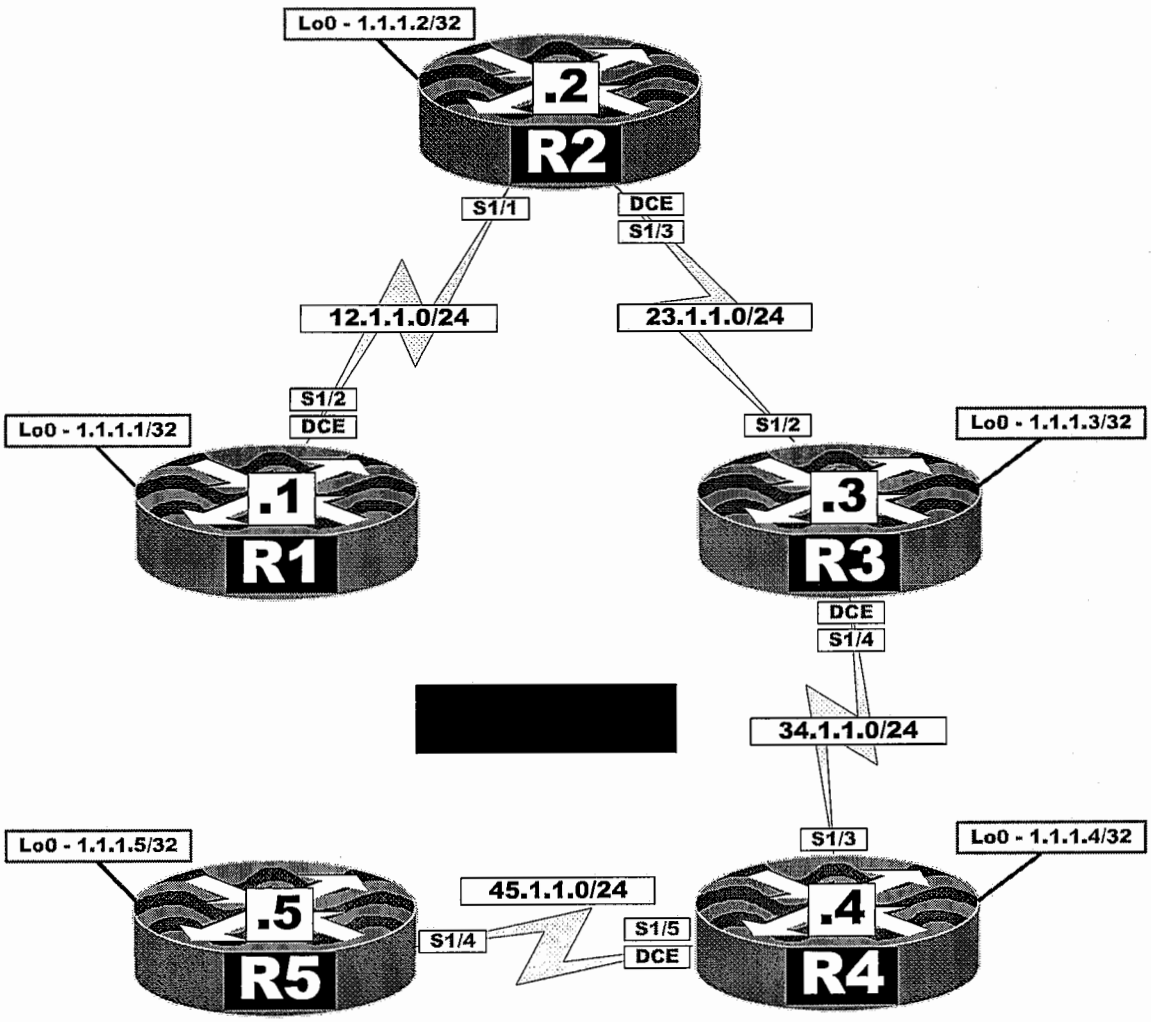

Task 1

Configure the routers based on the above diagram. DO NOT configure OSPF.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R1:

```
R1(config)#Int S1/2
R1(config-if)#clock rate 64000
R1(config-if)#IP address 12.1.1.1 255.255.255.0
R1(config-if)#No shut

R1(config)#Int Lo0
R1(config-if)#Ip addr 1.1.1.1 255.255.255.255
```

On R2:

```
R2(config)#Int S1/1
R2(config-if)#IP address 12.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int S1/3
R2(config-if)#clock rate 64000
R2(config-if)#IP address 23.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int Lo0
R2(config-if)#IP address 1.1.1.2 255.255.255.255
```

On R3:

```
R3(config)#Int S1/2
R3(config-subif)#IP address 23.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Int S1/4
R3(config-if)#clock rate 64000
R3(config-if)#IP address 34.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config-if)#Int Lo0
R3(config-if)#Ip addres 1.1.1.3 255.255.255.255
```

On R4:

```
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 S1/5
```


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

```
R4(config-if)#clock rate 64000
R4(config-if)#IP address 45.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int Lo0
R4(config-if)#IP address 1.1.1.4 255.255.255.255
```

On R5:

```
R5(config)#Int S1/4
R5(config-if)#IP address 45.1.1.5 255.255.255.0
R5(config-if)#No shut

R5(config)#Int Lo0
R5(config-if)#IP address 1.1.1.5 255.255.255.255
```

To verify the configuration:

On R2:

```
R2#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

On R4:

```
R4#Ping 34.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R4#Ping 45.1.1.5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 45.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/51/52 ms
```

Task 2

Configure the directly connected interfaces on all routers in area 0. The router-id of the routers in this area should NOT be based on any IP addressing.

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 0
R1(config-router)#netw 12.1.1.1 0.0.0.0 are 0
```

On R2:

```
R2(config-if)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#netw 1.1.1.2 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
```

On R3:

```
R3(config-if)#router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#netw 1.1.1.3 0.0.0.0 area 0
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:

```
R4(config-if)#router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#netw 1.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 34.1.1.4 0.0.0.0 area 0
R4(config-router)#netw 45.1.1.4 0.0.0.0 area 0
```


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

On R5:

```
R5(config-if)#router ospf 1
R5(config-router)#router-id 0.0.0.5
R5(config-router)#netw 45.1.1.5 0.0.0.0 area 0
R5(config-router)#netw 1.1.1.5 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 5 subnets
O        1.1.1.2 [110/782] via 12.1.1.2, 00:01:52, Serial1/2
O        1.1.1.3 [110/1563] via 12.1.1.2, 00:01:19, Serial1/2
O        1.1.1.4 [110/2344] via 12.1.1.2, 00:01:03, Serial1/2
O        1.1.1.5 [110/3125] via 12.1.1.2, 00:00:39, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/1562] via 12.1.1.2, 00:01:42, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/2343] via 12.1.1.2, 00:01:19, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/3124] via 12.1.1.2, 00:00:53, Serial1/2
```

On R3:

```
R3#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 [110/1563] via 23.1.1.2, 00:02:01, Serial1/2
O        1.1.1.2 [110/782] via 23.1.1.2, 00:02:01, Serial1/2
O        1.1.1.4 [110/782] via 34.1.1.4, 00:01:39, Serial1/4
O        1.1.1.5 [110/1563] via 34.1.1.4, 00:01:16, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/1562] via 23.1.1.2, 00:02:01, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/1562] via 34.1.1.4, 00:01:29, Serial1/4
```

On R5:

```
R5#Show ip route ospf | Inc 45.1.1.4
Gateway of last resort is not set
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
      1.0.0.0/32 is subnetted, 5 subnets
O        1.1.1.1 [110/3125] via 45.1.1.4, 00:01:42, Serial1/4
O        1.1.1.2 [110/2344] via 45.1.1.4, 00:01:42, Serial1/4
O        1.1.1.3 [110/1563] via 45.1.1.4, 00:01:42, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:01:42, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/3124] via 45.1.1.4, 00:01:42, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/2343] via 45.1.1.4, 00:01:42, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 45.1.1.4, 00:01:42, Serial1/4
```

Task 3

Configure plain text authentication on all the Serial links connecting the routers in this area. You MUST use a router configuration command as part of the solution to this task. Use "**Cisco**" as the password for this authentication.

OSPF supports two types of authentication, plain text (64 bit password) and MD5 (Which consists of a key ID and 128 bit password). In OSPF, authentication must be enabled and then applied.

In OSPF, enabling authentication can be configured in two different ways; one way to enable OSPF authentication is to configure it in the router configuration mode, in which case authentication is enabled globally on all OSPF enabled interfaces in the specified area. The second choice is to enable authentication directly on the interface for which authentication is required.

Since this task states that a router configuration mode must be used, OSPF authentication is enabled in the router configuration mode:

To understand OSPF's authentication, let's enable "Debug IP ospf packet":

On R1:

```
R1#Debug ip ospf packet
OSPF packet debugging is on
```

You should see the following debug messages:

```
OSPF-1 PAK  : rcv. v:2 t:1 l:48 rid:0.0.0.2 aid:0.0.0.0 chk:EC97 aut:0
auk: from Serial1/2
```


©https://t.me/learningnets

The output of the above debug message states the following:

- V:2 – OSPF Version 2
- T:1 – TTL of these messages are set to 1
- l:48 – The length of these messages are 48 Bytes
- rid:0.0.0.2 – This is the router-id of R2, the sending router
- aid:0.0.0.0 – This is the area id
- aut:0 – This means that there is no authentication
- auk: - No authentication key is defined
- from Serial1/2 – The packet is received through the local router's S1/2 interface

```
R1(config)#router ospf 1
R1(config-router)#area 0 authentication

R1(config-router)#int S1/2
R1(config-subif)#ip ospf authentication-key Cisco
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 0 authentication

R2(config-router)#int S1/1
R2(config-subif)#ip ospf authentication-key Cisco
```

On R1:

You should see that the output of the OSPF debug packets have their authentication type set to 1, this means clear text authentication; we will see MD5 authentication type later in this lab.

```
OSPF-1 PAK  : rcv. v:2 t:1 l:48 rid:0.0.0.2 aid:0.0.0.0 chk:EC96 aut:1
auk: from Serial1/2
```

Let's continue with R2's configuration:

On R2:

```
R2(config-if)#int S1/3
R2(config-if)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To turn off the debugs:

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

R2#Show ip ospf interface S1/1 | Inc auth

  Simple password authentication enabled
```

Note the output of the above "Show" command verifies that a simple password authentication is enabled and applied to this interface.

```
R2#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.1           0   FULL/  -        00:00:34    12.1.1.1        Serial1/1

R2#Show ip route ospf | Inc O
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:06:32, Serial1/1
```

Let's configure R3 and R4:

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 0 authentication

R3(config)#int S1/2
R3(config-if)#ip ospf authentication-key Cisco

R3(config)#int S1/4
R3(config-if)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R3:

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

      1.0.0.0/32 is subnetted, 3 subnets
O        1.1.1.1 [110/1563] via 23.1.1.2, 00:00:29, Serial1/2
```


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

```
O        1.1.1.2 [110/782] via 23.1.1.2, 00:00:29, Serial1/2
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/1562] via 23.1.1.2, 00:00:29, Serial1/2
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#area 0 authentication

R4(config)#int S1/3
R4(config-if)#ip ospf authentication-key Cisco

R4(config-if)#int S1/5
R4(config-if)#ip ospf authentication-key Cisco
```

To verify the configuration:

On R4:

You should NOT see 1.1.1.5/32 prefix in R4's routing table, if you still see this prefix in R4's routing table, you may have to wait for the adjacency to R5 to go down before entering the following show command:

```
R4#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 4 subnets
O        1.1.1.1 [110/2344] via 34.1.1.3, 00:00:48, Serial1/3
O        1.1.1.2 [110/1563] via 34.1.1.3, 00:00:48, Serial1/3
O        1.1.1.3 [110/782] via 34.1.1.3, 00:00:48, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/2343] via 34.1.1.3, 00:00:48, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/1562] via 34.1.1.3, 00:00:48, Serial1/3
```

Let's configure R5:

On R5:

```
R5(config)#Router ospf 1
R5(config-router)#area 0 authentication

R5(config-router)#int S1/4
R5(config-if)#ip ospf authentication-key Cisco
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify 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 [110/3125] via 45.1.1.4, 00:00:30, Serial1/4
O        1.1.1.2 [110/2344] via 45.1.1.4, 00:00:30, Serial1/4
O        1.1.1.3 [110/1563] via 45.1.1.4, 00:00:30, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:00:30, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/3124] via 45.1.1.4, 00:00:30, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/2343] via 45.1.1.4, 00:00:30, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 45.1.1.4, 00:00:30, Serial1/4
```

Task 4

Remove the authentication configuration from the previous task and ensure that every router sees every route advertised in area 0.

On All Routers:

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication
```

On R1:

```
R1(config)#int S1/2
R1(config-if)#No ip ospf authentication-key Cisco
```

On R2:

```
R2(config)#int S1/1
R2(config-if)#No ip ospf authentication-key Cisco

R2(config-if)#int S1/3
```


https://t.me/learningnets
© 2015 Narbik Kocharians. All rights reserved

```
R2(config-if)#No ip ospf authentication-key Cisco
```

On R3:

```
R3(config-router)#int S1/2
R3(config-if)#No ip ospf authentication-key Cisco

R3(config-if)#int S1/4
R3(config-if)#No ip ospf authentication-key Cisco
```

On R4:

```
R4(config)#int S1/3
R4(config-if)#No ip ospf authentication-key Cisco

R4(config)#int S1/5
R4(config-if)#No ip ospf authentication-key Cisco
```

On R5:

```
R5(config)#int S1/4
R5(config-if)#No ip ospf authentication-key Cisco
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | Inc O
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 5 subnets
O        1.1.1.2 [110/782] via 12.1.1.2, 00:17:46, Serial1/2
O        1.1.1.3 [110/1563] via 12.1.1.2, 00:09:36, Serial1/2
O        1.1.1.4 [110/2344] via 12.1.1.2, 00:07:31, Serial1/2
O        1.1.1.5 [110/3125] via 12.1.1.2, 00:05:36, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/1562] via 12.1.1.2, 00:17:46, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/2343] via 12.1.1.2, 00:09:36, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/3124] via 12.1.1.2, 00:07:31, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 5

Configure **MD5** authentication on all the Serial links in this area. You should use a router configuration command as part of the solution to this task. Use "**Cisco**" as the password for this authentication.

The following command enables MD5 authentication on the routers using the router configuration mode:

On All Routers:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

On R1:

```
R1(config)#int S1/2
R1(config-if)#ip ospf message-digest-key 1 MD5 Cisco
```

On R2:

```
R2(config)#int S1/1
R2(config-if)#ip ospf message-digest-key 1 MD5 Cisco
```

Let's see the Debug output and verify the authentication type and key:

On R1:

```
R1#Debug ip ospf packet
OSPF packet debugging is on
```

You should see the following debug output on your console:

```
OSPF-1 PAK  : rcv. v:2 t:1 l:48 rid:0.0.0.2 aid:0.0.0.0 chk:0 aut:2
keyid:1 seq:0x536538E9 from Serial1/2
```

You can clearly see the "aut: 2", this is identifying the authentication type which is set to 2, meaning that it's MD5 authentication, and the "keyid: 1" which means that the key value used in the configuration is 1.

On R2:

```
R2(config-if)#int S1/3
R2(config-if)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

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

Before we verify the configuration, let's disable the debug on R1

On R1:

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

On R2:

```
R2#Show ip ospf interface S0/0.21 | B Message

  Message digest authentication enabled
    Youngest key id is 1
```

NOTE: The output of the above show command reveals that MD5 authentication is enabled and applied and the key id is set to 1.

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

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:25:46, Serial1/1
```

On R3:

```
R3(config)#int S1/2
R3(config-if)#ip ospf message-digest-key 1 MD5 Cisco

R3(config)#int S1/4
R3(config-if)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

On R3:

```
R3#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 [110/1563] via 23.1.1.2, 00:00:11, Serial1/2
O        1.1.1.2 [110/782] via 23.1.1.2, 00:00:11, Serial1/2
O        1.1.1.4 [110/782] via 34.1.1.4, 00:16:51, Serial1/4
O        1.1.1.5 [110/1563] via 34.1.1.4, 00:14:46, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O        12.1.1.0 [110/1562] via 23.1.1.2, 00:00:11, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/1562] via 34.1.1.4, 00:16:51, Serial1/4
```

On R4:

```
R4(config)#int S1/3
R4(config-if)#ip ospf message-digest-key 1 MD5 Cisco

R4(config)#int S1/5
R4(config-if)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify the configuration:

On R4:

```
R4#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 [110/2344] via 34.1.1.3, 00:00:11, Serial1/3
O        1.1.1.2 [110/1563] via 34.1.1.3, 00:00:11, Serial1/3
O        1.1.1.3 [110/782] via 34.1.1.3, 00:00:11, Serial1/3
O        1.1.1.5 [110/782] via 45.1.1.5, 00:16:12, Serial1/5
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/2343] via 34.1.1.3, 00:00:11, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/1562] via 34.1.1.3, 00:00:11, Serial1/3
```

On R5:

```
R5(config)#int S1/4
R5(config-subif)#ip ospf message-digest-key 1 MD5 Cisco
```

To verify 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 [110/3125] via 45.1.1.4, 00:00:07, Serial1/4
O        1.1.1.2 [110/2344] via 45.1.1.4, 00:00:07, Serial1/4
O        1.1.1.3 [110/1563] via 45.1.1.4, 00:00:07, Serial1/4
```


https://t.me/learningnets
© 2015 Narbik Kocharians All rights reserved

```
O        1.1.1.4 [110/782] via 45.1.1.4, 00:00:07, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/3124] via 45.1.1.4, 00:00:07, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/2343] via 45.1.1.4, 00:00:07, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 45.1.1.4, 00:00:07, Serial1/4
```

Task 6

Remove the authentication configuration from the previous task and ensure that every router sees every route advertised in area 0.

On All Routers:

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication message-digest
```

On R1:

```
R1(config)#int S1/2
R1(config-if)#No ip ospf message-digest-key 1 MD5 Cisco
```

On R2:

```
R2(config)#int S1/1
R2(config-if)#No ip ospf message-digest-key 1 MD5 Cisco

R2(config)#int S1/3
R2(config-if)#No ip ospf message-digest-key 1 MD5 Cisco
```

On R3:

```
R3(config)#int S1/2
R3(config-if)#No ip ospf message-digest-key 1 MD5 Cisco

R3(config)#int S1/4
R3(config-if)#No ip ospf message-digest-key 1 MD5 Cisco
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#int S1/3
R4(config-if)#No ip ospf message-digest-key 1 MD5 Cisco

R4(config)#int S1/5
R4(config-if)#No ip ospf message-digest-key 1 MD5 Cisco
```

On R5:

```
R5(config)#int S1/4
R5(config-if)#No ip ospf message-digest-key 1 MD5 Cisco
```

To verify 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 [110/3125] via 45.1.1.4, 00:04:50, Serial1/4
O        1.1.1.2 [110/2344] via 45.1.1.4, 00:04:50, Serial1/4
O        1.1.1.3 [110/1563] via 45.1.1.4, 00:04:50, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:04:50, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/3124] via 45.1.1.4, 00:04:50, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/2343] via 45.1.1.4, 00:04:50, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 45.1.1.4, 00:04:50, Serial1/4
```

Task 7

Configure MD5 authentication on the Serial link connecting R1 to R2, you should use a router configuration command as part of the solution to this task. The password should be "ccie".

On Both Routers:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```


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

On R1:

```
R1(config)#int S1/2
R1(config-if)#ip ospf message-digest-key 1 MD5 ccie
```

On R2:

```
R2(config)#int S1/1
R2(config-if)#ip ospf message-digest-key 1 MD5 ccie
```

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

And then, you should see the following console message stating that the local router no longer has an adjacency with R3 with a router id of 0.0.0.3.

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

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       1.1.1.1 [110/782] via 12.1.1.1, 00:36:55, Serial1/1
```

Note because authentication is enabled in the router configuration mode, it is applied to every interface that is running in area 0, therefore, every router in area 0 MUST have the "Area 0 authentication message-digest" command configured. Since R3 does NOT have authentication enabled, these routers will drop their adjacency.

To verify the configuration:

On R2:

```
R2#Sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.1           0   FULL/  -        00:00:39    12.1.1.1        Serial1/1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

There are two solutions to fix this problem:

1. Enable authentication on R3, but if authentication is enabled on R3 under router ospf, then R4 will drop the adjacency, therefore, if router configuration mode MUST be used as part of the solution (Based on the task), authentication needs to be enabled on R3, R4 and R5.

2. Disable authentication under the S1/3 interface. If authentication is disabled on the interface facing R3, then R3, R4 and R5 won't need to have authentication enabled.

Let's configure the above solutions and verify:

Solution 1:

On R3, R4 and R5:

```
Rx(config)#Router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

You should see the following console message on R3:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Serial1/2 from LOADING to FULL,
Loading Done
```

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, 5 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:43:45, Serial1/1
O        1.1.1.3 [110/782] via 23.1.1.3, 00:00:57, Serial1/3
O        1.1.1.4 [110/1563] via 23.1.1.3, 00:00:57, Serial1/3
O        1.1.1.5 [110/2344] via 23.1.1.3, 00:00:57, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 23.1.1.3, 00:00:57, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/2343] via 23.1.1.3, 00:00:57, Serial1/3
```

Solution 2:

On R3, R4 and R5:

https://t.me/learningnets
© 2019 Narbik Kocharians. All rights reserved

```
Rx(config)#Router ospf 1
Rx(config-router)#No area 0 authentication message-digest
```

You should see the following console message after the dead interval expires:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 from FULL to DOWN,
Neighbor Down: Dead timer expired
```

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        1.1.1.1 [110/782] via 12.1.1.1, 00:45:32, Serial1/1
```

In this solution, authentication is disabled on R2's interface facing R3 using the "IP OSPF authentication null" interface configuration command, meaning that there is no need to have authentication downstream to S1/3 interface of R2. Therefore, R3, R4 and R5 DON'T need to have authentication enabled.

On R2:

```
R2(config)#Int S1/3
R2(config-if)#IP Ospf authentication null
```

You should see the following console message on R2:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | Inc O
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 5 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:47:16, Serial1/1
O        1.1.1.3 [110/782] via 23.1.1.3, 00:00:20, Serial1/3
O        1.1.1.4 [110/1563] via 23.1.1.3, 00:00:20, Serial1/3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O        1.1.1.5 [110/2344] via 23.1.1.3, 00:00:20, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 23.1.1.3, 00:00:20, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/2343] via 23.1.1.3, 00:00:20, Serial1/3
```

Task 8

Re-configure the authentication password on R1 and R2 to be "**CCIE12**" without interrupting the links operation.

To see the current configuration:

On R1:

```
R1#Show ip ospf int S1/2 | B Mess

  Message digest authentication enabled
    Youngest key id is 1

R1#Show run int S1/2 | Inc ip ospf

ip ospf message-digest-key 1 md5 ccie
```

On R2:

```
R2#Sh ip ospf int s1/1 | B Mess

  Message digest authentication enabled
    Youngest key id is 1

R2#Show run int s1/1 | Inc ip ospf

ip ospf message-digest-key 1 md5 ccie

R2#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 [110/782] via 12.1.1.1, 00:50:19, Serial1/1
O        1.1.1.3 [110/782] via 23.1.1.3, 00:03:23, Serial1/3
```


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

```
O        1.1.1.4 [110/1563] via 23.1.1.3, 00:03:23, Serial1/3
O        1.1.1.5 [110/2344] via 23.1.1.3, 00:03:23, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 23.1.1.3, 00:03:23, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/2343] via 23.1.1.3, 00:03:23, Serial1/3
```

In order to change the password without any interruption to the link, the second key is entered with the required password.

On R1:

```
R1(config)#int S1/2
R1(config-if)# ip ospf message-digest-key 2 md5 CCIE12
```

To verify the configuration:

On R1:

```
R1#Show run int S1/2 | Inc ip ospf

ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIE12

R1#Show ip ospf inter S1/2 | B Message

  Message digest authentication enabled
    Youngest key id is 2
    Rollover in progress, 1 neighbor(s) using the old key(s):
      key id 1
```

Even though the second key (key 2) is only configured on R1, R1 and R2 are still authenticating based on the first key (key 1), this is revealed in the second line of the above show command.

But the R1 knows that the second key is configured (The second line in the above display) and it knows that the rollover is in progress (The third line), but the other end (R2) has not been configured yet.

On R2:

```
R2(config-subif)#int S1/1
R2(config-if)# ip ospf message-digest-key 2 md5 CCIE12
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Sh ip ospf inter S0/0.21 | b Message

Message digest authentication enabled
    Youngest key id is 2
```

NOTE: Once R2 is configured, both routers (R1 and R2) will switchover and use the second key for their authentication.

On R1:

```
R1#Show ip ospf interface S1/2 | b Message

  Message digest authentication enabled
    Youngest key id is 2
```

Once R1 and R2's key rollover is completed and both routers display the same youngest key without the "rollover in progress" message, we can safely remove the prior key, in this case key id 1. Remember that the newest key is NOT determined based on the numerically higher value.

On R1:

```
R1#Show run int S1/2 | Inc ip ospf

 ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIE12

R1(config)#int S1/2
R1(config-subif)#No ip ospf message-digest-key 1 md5 ccie
```

On R2:

```
R2#Show run int S1/1 | Inc ip ospf

 ip ospf message-digest-key 1 md5 ccie
 ip ospf message-digest-key 2 md5 CCIE12

R2(config)#int S1/1
R2(config-subif)#No ip ospf message-digest-key 1 md5 ccie
```


https://t.me/learningnets
© 2018 Narbik Kocharians. All Rights Reserved

Task 9

Configure MD5 authentication on the link that connects R4 to R5 using "**Cisco45**" as the password. You should NOT use a router configuration mode to accomplish this task.

On R5:

```
R5(config)#Int S1/4
R5(config-if)#IP Ospf authentication message-digest
R5(config-if)#IP Ospf message-digest-key 1 md5 Cisco45
```

On R4:

```
R4(config)#Int S1/5
R4(config-if)#IP Ospf authentication message-digest
R4(config-if)#IP Ospf message-digest-key 1 md5 Cisco45
```

NOTE: The authentication is enabled and applied directly under the interface for which authentication was required. When authentication is enabled directly under a given interface, it enables authentication on that given interface ONLY, therefore, ONLY the neighbor/s through that interface should have authentication enabled. This is called per-interface authentication.

To verify 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 [110/3125] via 45.1.1.4, 00:00:09, Serial1/4
O        1.1.1.2 [110/2344] via 45.1.1.4, 00:00:09, Serial1/4
O        1.1.1.3 [110/1563] via 45.1.1.4, 00:00:09, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:00:09, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/3124] via 45.1.1.4, 00:00:09, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/2343] via 45.1.1.4, 00:00:09, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/1562] via 45.1.1.4, 00:00:09, Serial1/4
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 10

Re-configure OSPF Areas based on the following chart and remove all the authentications configured on the routers, these routers should see all the routes advertised in this routing domain.

Router	Interface	Area
R1	S1/2 Loopback 0	0 0
R2	S1/1 S1/3 Loopback 0	0 1 1
R3	S1/2 S1/4 Loopback 0	1 2 2
R4	S1/3 S1/5 Loopback 0	2 3 3
R5	S1/4 Loopback 0	3 3

On All Routers:

```
Rx(config)#No Router ospf 1
```

On R1:

```
R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
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)#Int S1/2
R1(config-subif)#No ip ospf message-digest-key 2 md5 CCIE12
```

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2
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 1
R2(config-router)#Netw 1.1.1.2 0.0.0.0 area 1
```


©https://t.me/learningnets

```
R2(config)#Int S1/1
R2(config-subif)#No ip ospf message-digest-key 2 md5 CCIE12

R2(config)#Int S1/3
R2(config-subif)#No ip ospf authentication null
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#Netw 1.1.1.3 0.0.0.0 area 2
R3(config-router)#Netw 34.1.1.3 0.0.0.0 area 2
R3(config-router)#Netw 23.1.1.3 0.0.0.0 area 1
```

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#Netw 1.1.1.4 0.0.0.0 area 3
R4(config-router)#Netw 45.1.1.4 0.0.0.0 area 3
R4(config-router)#Netw 34.1.1.4 0.0.0.0 area 2

R4(config)#Int S1/5
R4(config-if)#No ip ospf message-digest-key 1 md5 Cisco45
R4(config-if)#No ip ospf authentication message-digest
```

On R5:

```
R5(config)#Router ospf 1
R5(config-router)#router-id 0.0.0.5
R5(config-router)#Netw 1.1.1.5 0.0.0.0 area 3
R5(config-router)#Netw 45.1.1.5 0.0.0.0 area 3

R5(config)#Int S1/4
R5(config-if)#No ip ospf message-digest-key 1 md5 Cisco45
R5(config-if)#No ip ospf authentication message-digest
```

In order for these routers to see all the routes advertised in this routing domain, we MUST configure virtual-links because NOT all areas have connectivity to area 0.

Area 1 has a connection to area 0, but areas 2 and 3 do not. Let's begin with area 2:

On R2:

```
R2(config)#Router ospf 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Area 1 virtual-link 0.0.0.3
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Area 1 virtual-link 0.0.0.2
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on OSPF_VL0 from LOADING to FULL,
Loading Done
```

To connect area 3 to area 0:

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Area 2 virtual-link 0.0.0.4
```

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#Area 2 virtual-link 0.0.0.3
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on OSPF_VL2 from LOADING to FULL,
Loading Done
```

To verify 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 IA     1.1.1.1 [110/3125] via 45.1.1.4, 00:00:40, Serial1/4
O IA     1.1.1.2 [110/2344] via 45.1.1.4, 00:00:40, Serial1/4
O IA     1.1.1.3 [110/1563] via 45.1.1.4, 00:00:45, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:03:17, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/3124] via 45.1.1.4, 00:00:40, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
```


©https://t.me/learningnets

```
O IA     23.1.1.0 [110/2343] via 45.1.1.4, 00:00:40, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 45.1.1.4, 00:00:45, Serial1/4
```

Task 11

Configure MD5 authentication on the link between R1 and R2 in area 0, the password for this authentication should be set to "**Micronics**", you should use router configuration mode to accomplish this task.

On R1 and R2:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

On R1:

```
R1(config)#Int S1/2
R1(config-subif)#ip ospf message-digest-key 1 md5 Micronics
```

On R2:

```
R2(config)#int S1/1
R2(config-subif)#ip ospf message-digest-key 1 md5 Micronics
```

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, 5 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:07:10, Serial1/1
O IA     1.1.1.3 [110/782] via 23.1.1.3, 00:02:49, Serial1/3
O IA     1.1.1.4 [110/1563] via 23.1.1.3, 00:02:02, Serial1/3
O IA     1.1.1.5 [110/2344] via 23.1.1.3, 00:02:02, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 23.1.1.3, 00:02:49, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/2343] via 23.1.1.3, 00:02:02, Serial1/3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Why do we see all the routes?
Let's shutdown the lo0 interface of R2, and then "No shut" the interface, and you should see the following console message within 40 seconds:

```
R2(config)#int lo0
R2(config-if)#Shut
```

Wait for the link to go down before entering the following command:

```
R2(config-if)#No shut

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on OSPF_VL0 from FULL to DOWN,
Neighbor Down: Dead timer expired

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

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:24:18, Serial1/1
```

The reason we had to "Shut" and then "No Shut" an advertised route is because virtual-links are demand circuits, and when a link is demand circuit, OSPF suppresses the OSPF Hellos and Refresh messages. Demand circuits are typically configured on SVCs such as ISDN, so when OSPF is enabled on a demand circuit, OSPF hello messages will keep that link up indefinetly, to handle this issue the "IP ospf demand-circuit" command is configured, with this command configured, OSPF will form an adjacency and then the link goes down but the OSPF adjacency stays up, and since hellos and refresh messages are suppressed, the link can stay down.

Question:

When does this link ever come up?

When there is a topology change, enabling authentication is NOT a topology change, and this is the reason we had to "Shutdown" the interface and then "No Shut" the interface, this triggers a topology change. When a topology change is detected, the link comes up, and when the link comes up and you have enabled authentication on one end of the link and not the other, the virtual-link goes down and stays down until authentication is enabled on the other end of the link.

NOTE: R2 does not have any other prefix in its routing table; this is because authentication is enabled directly under the router configuration mode of R1 and R2, when authentication is enabled in the router configuration mode, it is enabled on all links in the configured area, in this case area 0, and since virtual-links are always in area 0, authentication must also be enabled on those links. There are two ways to fix this problem:

1. Enable authentication on R3, and R4 in their router configuration mode. Remember R5 does not have a virtual-link configured.

https://t.me/learningnets

2. Enable authentication directly on the virtual-links that are configured on R2, R3 and R4.

3. Disable authentication on R2's virtual-link.

Let's implement the first solution:

On R3 and R4:

```
Rx(config)#router ospf 1
Rx(config-router)#area 0 authentication message-digest
```

To verify 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 IA     1.1.1.1 [110/3125] via 45.1.1.4, 00:00:17, Serial1/4
O IA     1.1.1.2 [110/2344] via 45.1.1.4, 00:08:25, Serial1/4
O IA     1.1.1.3 [110/1563] via 45.1.1.4, 00:08:30, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:11:02, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/3124] via 45.1.1.4, 00:00:17, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/2343] via 45.1.1.4, 00:08:25, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 45.1.1.4, 00:08:30, Serial1/4
```

On R2:

```
R2#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 [110/782] via 12.1.1.1, 00:14:03, Serial1/1
O IA     1.1.1.3 [110/782] via 23.1.1.3, 00:01:07, Serial1/3
O IA     1.1.1.4 [110/1563] via 23.1.1.3, 00:01:07, Serial1/3
O IA     1.1.1.5 [110/2344] via 23.1.1.3, 00:01:07, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 23.1.1.3, 00:01:07, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/2343] via 23.1.1.3, 00:01:07, Serial1/3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Remember....when authentication is enabled in router configuration mode, authentication is enabled on all links/interfaces in the spcified area, since virtual-links are always in area 0, authentication will be enabled on all virtual-links.

Let's implement the second solution:

Before the second option is configured and verified, the configuration from the previous solution should be removed:

On R3 and R4:

```
Rx(config)#router ospf 1
Rx(config-router)#No area 0 authentication message-digest

Rx#Clear ip ospf process
Reset ALL OSPF processes? [no]: y
```

To verify the configuration:

On R2:

```
R2#Sh ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
O        1.1.1.1 [110/782] via 12.1.1.1, 00:16:26, Serial1/1
```

To enable authentication on the virtual-links:

```
R2(config)#router ospf 1
R2(config-router)#Area 1 virtual-link 0.0.0.3 authen mess
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#Area 1 virtual-link 0.0.0.2 authentication message-digest
R3(config-router)#Area 2 virtual-link 0.0.0.4 authentication message-digest
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on OSPF_VL0 from LOADING to FULL,
Loading Done
```

On R4:

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

```
R4(config)#Router ospf 1
R4(config-router)#Area 2 virtual-link 0.0.0.3 authentication message-digest
```

To verify 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 IA     1.1.1.1 [110/3125] via 45.1.1.4, 00:01:22, Serial1/4
O IA     1.1.1.2 [110/2344] via 45.1.1.4, 00:04:19, Serial1/4
O IA     1.1.1.3 [110/1563] via 45.1.1.4, 00:04:24, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:04:24, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/3124] via 45.1.1.4, 00:01:22, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/2343] via 45.1.1.4, 00:04:09, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 45.1.1.4, 00:04:24, Serial1/4
```

Let's implement the third solution:

Before the third option is configured and verified, the configuration from the previous solution is removed:

On R2:

```
R2(config)#router ospf 1
R2(config-router)#No Area 1 virtual-link 0.0.0.3
R2(config-router)#Area 1 virtual-link 0.0.0.3
```

On R3:

```
R3(config)#Router ospf 1
R3(config-router)#No area 1 virtual-link 0.0.0.2
R3(config-router)#No area 2 virtual-link 0.0.0.4

R3(config-router)#Area 1 virtual-link 0.0.0.2
R3(config-router)#Area 2 virtual-link 0.0.0.4
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#Router ospf 1
R4(config-router)#No area 2 virtual-link 0.0.0.3
R4(config-router)#Area 2 virtual-link 0.0.0.3
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 2 subnets
O IA     1.1.1.2 [110/782] via 12.1.1.2, 00:15:54, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/1562] via 12.1.1.2, 00:23:52, Serial1/2
```

To implement the third solution:

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#Area 1 virtual-link 0.0.0.3 authentication null
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on OSPF_VL2 from LOADING to FULL,
Loading Done
```

On R2:

```
R2#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 [110/782] via 12.1.1.1, 00:25:40, Serial1/1
O IA     1.1.1.3 [110/782] via 23.1.1.3, 00:00:48, Serial1/3
O IA     1.1.1.4 [110/1563] via 23.1.1.3, 00:00:48, Serial1/3
O IA     1.1.1.5 [110/2344] via 23.1.1.3, 00:00:48, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 23.1.1.3, 00:00:48, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/2343] via 23.1.1.3, 00:00:48, Serial1/3
```

On R5:

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

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 5 subnets
O IA     1.1.1.1 [110/3125] via 45.1.1.4, 00:01:10, Serial1/4
O IA     1.1.1.2 [110/2344] via 45.1.1.4, 00:04:02, Serial1/4
O IA     1.1.1.3 [110/1563] via 45.1.1.4, 00:04:07, Serial1/4
O        1.1.1.4 [110/782] via 45.1.1.4, 00:10:34, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/3124] via 45.1.1.4, 00:01:10, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/2343] via 45.1.1.4, 00:04:02, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1562] via 45.1.1.4, 00:04:07, Serial1/4
```

Task 12

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
OSPF Summarization

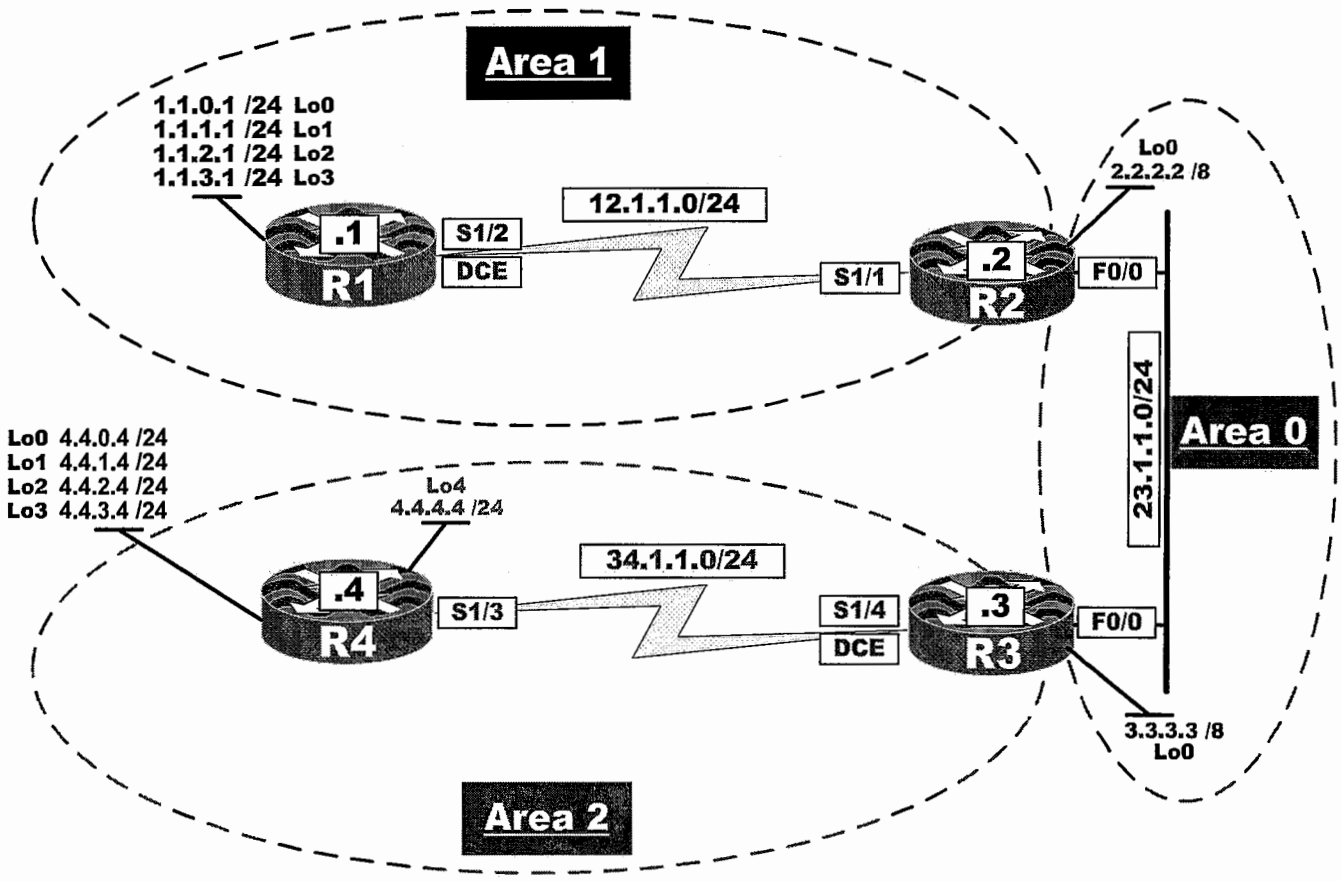

Task 1

Configure the Routers and the Switches based on the above diagram. DO NOT configure OSPF.

On SW1:

```
SW1(config)#Int Range f0/2-3
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc vlan 23
SW1(config-if-range)#No shut
```


https://t.me/learningnets

On R1:

```
R1(config)#Int S1/2
R1(config-if)#IP address 12.1.1.1 255.255.255.0
R1(config-if)#clock rate 64000
R1(config-if)#No shut

R1(config)#Int Lo0
R1(config-if)#IP address 1.1.0.1 255.255.255.0

R1(config)#Int Lo1
R1(config-if)#IP address 1.1.1.1 255.255.255.0

R1(config)#Int Lo2
R1(config-if)#IP address 1.1.2.1 255.255.255.0

R1(config)#Int Lo3
R1(config-if)#IP address 1.1.3.1 255.255.255.0
```

On R2:

```
R2(config)#Int S1/1
R2(config-if)#IP address 12.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int F0/0
R2(config-if)#IP address 23.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int Lo0
R2(config-if)#IP address 2.2.2.2 255.0.0.0
```

On R3:

```
R3(config)#Int S1/4
R3(config-subif)#IP address 34.1.1.3 255.255.255.0
R3(config-subif)#clock rate 64000
R3(config-if)#No shut

R3(config)#Int F0/0
R3(config-if)#IP address 23.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config)#Int Lo0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-if)#IP address 3.3.3.3 255.0.0.0
```

On R4:

```
R4(config)#Int S1/3
R4(config-subif)#IP address 34.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int Lo0
R4(config-if)#IP addr 4.4.0.4 255.255.255.0

R4(config)#Int Lo1
R4(config-if)#IP addr 4.4.1.4 255.255.255.0

R4(config)#Int Lo2
R4(config-if)#IP addr 4.4.2.4 255.255.255.0

R4(config)#Int Lo3
R4(config-if)#IP addr 4.4.3.4 255.255.255.0

R4(config)#Int Lo4
R4(config-if)#IP addr 4.4.4.4 255.255.255.0
```

To verify the configuration:

On R1:

```
R2#Ping 12.1.1.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

On R2:

```
R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

On R3:

©https://t.me/learningnets

```
R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
```

Task 2

Configure R4 based on the following policy:

- R4 should run OSPF area 2 on its Loopback 4 and its S1/3 interfaces.
- R4 should redistribute Lo0, Lo1, Lo2 and Lo3 in this routing domain.
- R4's router id should be configured to be 0.0.0.4

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 2
R4(config-router)#Network 4.4.4.4 0.0.0.0 area 2

R4(config)#Route-map tst
R4(config-route-map)#Match interface Lo0 Lo1 Lo2 Lo3

R4(config)#Router ospf 1
R4(config-router)#Redistribute connected route-map tst subnets
```

Task 3

Configure R3 based on the following policy:

- R3 should run OSPF Area 0 on its Loopback0 and its F0/0 interfaces, and run OSPF Area 2 on its S1/4 interface.

- R3's router id should be configured to be 0.0.0.3

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 23.1.1.3 0.0.0.0 area 0
R3(config-router)#Network 34.1.1.3 0.0.0.0 area 2
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Serial1/4 from LOADING to FULL,
Loading Done
```

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, 4 subnets
O E2     4.4.0.0 [110/20] via 34.1.1.4, 00:00:51, Serial1/4
O E2     4.4.1.0 [110/20] via 34.1.1.4, 00:00:51, Serial1/4
O E2     4.4.2.0 [110/20] via 34.1.1.4, 00:00:51, Serial1/4
O E2     4.4.3.0 [110/20] via 34.1.1.4, 00:00:51, Serial1/4
```

Task 4

Configure R2 based on the following policy:

- R2 should advertise its Loopback0 and its F0/0 interfaces in area 0 and its S1/1 interface in area 1.
- R2's router id should be configured to be 0.0.0.2

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 23.1.1.2 0.0.0.0 area 0
```


©https://t.me/learningnets

```
R2(config-router)#Network 12.1.1.2 0.0.0.0 area 1
```

To verify the configuration:

On R2:

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

      4.0.0.0/24 is subnetted, 4 subnets
O E2     4.4.0.0 [110/20] via 23.1.1.3, 00:00:39, FastEthernet0/0
O E2     4.4.1.0 [110/20] via 23.1.1.3, 00:00:39, FastEthernet0/0
O E2     4.4.2.0 [110/20] via 23.1.1.3, 00:00:39, FastEthernet0/0
O E2     4.4.3.0 [110/20] via 23.1.1.3, 00:00:39, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 23.1.1.3, 00:00:39, FastEthernet0/0
```

Task 5

Configure R1 based on the following policy:

- R1 should run OSPF Area 1 on all of its directly connected interfaces, you should use minimum number of "Network" statements to accomplish this task.
- R1's router id should be configured to be 0.0.0.1.

On R1:

```
R1(config)#router ospf 1
R1(config-router)#router-id 1.1.1.1
R1(config-router)#netw 0.0.0.0 0.0.0.0 area 1
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Serial1/2 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show ip route ospf | Inc O
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/782] via 12.1.1.2, 00:02:14, Serial1/2
      4.0.0.0/24 is subnetted, 4 subnets
O E2     4.4.0.0 [110/20] via 12.1.1.2, 00:02:14, Serial1/2
O E2     4.4.1.0 [110/20] via 12.1.1.2, 00:02:14, Serial1/2
O E2     4.4.2.0 [110/20] via 12.1.1.2, 00:02:14, Serial1/2
O E2     4.4.3.0 [110/20] via 12.1.1.2, 00:02:14, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/782] via 12.1.1.2, 00:02:14, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1563] via 12.1.1.2, 00:02:14, Serial1/2
```

On R2:

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

      1.0.0.0/32 is subnetted, 4 subnets
O        1.1.0.1 [110/782] via 12.1.1.1, 00:02:51, Serial1/1
O        1.1.1.1 [110/782] via 12.1.1.1, 00:02:51, Serial1/1
O        1.1.2.1 [110/782] via 12.1.1.1, 00:02:51, Serial1/1
O        1.1.3.1 [110/782] via 12.1.1.1, 00:02:51, Serial1/1
      4.0.0.0/24 is subnetted, 4 subnets
O E2     4.4.0.0 [110/20] via 23.1.1.3, 00:05:19, FastEthernet0/0
O E2     4.4.1.0 [110/20] via 23.1.1.3, 00:05:19, FastEthernet0/0
O E2     4.4.2.0 [110/20] via 23.1.1.3, 00:05:19, FastEthernet0/0
O E2     4.4.3.0 [110/20] via 23.1.1.3, 00:05:19, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 23.1.1.3, 00:05:19, FastEthernet0/0
```

Task 6

Configure the appropriate router in area 2 to summarize all the external (E2) routes.

In OSPF, summarization can be configured on two types of routers: An ABR/s and/or an ASBR/s. The OSPF intra-area routes can only be summarized on the ABR/s, whereas, the external (redistributed) routes can be summarized on the router that originates the external routes.

©https://t.me/learningnets

If the external routes are conveyed in LSA type 7s (N1 or N2 routes), they MUST be summarized on the ASBR, because ONLY the ASBR can originate the "N" routes; the ABR of a given NSSA area receives the "N" routes and originates an "E" route for every "N" route that it receives and injects them into OSPF area 0; in this case the ABR is the router that originates the "E" route/s, therefore, the ABR can summarize the external routes (The "E" routes).

When summarizing internal routes on an ABR, the "Area xx range" command must be used, where xx is the area id where the specific routes were originated.

Summarization of external routes can be accomplished by using the "summary-address" command. In this case since R4 is the router that originates type-5s, summarization can ONLY be performed on R4.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#summary-address 4.4.0.0 255.255.252.0
```

NOTE: Whenever summarization is performed on a given router, the IOS will auto-inject a summary route to Null0 to avoid forwarding loops.

To see the Null0 route for the summary:

On R4:

```
R4#Show ip route | Inc Null0

O        4.4.0.0/22 is a summary, 00:09:31, Null0
```

To verify the configuration:

On R1:

```
R1#Show ip route OSPF | B Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/782] via 12.1.1.2, 00:10:56, Serial1/2
      4.0.0.0/22 is subnetted, 1 subnets
O E2     4.4.0.0 [110/20] via 12.1.1.2, 00:02:14, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/782] via 12.1.1.2, 00:10:56, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1563] via 12.1.1.2, 00:10:56, Serial1/2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

NOTE: The external routes are summarized.

Task 7

Configure the Loopback interfaces of R1 such that they are advertised with their correct mask.

On R1:

```
R1(config)#Interface Range Lo0 - 3
R1(config-if-range)#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

      1.0.0.0/24 is subnetted, 4 subnets
O IA     1.1.0.0 [110/783] via 23.1.1.2, 00:00:17, FastEthernet0/0
O IA     1.1.1.0 [110/783] via 23.1.1.2, 00:00:17, FastEthernet0/0
O IA     1.1.2.0 [110/783] via 23.1.1.2, 00:00:17, FastEthernet0/0
O IA     1.1.3.0 [110/783] via 23.1.1.2, 00:00:17, FastEthernet0/0
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/2] via 23.1.1.2, 00:15:21, FastEthernet0/0
      4.0.0.0/22 is subnetted, 1 subnets
O E2     4.4.0.0 [110/20] via 34.1.1.4, 00:04:14, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/782] via 23.1.1.2, 00:15:15, FastEthernet0/0
```

Task 8

Configure the appropriate router in area 1 to summarize networks 1.1.0.0 /24, 1.1.1.0 /24, 1.1.2.0 /24 and 1.1.3.0 /24) into the OSPF routing domain.

The routes that are identified in this task are originated by R1, but they can only be summarized by the

©https://t.me/learningnets

ABR of area 1, in this topology the ABR for area 1 is R2.

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 range 1.1.0.0 255.255.252.0
```

NOTE: Whenever summarization is performed on a given router, the IOS will auto-inject a summary route to Null0 to avoid forwarding loops.

To see the Null0 route for the summary:

On R2:

```
R2#Show ip route ospf | Inc Null0

O        1.1.0.0/22 is a summary, 00:00:15, Null0
```

To verify the configuration:

On R3:

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

      1.0.0.0/22 is subnetted, 1 subnets
O IA     1.1.0.0 [110/783] via 23.1.1.2, 00:01:20, FastEthernet0/0
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/2] via 23.1.1.2, 00:25:40, FastEthernet0/0
      4.0.0.0/22 is subnetted, 1 subnets
O E2     4.4.0.0 [110/20] via 34.1.1.4, 00:14:33, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/782] via 23.1.1.2, 00:25:34, FastEthernet0/0
```

Task 9

The routers should NOT install a null 0 route in the routing table when they summarize internal or external routes. DO NOT negate any command/s on R4 to accomplish this task.

In OSPF, the discard route is created automatically whenever summarization is configured; there are two

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

types of summary routes: Internal and External.

When summarization is configured for intra-area routes, OSPF will auto-inject an internal discard route, and when summarization is configured for external OSPF route, OSPF process will create an external discard route. The discard routes are created to stop forwarding loops.

To remove the internal discard route:

On R2:

```
R2(config)#router ospf 1
R2(config-router)#No discard-route internal
```

To verify the configuration:

On R2:

```
R2#Show ip route ospf | Inc Null0
R2#
```

To remove the external discard route:

On R4:

```
R4#Show ip route ospf | Inc summary

O        4.4.0.0/22 is a summary, 00:06:13, Null0
```

NOTE: A route with an administrative distance of 255 will NEVER be injected into the routing table of the local router:

On R4:

```
R4(config)#router ospf 1
R4(config-router)#discard-route external 255
```

To Verify the configuration:

On R4:

```
R4#Show ip route ospf | Inc Null0
R4#
```


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

Task 10

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 9 – OSPF Filtering

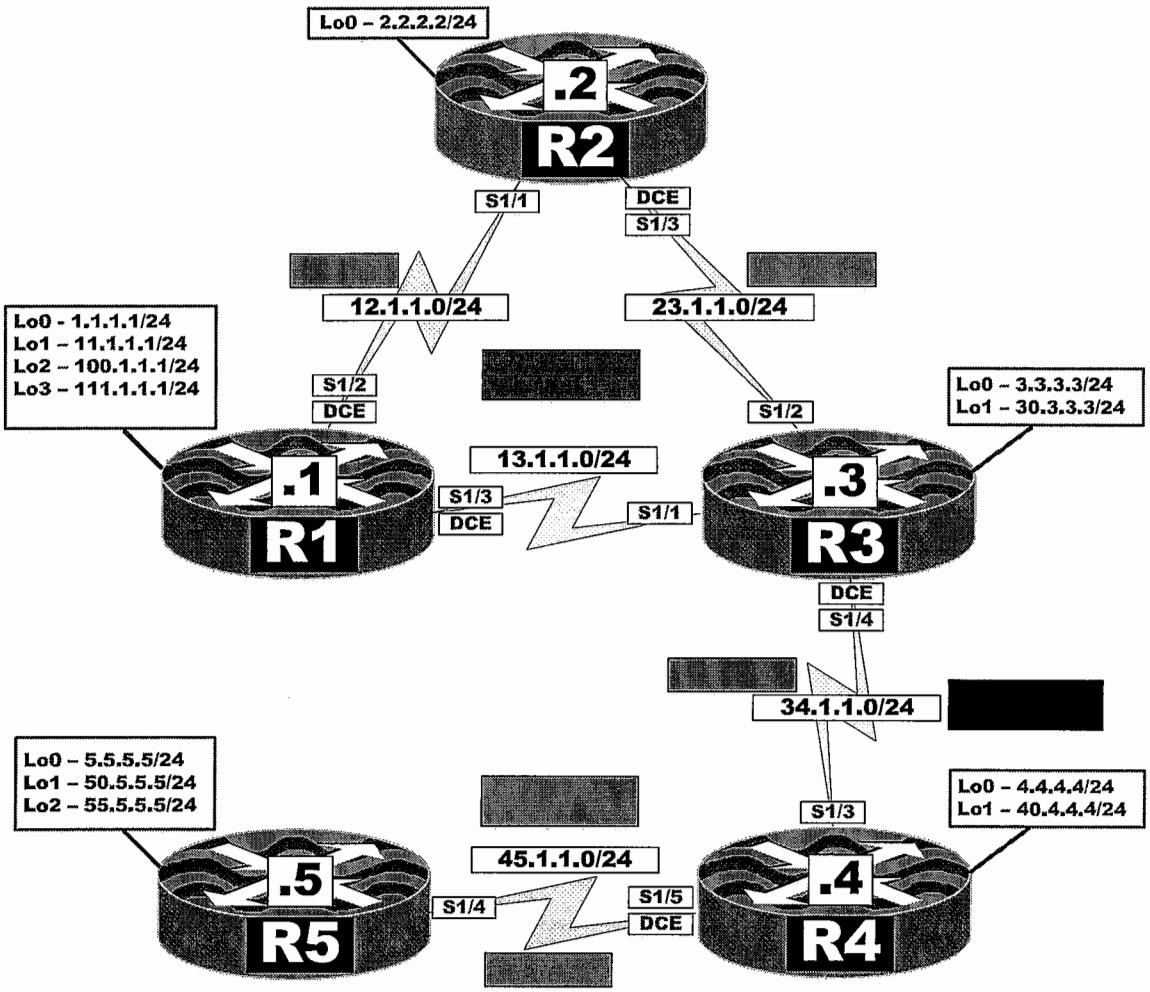

Task 1

Configure the Routers and the Switches based on the above diagram. DO NOT configure any routing protocol.

On R1:

©[illegible] https://t.me/learningnets

```
R1(config)#Int S1/2
R1(config-if)#IP address 12.1.1.1 255.255.255.0
R1(config-if)#Clock rate 64000
R1(config-if)#bandwidth 128
R1(config-if)#No shut

R1(config-if)#Int S1/3
R1(config-if)#IP address 13.1.1.1 255.255.255.0
R1(config-if)#Clock rate 64000
R1(config-if)#bandwidth 1544
R1(config-if)#No shut

R1(config-if)#Int Lo0
R1(config-if)#IP address 1.1.1.1 255.255.255.0

R1(config-if)#Int Lo1
R1(config-if)#IP address 11.1.1.1 255.255.255.0

R1(config-if)#Int Lo2
R1(config-if)#IP address 100.1.1.1 255.255.255.0

R1(config-if)#Int Lo3
R1(config-if)#IP address 111.1.1.1 255.255.255.0
```

On R2:

```
R2(config)#Int S1/1
R2(config-if)#IP address 12.1.1.2 255.255.255.0
R2(config-if)#bandwidth 128
R2(config-if)#No shut

R2(config)#Int S1/3
R2(config-if)#IP address 23.1.1.2 255.255.255.0
R2(config-if)#Clock rate 64000
R2(config-if)#bandwidth 1544
R2(config-if)#No shut

R2(config)#Int Lo0
R2(config-if)#IP address 2.2.2.2 255.255.255.0
```

On R3:

```
R3(config)#Int S1/2
R3(config-subif)#IP address 23.1.1.3 255.255.255.0
R3(config-if)#bandwidth 1544
R3(config-if)#No shut
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config)#Int S1/1
R3(config-if)#IP address 13.1.1.3 255.255.255.0
R3(config-if)#bandwidth 1544
R3(config-if)#No shut

R3(config)#Int S1/4
R3(config-if)#IP address 34.1.1.3 255.255.255.0
R3(config-if)#Clock rate 64000
R3(config-if)#bandwidth 1544
R3(config-if)#No shut

R3(config-if)#Int Lo0
R3(config-if)#IP address 3.3.3.3 255.255.255.0

R3(config-if)#Int Lo1
R3(config-if)#IP address 30.3.3.3 255.255.255.0
```

On R4:

```
R4(config)#Int S1/3
R4(config-if)#IP address 34.1.1.4 255.255.255.0
R4(config-if)#bandwidth 1544
R4(config-if)#No shut

R4(config)#Int S1/5
R4(config-if)#IP address 45.1.1.4 255.255.255.0
R4(config-if)#Clock rate 64000
R4(config-if)#bandwidth 1544
R4(config-if)#No shut

R4(config-if)#Int Lo0
R4(config-if)#IP addr 4.4.4.4 255.255.255.0

R4(config-if)#int Lo1
R4(config-if)#IP address 40.4.4.4 255.255.255.0
```

On R5:

```
R5(config)#Int S1/4
R5(config-if)#IP address 45.1.1.5 255.255.255.0
R5(config-if)#bandwidth 1544
R5(config-if)#No shut

R5(config-if)#Int Lo0
R5(config-if)#IP address 5.5.5.5 255.255.255.0
```


© [illegible] All rights reserved
https://t.me/learningnets

```
R5(config-if)#Int Lo1
R5(config-if)#IP address 50.5.5.5 255.255.255.0

R5(config-if)#Int Lo2
R5(config-if)#IP address 55.5.5.5 255.255.255.0
```

To verify the configuration:

On R2:

```
R2#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

On R3:

```
R3#Ping 13.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 13.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms

R3#Ping 34.1.1.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms
```

On R4:

```
R4#Ping 45.1.1.5

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 45.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/50/52 ms
```

Task 2

Configure R1 and R2's directly connected interfaces, and R3's Serial connections to R2 and R1 in area 1. Configure 0.0.0.1, 0.0.0.2, 0.0.0.3 to be the OSPF RID of R1, R2 and R3 respectively.

On R1:

```
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 1
R1(config-router)#Network 13.1.1.1 0.0.0.0 area 1
R1(config-router)#Netw 1.1.1.1 0.0.0.0 area 1
R1(config-router)#Netw 11.1.1.1 0.0.0.0 area 1
R1(config-router)#Netw 100.1.1.1 0.0.0.0 area 1
R1(config-router)#Netw 111.1.1.1 0.0.0.0 area 1
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Netw 12.1.1.2 0.0.0.0 area 1
R2(config-router)#Netw 23.1.1.2 0.0.0.0 area 1
R2(config-router)#Netw 2.2.2.2 0.0.0.0 area 1
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

On R3:

```
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 1
R3(config-router)#Netw 13.1.1.3 0.0.0.0 area 1
```


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

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Serial1/2 from LOADING to FULL,
Loading Done

%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 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

      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/129] via 13.1.1.3, 00:00:49, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/128] via 13.1.1.3, 00:00:49, Serial1/3
```

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 [110/65] via 13.1.1.1, 00:01:23, Serial1/1
      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/65] via 23.1.1.2, 00:01:33, Serial1/2
      11.0.0.0/32 is subnetted, 1 subnets
O        11.1.1.1 [110/65] via 13.1.1.1, 00:01:23, Serial1/1
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/845] via 23.1.1.2, 00:01:33, Serial1/2
                  [110/845] via 13.1.1.1, 00:01:23, Serial1/1
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.1 [110/65] via 13.1.1.1, 00:01:23, Serial1/1
      111.0.0.0/32 is subnetted, 1 subnets
O        111.1.1.1 [110/65] via 13.1.1.1, 00:01:23, Serial1/1
```

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 [110/129] via 23.1.1.3, 00:02:04, Serial1/3
      11.0.0.0/32 is subnetted, 1 subnets
O        11.1.1.1 [110/129] via 23.1.1.3, 00:02:04, Serial1/3
      13.0.0.0/24 is subnetted, 1 subnets
O        13.1.1.0 [110/128] via 23.1.1.3, 00:02:04, Serial1/3
      100.0.0.0/32 is subnetted, 1 subnets
O        100.1.1.1 [110/129] via 23.1.1.3, 00:02:04, Serial1/3
      111.0.0.0/32 is subnetted, 1 subnets
O        111.1.1.1 [110/129] via 23.1.1.3, 00:02:04, Serial1/3
```

Task 3

Configure the Serial connection between R3 and R4 in Area 0. R4's OSPF RID should be set to 0.0.0.4.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#Netw 34.1.1.3 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)#Netw 34.1.1.4 0.0.0.0 area 0
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 from LOADING to FULL,
Loading Done
```

To verify the configuration:

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 IA     1.1.1.1 [110/129] via 34.1.1.3, 00:00:38, Serial1/3
```


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

```
      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/129] via 34.1.1.3, 00:00:38, Serial1/3
      11.0.0.0/32 is subnetted, 1 subnets
O IA     11.1.1.1 [110/129] via 34.1.1.3, 00:00:38, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/909] via 34.1.1.3, 00:00:38, Serial1/3
      13.0.0.0/24 is subnetted, 1 subnets
O IA     13.1.1.0 [110/128] via 34.1.1.3, 00:00:38, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/128] via 34.1.1.3, 00:00:38, Serial1/3
      100.0.0.0/32 is subnetted, 1 subnets
O IA     100.1.1.1 [110/129] via 34.1.1.3, 00:00:38, Serial1/3
      111.0.0.0/32 is subnetted, 1 subnets
O IA     111.1.1.1 [110/129] via 34.1.1.3, 00:00:38, Serial1/3
```

Task 4

Configure the Serial connection between R4 and R5 in Area 2. R5's OSPF RID should be set to 0.0.0.5.

On R4:

```
R4(config)#router ospf 1
R4(config-router)#Netw 45.1.1.4 0.0.0.0 area 2
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#router-id 0.0.0.5
R5(config-router)#Netw 45.1.1.5 0.0.0.0 area 2
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Serial1/4 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | B Gate
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/193] via 45.1.1.4, 00:00:43, Serial1/4
      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/193] via 45.1.1.4, 00:00:43, Serial1/4
      11.0.0.0/32 is subnetted, 1 subnets
O IA     11.1.1.1 [110/193] via 45.1.1.4, 00:00:43, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/973] via 45.1.1.4, 00:00:43, Serial1/4
      13.0.0.0/24 is subnetted, 1 subnets
O IA     13.1.1.0 [110/192] via 45.1.1.4, 00:00:43, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/192] via 45.1.1.4, 00:00:43, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/128] via 45.1.1.4, 00:00:43, Serial1/4
      100.0.0.0/32 is subnetted, 1 subnets
O IA     100.1.1.1 [110/193] via 45.1.1.4, 00:00:43, Serial1/4
      111.0.0.0/32 is subnetted, 1 subnets
O IA     111.1.1.1 [110/193] via 45.1.1.4, 00:00:43, Serial1/4
```

Task 5

Configure the loopback interfaces of R1 and R2 with their correct mask.

On R1:

```
R1(config)#Int range Lo0 - 3
R1(config-if)#IP Ospf network point-to-point
```

On R2:

```
R2(config)#Int Lo0
R2(config-if)#IP Ospf network point-to-point
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set
```


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

```
      1.0.0.0/24 is subnetted, 1 subnets
O IA     1.1.1.0 [110/193] via 45.1.1.4, 00:00:37, Serial1/4
      2.0.0.0/24 is subnetted, 1 subnets
O IA     2.2.2.0 [110/193] via 45.1.1.4, 00:00:20, Serial1/4
      11.0.0.0/24 is subnetted, 1 subnets
O IA     11.1.1.0 [110/193] via 45.1.1.4, 00:00:37, Serial1/4
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/973] via 45.1.1.4, 00:02:26, Serial1/4
      13.0.0.0/24 is subnetted, 1 subnets
O IA     13.1.1.0 [110/192] via 45.1.1.4, 00:02:26, Serial1/4
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/192] via 45.1.1.4, 00:02:26, Serial1/4
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/128] via 45.1.1.4, 00:02:26, Serial1/4
      100.0.0.0/24 is subnetted, 1 subnets
O IA     100.1.1.0 [110/193] via 45.1.1.4, 00:00:37, Serial1/4
      111.0.0.0/24 is subnetted, 1 subnets
O IA     111.1.1.0 [110/193] via 45.1.1.4, 00:00:37, Serial1/4
```

Task 6

Configure R2 to filter network 1.1.1.0 /24 from its routing table. Ensure that R3, R4 and R5 have full reachability to this network. You should use an "IP prefix-list" to accomplish this task.

In this case, the following solution only affects the router that it's configured on.

On R2:

```
R2(config)#ip prefix-list TST seq 5 deny 1.1.1.0/24
R2(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R2(config)#router ospf 1
R2(config-router)#distribute-list Prefix TST in
```

To verify the configuration:

On R2:

```
R2#Show ip route 1.1.1.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
% Network not in table
```

Note the "distribute-list in" router configuration mode command can be used when filtering any route on a given router, this command ONLY filters the prefix/es from the local router's routing table and NOT the database. The output of the following show command reveals that R2 still has the prefix in its database:

```
R2#Show ip ospf database router adv-router 0.0.0.1

            OSPF Router with ID (0.0.0.2) (Process ID 1)

                Router Link States (Area 1)

  LS age: 383
  Options: (No TOS-capability, DC)
  LS Type: Router Links
  Link State ID: 0.0.0.1
  Advertising Router: 0.0.0.1
  LS Seq Number: 80000009
  Checksum: 0xD11D
  Length: 120
  Number of Links: 8

    Link connected to: a Stub Network
     (Link ID) Network/subnet number: 1.1.1.0
     (Link Data) Network Mask: 255.255.255.0
      Number of MTID metrics: 0
       TOS 0 Metrics: 1
(The rest of the output is omitted for brevity)
```

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 = 28/28/32 ms
```

On R4:

```
R4#Ping 1.1.1.1

Type escape sequence to abort.
```


© 2013 Narbik Kocharians All rights reserved
https://t.me/learningnets

```
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 = 80/81/84 ms
```

On R5:

```
R5#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 = 128/130/132 ms
```

Task 7

Configure filtering on the appropriate router/s such that the existing and future routers in area 2 do not see network 11.1.1.0 /24 in their routing table and/or database.

The following method ONLY works for filtering Type-3 LSAs, filtering type-3 LSAs can only work when it's configured on the ABR/s; in this case it can only work on R4.

The first step is to configure a prefix-list to deny network 11.1.1.0/24 and permit everything else:

On R4:

```
R4(config)#ip prefix-list TST seq 5 deny 11.1.1.0/24
R4(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32
```

Once the prefix-list is configured, it can be applied to the area that it must be filtered from, in this case area 2. In the following configuration, the "Prefix-list" filter's network 11.1.1.0 /24 from getting into area 2.

```
R4(config)#router ospf 1
R4(config-router)#area 2 filter-list prefix TST in
```

To verify the configuration:

On R5:

```
R5#Sh ip route 11.1.1.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
% Network not in table
```

The above show command reveals that R5 does NOT have the route in its routing table, and the following command verifies that R5 does not have the prefix in its database.

```
R5#Show ip ospf database summary 11.1.1.0

            OSPF Router with ID (0.0.0.5) (Process ID 1)
```

On R4:

```
R4#Show ip route 11.1.1.0

Routing entry for 11.1.1.0/24
  Known via "ospf 1", distance 110, metric 129, type inter area
  Last update from 34.1.1.3 on Serial1/3, 00:21:56 ago
  Routing Descriptor Blocks:
  * 34.1.1.3, from 0.0.0.3, 00:21:56 ago, via Serial1/3
      Route metric is 129, traffic share count is 1
```

Note even though the output of the above show command reveals that network 11.1.1.0 /24 is in R4's routing table, the output of the following show command clearly shows that it's in the database of area 0 and NOT in the database that belongs to area 2.

```
R4#Show ip ospf database summary 11.1.1.0

            OSPF Router with ID (0.0.0.4) (Process ID 1)

                Summary Net Link States (Area 0)

  Routing Bit Set on this LSA in topology Base with MTID 0
  LS age: 1369
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 11.1.1.0 (summary Network Number)
  Advertising Router: 0.0.0.3
  LS Seq Number: 80000001
  Checksum: 0x5D8E
  Length: 28
  Network Mask: /24
        MTID: 0         Metric: 65
```

Let's check another prefix like 1.1.1.0/24 and see the difference, this prefix should be in the database of area 0 and area 2:

© https://t.me/learningnets served

```
R4#Sh ip ospf da summ 1.1.1.0

            OSPF Router with ID (0.0.0.4) (Process ID 1)

                Summary Net Link States (Area 0)

  Routing Bit Set on this LSA in topology Base with MTID 0
  LS age: 1406
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 1.1.1.0 (summary Network Number)
  Advertising Router: 0.0.0.3
  LS Seq Number: 80000001
  Checksum: 0xDF16
  Length: 28
  Network Mask: /24
        MTID: 0         Metric: 65


                Summary Net Link States (Area 2)

  LS age: 1405
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 1.1.1.0 (summary Network Number)
  Advertising Router: 0.0.0.4
  LS Seq Number: 80000001
  Checksum: 0x5C58
  Length: 28
  Network Mask: /24
        MTID: 0         Metric: 129
```

We can see that from R4's perspective network 1.1.1.0/24 shows up in the database of areas 0 and 2, and from area 0's perspective the cost is 65 and from area 2's perspective the cost os 129.

Task 8

Configure the appropriate router such that the routers in area 0 or area 2 do not see network 11.1.1.0 /24 in their routing table or Link state database. You should use the same solution as the one in the previous task, but it should be implemented in the OUT bound direction.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

The following method is used for filtering LSA type 3s. A prefix-list is configured to deny network 11.1.1.0 /24 on the ABR, in this case R3, but in the following case the "area filter-list" command is filtering network 11.1.1.0 /24 as it's advertised OUT of area 1. This will affect areas 0 and 2. To see the existing routing table of R4:

On R4:

```
R4#Show ip route 11.1.1.0

Routing entry for 11.1.1.0/24
  Known via "ospf 1", distance 110, metric 129, type inter area
  Last update from 34.1.1.3 on Serial1/3, 00:33:58 ago
  Routing Descriptor Blocks:
  * 34.1.1.3, from 0.0.0.3, 00:33:58 ago, via Serial1/3
      Route metric is 129, traffic share count is 1
```

On R3:

```
R3(config)#ip prefix-list TST seq 5 deny 11.1.1.0/24
R3(config)#ip prefix-list TST seq 10 permit 0.0.0.0/0 LE 32

R3(config)#router ospf 1
R3(config-router)#area 1 filter-list prefix TST OUT
```

To verify the configuration:

Note the output of the following commands show that prefix 11.1.1.0 /24 is no longer in the routing table of R4.

On R4:

```
R4#Show ip route 11.1.1.0

% Network not in table
```

Note the prefix is still in the routing table of R3, where the filtering is performed, but the prefix is in the routing table of R3 as a route from area 1 (An intra-area route) and NOT a prefix from area 0 (From area 0's perspective this prefix is an inter-area route).

On R3:

```
R3#Show ip route 11.1.1.0
```


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

```
Routing entry for 11.1.1.0/24
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 13.1.1.1 on Serial1/1, 00:37:55 ago
  Routing Descriptor Blocks:
  * 13.1.1.1, from 0.0.0.1, 00:37:55 ago, via Serial1/1
      Route metric is 65, traffic share count is 1
```

To prove this further:

On R3:

```
R3#Sh ip ospf database summary 11.1.1.0

      OSPF Router with ID (0.0.0.3) (Process ID 1)
```

Note the output of the above show command reveals that network 11.1.1.0 /24 is NOT in area 0, if it was in area 0, it would have been in the Link State database of this router as a summary LSA or LSA type 3, whereas, the following show command reveals that the prefix is in area 1 as a router LSA or LSA type 1.

```
R3#Show ip ospf database router | Inc Area 1|_11.1.1.0

                Router Link States (Area 1)
    (Link ID) Network/subnet number: 11.1.1.0
```

Task 9

Configure the appropriate router/s to filter network 111.1.1.0/24 such that the routers in area 0 or area 2 DO NOT have this network in their routing table or database. Use the minimum number of commands to accomplish this task.

The output of the following show command verifies that network 111.1.1.0 /24 is in the database of R3 belonging to Area 1 as a router LSA or LSA type 1:

On R3:

```
R3#Show ip ospf database router | Inc Area 1|_111.1.1.0

                Router Link States (Area 1)
    (Link ID) Network/subnet number: 111.1.1.0
```

The output of the following show command reveals that prefix 111.1.1.0 /24 is also in the database of R3 as

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

a summary LSA or LSA type 3, but from Area 0's perspective:

```
R3#Show ip ospf database summary 111.1.1.0

            OSPF Router with ID (0.0.0.3) (Process ID 1)

                Summary Net Link States (Area 0)

  LS age: 120
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 111.1.1.0 (summary Network Number)
  Advertising Router: 0.0.0.3
  LS Seq Number: 80000003
  Checksum: 0x4045
  Length: 28
  Network Mask: /24
        MTID: 0         Metric: 65
```

The following command reveals that prefix 111.1.1.0 /24 is in the routing table of R3 as an intra-area route.

```
R3#Show ip route ospf | inc 111.1.1.0

O        111.1.1.0 [110/65] via 13.1.1.1, 01:07:30, Serial1/1

R3#Show ip ospf route | Inc 111.1.1.0/24

*>  111.1.1.0/24, Intra, cost 65, area 1
```

The following OSPF filtering mechanism works ONLY on intra-area routes. It instructs the router NOT to generate an LSA Type 3 for this prefix.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#area 1 range 111.1.1.0 255.255.255.0 not-advertise
```

Note as stated earlier the "Area range Not-advertise" command stops the generation of Type-3 LSAs, and this is why you see the prefix in the routing table of R3 as an intra-area route.

```
R3#Show ip route 111.1.1.0

Routing entry for 111.1.1.0/24
  Known via "ospf 1", distance 110, metric 65, type intra area
  Last update from 13.1.1.1 on Serial1/1, 00:00:27 ago
```


https://t.me/learningnets
© 2017 Narbik Kocharians. All rights reserved

```
  Routing Descriptor Blocks:
  * 13.1.1.1, from 0.0.0.1, 00:00:27 ago, via Serial1/1
      Route metric is 65, traffic share count is 1
```

The prefix is no longer in the database of R3 as LSA type 3, which means that the routers in area 0 or any other area downstream to area 0 will not have this prefix in their routing table or their link state database.

```
R3#Show ip ospf database summary 111.1.1.0

            OSPF Router with ID (0.0.0.3) (Process ID 1)

R3#Show ip ospf database router | Inc Area 1|_111.1.1.0

                Router Link States (Area 1)
    (Link ID) Network/subnet number: 111.1.1.0
```

On R4 or R5:

```
Rx#Show ip route 111.1.1.0

% Network not in table
```

Task 10

Configure the appropriate router/s such that none of the routers except R1 can see network 100.1.1.0 /24 in their routing table; DO NOT stop advertising this network, You should use a distribute-list to accomplish this task.

On R2:

You should always display the existing access-list/s and distribute-list/s before configuring one.

```
R2#Sh ip prefix-list

ip prefix-list TST: 2 entries
   seq 5 deny 1.1.1.0/24
   seq 10 permit 0.0.0.0/0 le 32

R2#Show run | S router ospf 1

router ospf 1
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
 router-id 0.0.0.2
 log-adjacency-changes
 network 2.2.2.2 0.0.0.0 area 1
 network 12.1.1.2 0.0.0.0 area 1
 network 23.1.1.2 0.0.0.0 area 1
 distribute-list prefix TST in
```

NOTE: The above show command verifies that there is already a distribute-list configured in the router configuration mode, therefore, we should try to modify the existing prefix-list that is applied by the existing distribute-list.

```
R2(config)#ip prefix-list TST seq 7 deny 100.1.1.0/24
```

To see the prefix-list after the above addition:

On R2:

```
R2#Show run | Inc ip pref

ip prefix-list TST seq 5 deny 1.1.1.0/24
ip prefix-list TST seq 7 deny 100.1.1.0/24
ip prefix-list TST seq 10 permit 0.0.0.0/0 le 32
```

To verify the configuration:

On R2:

```
R2#Show ip route 100.1.1.0

% Network not in table
```

Let's check R3:

On R3:

```
R3#Show ip prefix-list

ip prefix-list TST: 2 entries
   seq 5 deny 11.1.1.0/24
   seq 10 permit 0.0.0.0/0 le 32

R3#Show run | s ospf

router ospf 1
```


https://t.me/learningnets

```
 router-id 0.0.0.3
 area 1 range 111.1.1.0 255.255.255.0 not-advertise
 area 1 filter-list prefix TST out
 network 13.1.1.3 0.0.0.0 area 1
 network 23.1.1.3 0.0.0.0 area 1
 network 34.1.1.3 0.0.0.0 area 0

R3#Show ip route | Inc 100.1.1.0

O        100.1.1.0 [110/65] via 13.1.1.1, 00:08:38, Serial1/1
```

To configure R3:

```
R3(config)#ip prefix-list NET seq 5 deny 100.1.1.0/24
R3(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0 LE 32

R3(config)#router ospf 1
R3(config-router)#distribute-list prefix NET in
```

To verify the configuration:

On R3:

```
R3#Show ip route ospf | Inc 100.1.1.0
R3#
```

On R4:

On R4 the distance command can be used; referencing network 100.1.1.0 /24 in access-list 1, and the distance command sets the AD of the network referenced in access-list 1 to 255 sourcing from R2 (0.0.0.3 0.0.0.0); AD of 255 is unreachable and won't be injected into the routing table:

```
R4(config)#access-list 1 permit 100.1.1.0 0.0.0.255

R4(config)#router ospf 1
R4(config-router)#distance 255 0.0.0.3 0.0.0.0 1
```

On R5:

```
R5#Show ip route 100.1.1.0 255.255.255.0

% Network not in table

R5#Show ip ospf da summ 100.1.1.0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
OSPF Router with ID (0.0.0.5) (Process ID 1)
```

Note the above solution utilizes the behavior of OSPF when prefixes are advertised from area 0 into a non-zero area/s; when inter-area prefixes are advertised from area 0 into other areas, the behavior is like an internal redistribution, basically if R4 does NOT have the route in its routing table, it will NOT redistribute the route into area 2 and therefore, R5 or any other router in area 2 will NOT see the prefixe/s in their routing table or database.

Task 11

Configure R5 to redistribute Loopback 0, 1 and 2 interfaces into OSPF routing domain using the default cost.

On R5:

```
R5(config)#route-map TST permit 10
R5(config-route-map)#match interface lo0 lo1 lo2

R5(config)#router ospf 1
R5(config-router)#redistribute connected subnets route-map TST
```

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     5.5.5.0 [110/20] via 45.1.1.5, 00:00:08, Serial1/5
O E2     50.5.5.0 [110/20] via 45.1.1.5, 00:00:08, Serial1/5
O E2     55.5.5.0 [110/20] via 45.1.1.5, 00:00:08, Serial1/5
```

Task 12

Configure the appropriate router such that none of the routers except R5 can see network 5.5.5.0 /24 in their routing table.

© All rights reserved https://t.me/learningnets

On R5:

```
R5(config)#ip prefix-list NET seq 5 deny 5.5.5.0/24
R5(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0 LE 32

R5(config)#router ospf 1
R5(config-router)#distribute-list prefix NET OUT
```

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     50.5.5.0 [110/20] via 45.1.1.5, 00:02:07, Serial1/5
O E2     55.5.5.0 [110/20] via 45.1.1.5, 00:02:07, Serial1/5
```

On R1:

```
R1#Show ip route ospf | Inc E2
       E1 - OSPF external type 1, E2 - OSPF external type 2

O E2     50.5.5.0 [110/20] via 13.1.1.3, 00:02:55, Serial1/3
O E2     55.5.5.0 [110/20] via 13.1.1.3, 00:02:55, Serial1/3
```

Note this is the ONLY scenario where the "distribute-list OUT" command works in OSPF. This command MUST be configured on the ASBR or else it will not have any effect whatsoever. This command filters external routes.

```
R5#Show ip ospf da ext | Inc 5

            OSPF Router with ID (0.0.0.5) (Process ID 1)
                Type-5 AS External Link States
  Link State ID: 50.5.5.0 (External Network Number )
  Advertising Router: 0.0.0.5
  Checksum: 0xC995
  Link State ID: 55.5.5.0 (External Network Number )
  Advertising Router: 0.0.0.5
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 13

Configure the appropriate router such that none of the routers except R5 can see network 50.5.5.0 /24 in their routing table or database. You should NOT use the solution that was implemented in the previous task.

The following command is used to filter The external routes ("E" and/or "N"), this command must be configured on an ASBR or the router that originated the external route/s, and when configured, it filters the specified prefix from the OSPF Link state database.

On R5:

```
R5(config)#router ospf 1
R5(config-router)#summary-address 50.5.5.0 255.255.255.0 not-advertise
```

The following show command reveals that network 50.5.5.0 /24 is NOT in the link state database of R5.

```
R5#Show ip ospf da ext | Inc 5

            OSPF Router with ID (0.0.0.5) (Process ID 1)
                Type-5 AS External Link States
  Link State ID: 55.5.5.0 (External Network Number )
  Advertising Router: 0.0.0.5
```

To verify the configuration:

On R4:

```
R4#Show ip route 50.5.5.0

% Network not in table
```

On R1:

```
R1#Show ip route 50.5.5.0

% Network not in table
```


https://t.me/learningnets

Task 14

Configure the appropriate router such that R1 does NOT have network 55.5.5.0 /24 in its routing table.

To see R1's existing routing table:

On R1:

```
R1#Show ip route 55.5.5.0

Routing entry for 55.5.5.0/24
  Known via "ospf 1", distance 110, metric 20, type extern 2, forward
metric 192
  Last update from 13.1.1.3 on Serial1/3, 00:10:17 ago
  Routing Descriptor Blocks:
  * 13.1.1.3, from 0.0.0.5, 00:10:17 ago, via Serial1/3
      Route metric is 20, traffic share count is 1
```

NOTE: There are no access-lists or ip prefix-lists configured on this router:

```
R1#Show access-list
R1#

R1#Show ip prefix-list
R1#

R1(config)#ip prefix-list NET seq 5 deny 55.5.5.0/24
R1(config)#ip prefix-list NET seq 10 permit 0.0.0.0/0 LE 32

R1(config)#router ospf 1
R1(config-router)#distribute-list prefix NET in
```

To verify the configuration:

On R1:

```
R1#Show ip route 55.5.5.0

% Network not in table
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 15

Remove any filtering used in the previous tasks, if this configuration is performed successfully, all the routers in this routing domain should have every route advertised and redistributed in this lab.

On R1:

```
R1(config)#No ip prefix NET

R1(config)#Router ospf 1
R1(config-router)#No distribute-list prefix NET in

R1#Sh run | s router ospf

router ospf 1
 router-id 0.0.0.1
 network 1.1.1.1 0.0.0.0 area 1
 network 11.1.1.1 0.0.0.0 area 1
 network 12.1.1.1 0.0.0.0 area 1
 network 13.1.1.1 0.0.0.0 area 1
 network 100.1.1.1 0.0.0.0 area 1
 network 111.1.1.1 0.0.0.0 area 1
```

On R2:

```
R2(config)#No ip prefix-list TST

R2(config)#router ospf 1
R2(config-router)#No distribute-list prefix TST in

R2#Show run | s ospf

 ip ospf network point-to-point
router ospf 1
 router-id 0.0.0.2
 network 2.2.2.2 0.0.0.0 area 1
 network 12.1.1.2 0.0.0.0 area 1
 network 23.1.1.2 0.0.0.0 area 1
```

On R3:

```
R3(config)#No ip prefix-list NET
R3(config)#No ip prefix-list TST
```


https://t.me/learningnets

```
R3(config)#Router ospf 1
R3(config-router)#No area 1 range 111.1.1.0 255.255.255.0 not-advertise
R3(config-router)#No distribute-list prefix NET in
R3(config-router)#No area 1 filter-list prefix TST out

R3#Sh run | s ospf

router ospf 1
 router-id 0.0.0.3
 network 13.1.1.3 0.0.0.0 area 1
 network 23.1.1.3 0.0.0.0 area 1
 network 34.1.1.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#No ip prefix-list TST
R4(config)#No access-list 1

R4(config)#Router ospf 1
R4(config-router)#No distance 255 0.0.0.3 0.0.0.0 1
R4(config-router)#No area 2 filter-list prefix TST in

R4#Show run | s ospf

router ospf 1
 router-id 0.0.0.4
 network 34.1.1.4 0.0.0.0 area 0
 network 45.1.1.4 0.0.0.0 area 2
```

On R5:

```
R5(config)#No ip prefix-list NET

R5(config)#Router ospf 1
R5(config-router)#No distribute-list prefix NET out
R5(config-router)#No summary-address 50.5.5.0 255.255.255.0 not-advertise

R5#Show run | s ospf

router ospf 1
 router-id 0.0.0.5
 summary-address 50.5.5.0 255.255.255.0
 redistribute connected subnets route-map TST
 network 45.1.1.5 0.0.0.0 area 2
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

To verify the configuration:

On R4:

```
R4#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
O IA     1.1.1.0 [110/129] via 34.1.1.3, 00:04:14, Serial1/3
      2.0.0.0/24 is subnetted, 1 subnets
O IA     2.2.2.0 [110/129] via 34.1.1.3, 00:04:14, Serial1/3
      5.0.0.0/24 is subnetted, 1 subnets
O E2     5.5.5.0 [110/20] via 45.1.1.5, 00:02:45, Serial1/5
      11.0.0.0/24 is subnetted, 1 subnets
O IA     11.1.1.0 [110/129] via 34.1.1.3, 00:04:14, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/909] via 34.1.1.3, 00:04:14, Serial1/3
      13.0.0.0/24 is subnetted, 1 subnets
O IA     13.1.1.0 [110/128] via 34.1.1.3, 00:04:14, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/128] via 34.1.1.3, 00:04:14, Serial1/3
      50.0.0.0/24 is subnetted, 1 subnets
O E2     50.5.5.0 [110/20] via 45.1.1.5, 00:02:12, Serial1/5
      55.0.0.0/24 is subnetted, 1 subnets
O E2     55.5.5.0 [110/20] via 45.1.1.5, 00:04:14, Serial1/5
      100.0.0.0/24 is subnetted, 1 subnets
O IA     100.1.1.0 [110/129] via 34.1.1.3, 00:04:14, Serial1/3
      111.0.0.0/24 is subnetted, 1 subnets
O IA     111.1.1.0 [110/129] via 34.1.1.3, 00:04:14, Serial1/3
```

On R3:

```
R5#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 [110/65] via 13.1.1.1, 00:07:02, Serial1/1
      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/65] via 23.1.1.2, 00:07:02, Serial1/2
      5.0.0.0/24 is subnetted, 1 subnets
O E2     5.5.5.0 [110/20] via 34.1.1.4, 00:03:48, Serial1/4
      11.0.0.0/24 is subnetted, 1 subnets
O        11.1.1.0 [110/65] via 13.1.1.1, 00:07:02, Serial1/1
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/845] via 23.1.1.2, 00:07:02, Serial1/2
```


© [illegible] All rights reserved
https://t.me/learningnets

```
                [110/845] via 13.1.1.1, 00:07:02, Serial1/1
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/128] via 34.1.1.4, 00:07:02, Serial1/4
      50.0.0.0/24 is subnetted, 1 subnets
O E2     50.5.5.0 [110/20] via 34.1.1.4, 00:03:15, Serial1/4
      55.0.0.0/24 is subnetted, 1 subnets
O E2     55.5.5.0 [110/20] via 34.1.1.4, 00:07:02, Serial1/4
      100.0.0.0/24 is subnetted, 1 subnets
O        100.1.1.0 [110/65] via 13.1.1.1, 00:07:02, Serial1/1
      111.0.0.0/24 is subnetted, 1 subnets
O        111.1.1.0 [110/65] via 13.1.1.1, 00:07:02, Serial1/1
```

On R2:

```
R3#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 [110/129] via 23.1.1.3, 00:10:15, Serial1/3
      5.0.0.0/24 is subnetted, 1 subnets
O E2     5.5.5.0 [110/20] via 23.1.1.3, 00:04:28, Serial1/3
      11.0.0.0/24 is subnetted, 1 subnets
O        11.1.1.0 [110/129] via 23.1.1.3, 00:10:15, Serial1/3
      13.0.0.0/24 is subnetted, 1 subnets
O        13.1.1.0 [110/128] via 23.1.1.3, 00:10:15, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/128] via 23.1.1.3, 00:10:15, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/192] via 23.1.1.3, 00:10:15, Serial1/3
      50.0.0.0/24 is subnetted, 1 subnets
O E2     50.5.5.0 [110/20] via 23.1.1.3, 00:03:56, Serial1/3
      55.0.0.0/24 is subnetted, 1 subnets
O E2     55.5.5.0 [110/20] via 23.1.1.3, 00:10:15, Serial1/3
      100.0.0.0/24 is subnetted, 1 subnets
O        100.1.1.0 [110/129] via 23.1.1.3, 00:10:15, Serial1/3
      111.0.0.0/24 is subnetted, 1 subnets
O        111.1.1.0 [110/129] via 23.1.1.3, 00:10:15, Serial1/3
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 16

Advertise Lo0 interface of R3 in area 1; Lo1 interfaces of R3 and R4 in area 0, and Lo0 interface of R4 in area 2. These Loopback interfaces MUST be advertised with their correct mask.

On R3:

```
R3(config)#int range lo0 - 1
R3(config-if-range)#ip ospf net point-to-point

R3(config-if)#router ospf 1
R3(config-router)#Netw 3.3.3.3 0.0.0.0 area 1
R3(config-router)#Netw 30.3.3.3 0.0.0.0 area 0
```

On R4:

```
R4(config)#int range lo0 - 1
R4(config-if-range)#ip ospf net point-to-point

R4(config-if)#router ospf 1
R4(config-router)#Netw 4.4.4.4 0.0.0.0 area 2
R4(config-router)#Netw 40.4.4.4 0.0.0.0 area 0
```

To verify the configuration:

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/129] via 13.1.1.3, 00:15:41, Serial1/3
      3.0.0.0/24 is subnetted, 1 subnets
O        3.3.3.0 [110/65] via 13.1.1.3, 00:01:35, Serial1/3
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/129] via 13.1.1.3, 00:00:19, Serial1/3
      5.0.0.0/24 is subnetted, 1 subnets
O E2     5.5.5.0 [110/20] via 13.1.1.3, 00:08:47, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/128] via 13.1.1.3, 00:15:41, Serial1/3
      30.0.0.0/24 is subnetted, 1 subnets
O IA     30.3.3.0 [110/65] via 13.1.1.3, 00:01:18, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
```


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

```
O IA     34.1.1.0 [110/128] via 13.1.1.3, 00:15:41, Serial1/3
      40.0.0.0/24 is subnetted, 1 subnets
O IA     40.4.4.0 [110/129] via 13.1.1.3, 00:00:14, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/192] via 13.1.1.3, 00:15:41, Serial1/3
      50.0.0.0/24 is subnetted, 1 subnets
O E2     50.5.5.0 [110/20] via 13.1.1.3, 00:08:14, Serial1/3
      55.0.0.0/24 is subnetted, 1 subnets
O E2     55.5.5.0 [110/20] via 13.1.1.3, 00:15:41, Serial1/3
```

Task 17

Configure the appropriate router/s such that the routers in area 2 do NOT see any of the networks advertised by any of the routers in this topology, but routers R1, R2 and R3 and R4 should see all the networks advertised by the existing and future router/s advertised in area 2.

By default, all outgoing LSAs are flooded to the interface. The following configuration prevents flooding of ALL OSPF LSAs out of a given interface, in this case R4's S1/5 interface.

On R4:

```
R4(config)#int S1/5
R4(config-subif)#ip ospf database-filter all out
```

For this filtering mechanism to work, the OSPF process must be cleared.

On R5:

```
R5#Clear ip ospf proc
Reset ALL OSPF processes? [no]: y
```

Note R4 and R5 are still maintaining their neighbor adjacency

```
R5#Sh ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.4           0   FULL/  -        00:00:39    45.1.1.4        Serial1/4
```

You can see that R5 does NOT have any of the routes from the other routers; this includes R4's routes which are advertised within the same area. Remember that the filtering is done on the S1/5 interface of

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R4, meaning that all LSAs (Regardless of their area) are filtered.

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set
```

Note R1 has all the routes including the ones advertised by R5.

On R1:

```
R1#Show ip route ospf | B Gate
Gateway of last resort is not set

      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/129] via 13.1.1.3, 00:20:22, Serial1/3
      3.0.0.0/24 is subnetted, 1 subnets
O        3.3.3.0 [110/65] via 13.1.1.3, 00:06:16, Serial1/3
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/129] via 13.1.1.3, 00:05:00, Serial1/3
      5.0.0.0/24 is subnetted, 1 subnets
O E2     5.5.5.0 [110/20] via 13.1.1.3, 00:01:59, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/128] via 13.1.1.3, 00:20:22, Serial1/3
      30.0.0.0/24 is subnetted, 1 subnets
O IA     30.3.3.0 [110/65] via 13.1.1.3, 00:05:59, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/128] via 13.1.1.3, 00:20:22, Serial1/3
      40.0.0.0/24 is subnetted, 1 subnets
O IA     40.4.4.0 [110/129] via 13.1.1.3, 00:04:55, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/192] via 13.1.1.3, 00:20:22, Serial1/3
      50.0.0.0/24 is subnetted, 1 subnets
O E2     50.5.5.0 [110/20] via 13.1.1.3, 00:01:59, Serial1/3
      55.0.0.0/24 is subnetted, 1 subnets
O E2     55.5.5.0 [110/20] via 13.1.1.3, 00:01:59, Serial1/3
```

Task 18

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 10
Virtual-links and GRE Tunnels

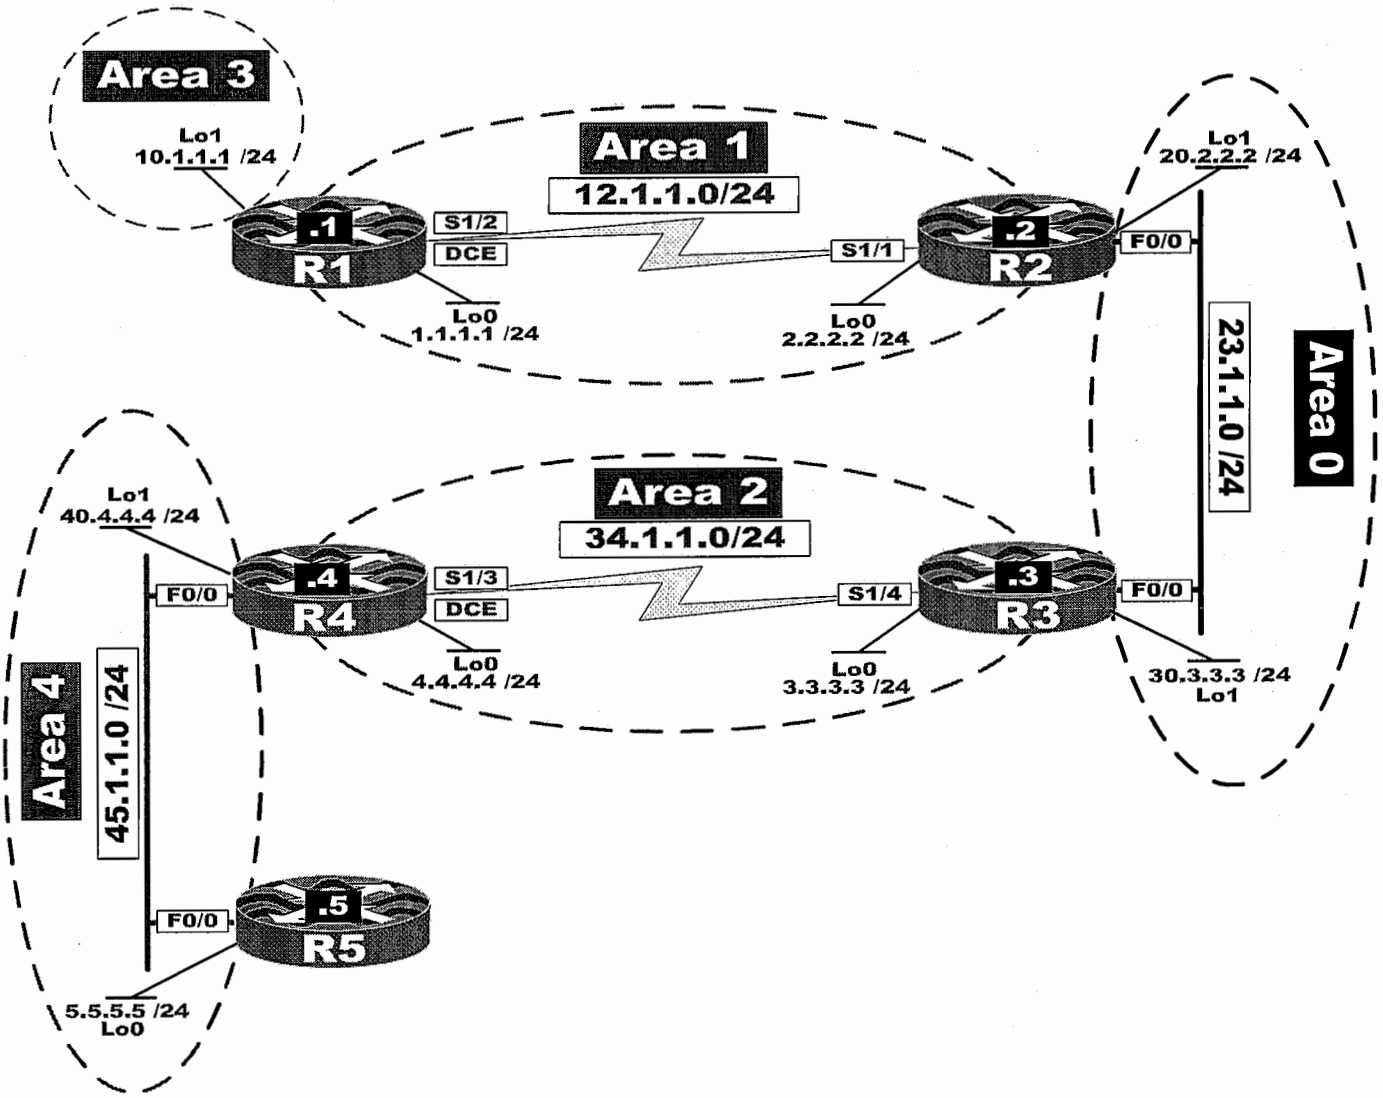

Task 1

Configure the Routers and the Switches based on the above diagram. DO NOT configure OSPF.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On SW1:

```
SW1(config)#Int range F0/2-3
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc vlan 23
SW1(config-if-range)#No shut

SW1(config)#Int range F0/4-5
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc vlan 45
SW1(config-if-range)#No shut
```

On R1:

```
R1(config)#Int Lo0
R1(config-if)#IP address 1.1.1.1 255.255.255.0

R1(config-if)#Int Lo1
R1(config-if)#IP address 10.1.1.1 255.255.255.0

R1(config)#Int S1/2
R1(config-if)#IP address 12.1.1.1 255.255.255.0
R1(config-if)#Clock rate 64000
R1(config-if)#No shut
```

On R2:

```
R2(config)#Int S1/1
R2(config-if)#IP address 12.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config)#Int F0/0
R2(config-if)#IP address 23.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config-if)#Int Lo0
R2(config-if)#IP address 2.2.2.2 255.255.255.0

R2(config-if)#Int Lo1
R2(config-if)#IP address 20.2.2.2 255.255.255.0
```

On R3:

```
R3(config)#Int S1/4
R3(config-subif)#IP address 34.1.1.3 255.255.255.0
```


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

```
R3(config-subif)#Clock rate 64000
R3(config-if)#No shut

R3(config-fr-dlci)#Int F0/0
R3(config-if)#IP address 23.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config-if)#Int Lo0
R3(config-if)#IP addr 3.3.3.3 255.255.255.0

R3(config-if)#Int Lo1
R3(config-if)#IP address 30.3.3.3 255.255.255.0
```

On R4:

```
R4(config)#Int S1/3
R4(config-subif)#IP address 34.1.1.4 255.255.255.0
R4(config-if)#No shut

R4(config)#Int Lo0
R4(config-if)#IP address 4.4.4.4 255.255.255.0

R4(config)#Int Lo1
R4(config-if)#IP address 40.4.4.4 255.255.255.0

R4(config)#Int F0/0
R4(config-if)#IP address 45.1.1.4 255.255.255.0
R4(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/0
R5(config-if)#IP address 45.1.1.5 255.255.255.0

R5(config-if)#No shut
R5(config-if)#Int Lo0
R5(config-if)#IP addr 5.5.5.5 255.255.255.0
```

To verify the configuration:

On R2:

```
R2#Ping 12.1.1.1

Type escape sequence to abort.
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

On R4:

```
R4#Ping 34.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R4#Ping 45.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 45.1.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Task 2

Configure OSPF based on the following chart:

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

Router	Interface	OSPF Area	Router-ID
R1	Lo0 S1/2 Lo1	1 1 3	0.0.0.1
R2	Lo0 S1/1 F0/0 Lo1	1 1 0 0	0.0.0.2
R3	Lo1 F0/0 S1/4 Lo0	0 0 2 2	0.0.0.3
R4	S1/3 Lo0 Lo1 F0/0	2 2 4 4	0.0.0.4
R5	F0/0 Lo0	4 4	0.0.0.5

On R1:

```
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 are 3
R1(config-router)#netw 12.1.1.1 0.0.0.0 are 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 are 1
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#netw 12.1.1.2 0.0.0.0 are 1
R2(config-router)#netw 2.2.2.2 0.0.0.0 are 1
R2(config-router)#netw 23.1.1.2 0.0.0.0 are 0
R2(config-router)#netw 20.2.2.2 0.0.0.0 are 0
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

On R3:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R3(config-if)#router ospf 1
R3(config-router)#router-id 0.0.0.3
R3(config-router)#netw 23.1.1.3 0.0.0.0 are 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 are 2
R3(config-router)#netw 34.1.1.3 0.0.0.0 are 2
R3(config-router)#netw 30.3.3.3 0.0.0.0 are 0
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

On R4:

```
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 2
R4(config-router)#netw 4.4.4.4 0.0.0.0 are 2
R4(config-router)#netw 45.1.1.4 0.0.0.0 are 4
R4(config-router)#netw 40.4.4.4 0.0.0.0 are 4
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 from LOADING to FULL,
Loading Done
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#router-id 0.0.0.5
R5(config-router)#netw 45.1.1.5 0.0.0.0 are 4
R5(config-router)#netw 5.5.5.5 0.0.0.0 are 4
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

Task 3

Ensure that the network/s from area 3 are reachable by R2, R3, and R4. Do NOT use a GRE Tunnel to accomplish this task.

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

In OSPF every non-zero area MUST be connected to area 0 (The backbone area), this connection can be a physical or a logical connection. The area through which the Virtual-link is configured is called the transit area, this area must have full routing information and can't be a stub area.

When the virtual-link is forst configured it is in down state because R1 (0.0.0.1) and R2 (0.0.0.2) don't have NLRI (Network Layer Reachability Information) to each other, therefore, all the LSAs need to be flooded and the SPF algorithm needs to be calculated so these routers know how to reach each other through area 1. Once the routers have NLRI to each other, they form an adjacency across the virtual-link.

It is important to note that the OSPF packets between the endpoints of the virtual-link are NOT Multicast packets; they are tunneled packets from 12.1.1.1 (R1) to 12.1.1.2 (R2) destination or vice versa.

The virtual-link configuration must be performed on the ABRs, and if one of the endpoints is not an ABR, the virtual-link will not transition into up state; to be a BB router, you must have a minimum of one link in the BB area; once the routers form an adjacency on the virtual-link, R1 becomes a Backbone router, and once R1 becomes a BB router, and since it has another connection to area 3, it becomes an ABR.

Since area 3 is NOT connected to area 0; the routers in the other areas won't be able to see the route/s advertised by this area, therefore, a virtual-link or a GRE tunnel must be configured to connect R1 (The router that connects area 3 to area 1) to R2 in area 0.

Use the following steps to configure a virtual-link:

Step 1:
Since the virtual-link is terminated on the ABRs, the Router-id of these routers MUST be determined:

On R1:

```
R1#Show ip ospf | Inc ID
 Routing Process "ospf 1" with ID 0.0.0.1
```

On R2:

```
R2#Show ip ospf | Inc ID

 Routing Process "ospf 1" with ID 0.0.0.2
```

Step 2:
Ensure that the router-ids are stable, meaning that they are statically configured on the routers, unless you are in the CCIE lab and you are told not to configure router-ids:

On R1:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R1#Show run | S router ospf

router ospf 1
 router-id 0.0.0.1
 network 1.1.1.1 0.0.0.0 area 1
 network 10.1.1.1 0.0.0.0 area 3
 network 12.1.1.1 0.0.0.0 area 1
```

On R2:

```
R2#Show run | S router ospf

router ospf 1
 router-id 0.0.0.2
 network 2.2.2.2 0.0.0.0 area 1
 network 12.1.1.2 0.0.0.0 area 1
 network 20.2.2.2 0.0.0.0 area 0
 network 23.1.1.2 0.0.0.0 area 0
```

Step 3:

Ensure the the transit area is NOT a stub:

Well based on the above show command we can see that these the transit area, area 1, is not a stub area.

Step 4:

Configure the virtual-links:

On R1:

```
R1(config)#router ospf 1
R1(config-router)#area 1 virtual-link 0.0.0.2
```

On R2:

You should see the following console message every hello interval, and they should stop once the virtual-link is configured successfully:

```
%OSPF-4-ERRRCV: Received invalid packet: mismatched area ID, from backbone
area must be virtual-link but not found from 12.1.1.1, Serial1/1

R2(config)#router ospf 1
R2(config-router)#area 1 virtual-link 0.0.0.1
```

To verify and test the configuration:

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

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 [110/782] via 12.1.1.1, 00:46:58, Serial1/1
      3.0.0.0/32 is subnetted, 1 subnets
O IA     3.3.3.3 [110/2] via 23.1.1.3, 00:45:18, FastEthernet0/0
      4.0.0.0/32 is subnetted, 1 subnets
O IA     4.4.4.4 [110/783] via 23.1.1.3, 00:44:25, FastEthernet0/0
      10.0.0.0/32 is subnetted, 1 subnets
O IA     10.1.1.1 [110/782] via 12.1.1.1, 00:00:18, Serial1/1
      30.0.0.0/32 is subnetted, 1 subnets
O        30.3.3.3 [110/2] via 23.1.1.3, 00:45:18, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 23.1.1.3, 00:45:18, FastEthernet0/0

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 = 48/50/52 ms

R2#Show ip ospf virtual-links

Virtual Link OSPF_VL0 to router 0.0.0.1 is up
  Run as demand circuit
  DoNotAge LSA allowed.
  Transit area 1, via interface Serial1/1
 Topology-MTID     Cost     Disabled     Shutdown      Topology Name
        0           781        no           no            Base
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
    Adjacency State FULL (Hello suppressed)
    Index 2/3, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

We can see that the virtual-link is up and the cost of this virtual-link is 781. Since the OSPF packets are tunneled packets from 12.1.1.1 (R1) to 12.1.1.2 (R2), the cost of the virtual-link is based on the OSPF cost of all links from R1's S1/2 to R2's S1/1 interface. Let's see the OSPF cost of the serial link between R1 and R2:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2#Show ip ospf interface S1/1 | Inc Cost

  Process ID 1, Router ID 0.0.0.2, Network Type POINT_TO_POINT, Cost: 781
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

On R3:

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

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/783] via 23.1.1.2, 00:54:56, FastEthernet0/0
      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/2] via 23.1.1.2, 00:54:56, FastEthernet0/0
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/782] via 34.1.1.4, 00:54:01, Serial1/4
      10.0.0.0/32 is subnetted, 1 subnets
O IA     10.1.1.1 [110/783] via 23.1.1.2, 00:09:53, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/782] via 23.1.1.2, 00:54:56, FastEthernet0/0
      20.0.0.0/32 is subnetted, 1 subnets
O        20.2.2.2 [110/2] via 23.1.1.2, 00:54:56, FastEthernet0/0

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 = 48/50/52 ms
```

Task 4

Ensure that the routers in area 4 can see and have NLRI to all routes advertised in this routing domain. You should use Loopback 0 IP addresses, DO NOT change the area in which they are already configured in, DO NOT use a virtual-link to accomplish this task.

Let's check the routing table of R5:

On R5:

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

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      40.0.0.0/32 is subnetted, 1 subnets
O        40.4.4.4 [110/2] via 45.1.1.4, 01:00:16, FastEthernet0/0
```

The reason R5 only sees network 40.4.4.0/24 is because Area 4 does not have a connection (Logical or Physical) to area 0. In order to rectify this problem we must configure a virtual-link, since configuring virtual-links is prohibited by the condition of this task, a GRE tunnel is configured instead; the task also states that we must use an existing IP address, let's use the loopback 0 IP addresses of R4 and R3:

On R4:

```
R4(config)#Int tunnel 43
R4(config-if)#ip unnumbered Lo0
R4(config-if)#tunnel source 34.1.1.4
R4(config-if)#tunnel destination 34.1.1.3
```

On R3:

```
R3(config)#int tunnel 34
R3(config-if)#ip unnumbered Lo0
R3(config-if)#tunnel source 34.1.1.3
R3(config-if)#tunnel destination 34.1.1.4
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Tunnel34 from LOADING to FULL,
Loading Done
```

When configuring a GRE tunnel the IP address of the tunnel MUST be configured in area 0. This can become a major problem when "IP unnumbered Loopback" interface configuration commands are used, especially when the loopback interfaces are configured in different areas.

To Verify the configuration:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      40.0.0.0/32 is subnetted, 1 subnets
O        40.4.4.4 [110/2] via 45.1.1.4, 01:07:31, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

R5 can not see any routes from the other areas, why?
The tunnel interface must be configured in area 0, let's verify this:

```
R4#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     2               4.4.4.4/24         1     LOOP  0/0
Tu43         1     2               Unnumbered Lo0     1000  P2P   1/1
Se1/3        1     2               34.1.1.4/24        781   P2P   1/1
Lo1          1     4               40.4.4.4/24        1     LOOP  0/0
Fa0/0        1     4               45.1.1.4/24        1     DR    1/1
```

We can see the problem, the tunnel interface is configured based on the loopback0 interface which is running in area 2, since the task states that we can not change the area assignment of any interface, we can go to the tunnel interface and configure OSPF directly under the tunnel interface:

On R4:

```
R4(config)#Int tunnel 43
R4(config-if)#IP ospf 1 area 0
```

On R3:

```
R3(config)#int tunnel 34
R3(config-if)#IP ospf 1 area 0
```

We should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on Tunnel34 from LOADING to FULL,
Loading Done
```

To verify and 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, 1 subnets
O IA     1.1.1.1 [110/1784] via 45.1.1.4, 00:00:50, FastEthernet0/0
      2.0.0.0/32 is subnetted, 1 subnets
O IA     2.2.2.2 [110/1003] via 45.1.1.4, 00:00:50, FastEthernet0/0
      3.0.0.0/32 is subnetted, 1 subnets
O IA     3.3.3.3 [110/783] via 45.1.1.4, 00:01:15, FastEthernet0/0
```


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

```
      4.0.0.0/32 is subnetted, 1 subnets
O IA     4.4.4.4 [110/2] via 45.1.1.4, 00:01:15, FastEthernet0/0
      10.0.0.0/32 is subnetted, 1 subnets
O IA     10.1.1.1 [110/1784] via 45.1.1.4, 00:00:50, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/1783] via 45.1.1.4, 00:00:50, FastEthernet0/0
      20.0.0.0/32 is subnetted, 1 subnets
O IA     20.2.2.2 [110/1003] via 45.1.1.4, 00:00:50, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/1002] via 45.1.1.4, 00:00:50, FastEthernet0/0
      30.0.0.0/32 is subnetted, 1 subnets
O IA     30.3.3.3 [110/1002] via 45.1.1.4, 00:00:50, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 45.1.1.4, 00:01:15, FastEthernet0/0
      40.0.0.0/32 is subnetted, 1 subnets
O        40.4.4.4 [110/2] via 45.1.1.4, 01:13:26, FastEthernet0/0

R5#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 = 64/64/64 ms

R4#Show ip ospf interface brief

Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Tu43         1     0               Unnumbered Lo0     1000  P2P   1/1
Lo0          1     2               4.4.4.4/24         1     LOOP  0/0
Se1/3        1     2               34.1.1.4/24        781   P2P   1/1
Lo1          1     4               40.4.4.4/24        1     LOOP  0/0
Fa0/0        1     4               45.1.1.4/24        1     DR    1/1
```

Unlike virtual-links that their cost is based on the OSPF cost of all links between the ABRs, the virtual-links cost of the GRE tunnel is NOT based on the cost of all links between the endpoints of the tunnel. Let's verify:

On R4:

```
R4#Show ip ospf inter tunnel 43 | Inc Cost

  Process ID 1, Router ID 0.0.0.4, Network Type POINT_TO_POINT, Cost: 1000
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 5

Erase the startup config and reload the routers before proceeding to the next lab.

https://t.me/learningnets

Lab 11
OSPF Stub, Totally Stubby, and NSSA Areas

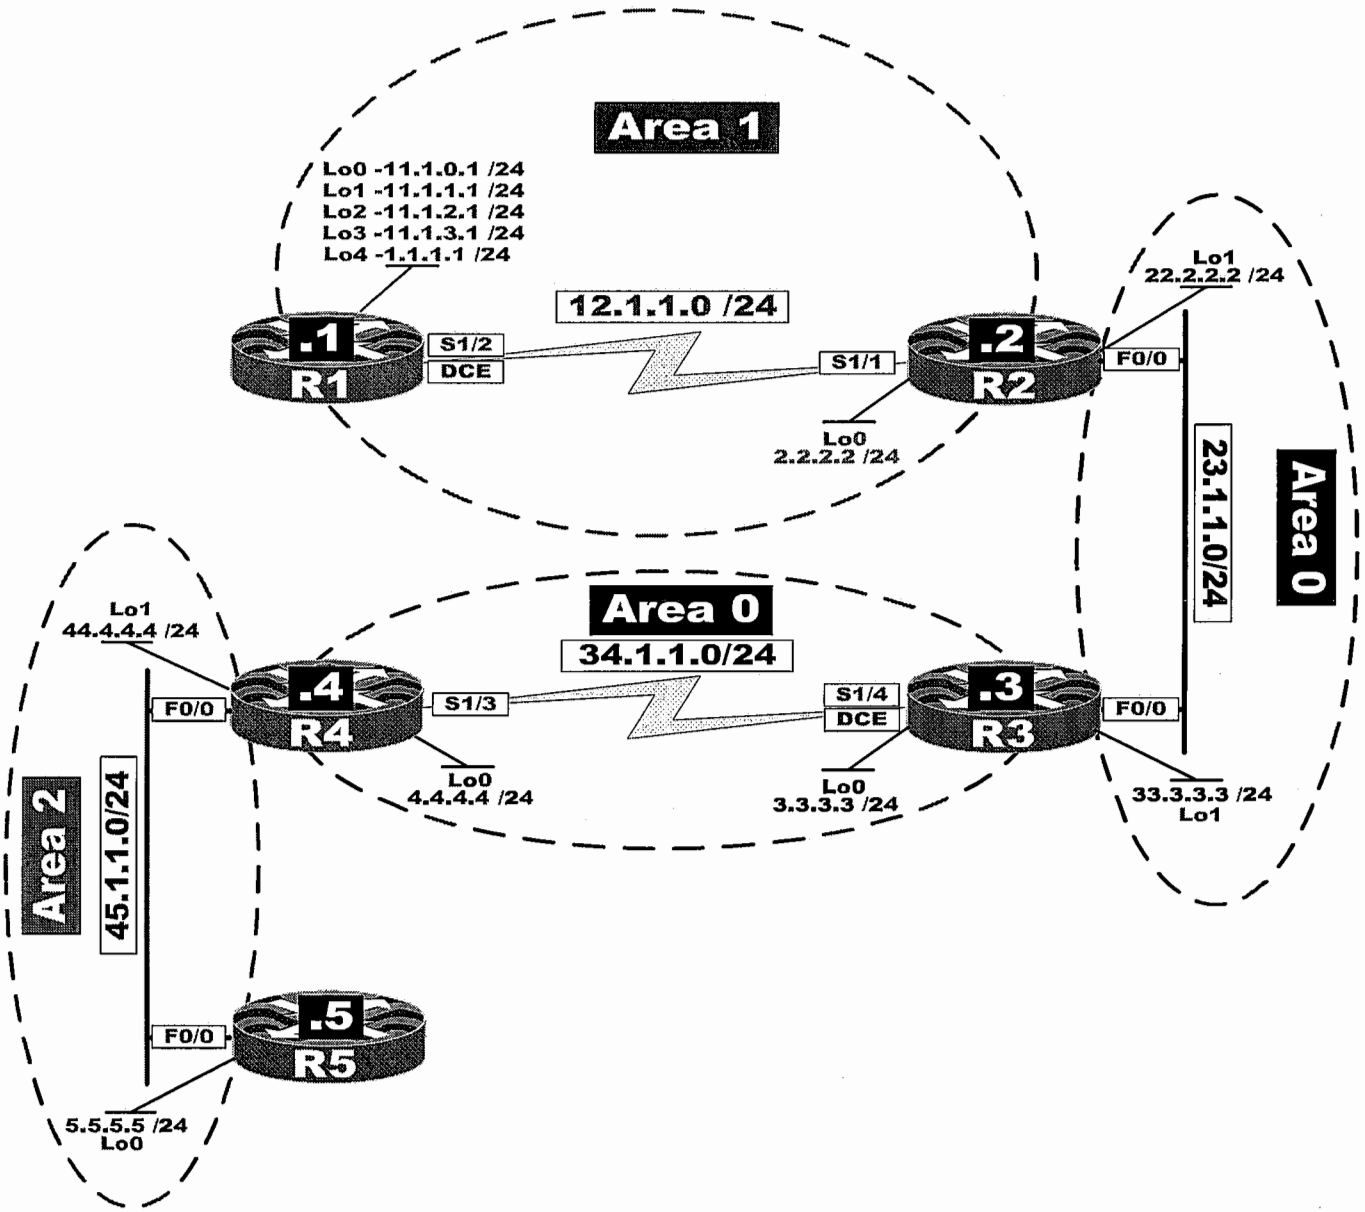

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 1

Configure the Routers and the Switches based on the above diagram. DO NOT configure OSPF.

On SW1:

```
SW1(config)#Int range F0/2-3
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc vlan 23
SW1(config-if-range)#No shut

SW1(config)#Int range F0/4-5
SW1(config-if-range)#Swi mode acc
SW1(config-if-range)#Swi acc vlan 45
SW1(config-if-range)#No shut
```

On R1:

```
R1(config)#Int Lo0
R1(config-if)#IP address 11.1.0.1 255.255.255.0

R1(config)#Int Lo1
R1(config-if)#IP address 11.1.1.1 255.255.255.0

R1(config)#Int Lo2
R1(config-if)#IP address 11.1.2.1 255.255.255.0

R1(config)#Int Lo3
R1(config-if)#IP address 11.1.3.1 255.255.255.0

R1(config)#Int Lo4
R1(config-if)#IP address 1.1.1.1 255.255.255.0

R1(config)#Int S1/2
R1(config-if)#IP address 12.1.1.1 255.255.255.0
R1(config-if)#Clock rate 64000
R1(config-if)#No shut
```

On R2:

```
R2(config)#Int S1/1
R2(config-if)#IP address 12.1.1.2 255.255.255.0
R2(config-if)#No shut
```


© [illegible] reserved
https://t.me/learningnets

```
R2(config)#Int F0/0
R2(config-if)#IP address 23.1.1.2 255.255.255.0
R2(config-if)#No shut

R2(config-if)#Int Lo0
R2(config-if)#IP address 2.2.2.2 255.255.255.0

R2(config-if)#Int Lo1
R2(config-if)#IP address 20.2.2.2 255.255.255.0
```

On R3:

```
R3(config)#Int S1/4
R3(config-if)#IP address 34.1.1.3 255.255.255.0
R3(config-if)#Clock rate 64000
R3(config-if)#No shut

R3(config)#Int F0/0
R3(config-if)#IP address 23.1.1.3 255.255.255.0
R3(config-if)#No shut

R3(config-if)#Int Lo0
R3(config-if)#IP addr 3.3.3.3 255.255.255.0

R3(config-if)#Int Lo1
R3(config-if)#IP address 30.3.3.3 255.255.255.0
```

On R4:

```
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 4.4.4.4 255.255.255.0

R4(config)#Int Lo1
R4(config-if)#IP address 40.4.4.4 255.255.255.0

R4(config)#Int F0/0
R4(config-if)#IP address 45.1.1.4 255.255.255.0
R4(config-if)#No shut
```

On R5:

```
R5(config)#Int F0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5(config-if)#IP address 45.1.1.5 255.255.255.0

R5(config-if)#No shut
R5(config-if)#Int Lo0
R5(config-if)#IP addr 5.5.5.5 255.255.255.0
```

To verify the configuration:

On R2:

```
R2#Ping 12.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R2#Ping 23.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

On R4:

```
R4#Ping 34.1.1.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R4#Ping 45.1.1.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 45.1.1.5, 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 OSPF Area 1 based on the following chart; the loopback interfaces should be advertised with their correct mask.

Router	Interface	Area	Router-ID
R1	Lo0 – Lo4 S1/2	1 1	0.0.0.1
R2	Lo0 S1/1	1 1	0.0.0.2

On R1:

```
R1(config)#int range lo0 - 4
R1(config-if-range)#IP ospf network point-to-point

R1(config)#Router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#Netw 11.1.0.1 0.0.0.0 area 1
R1(config-router)#Netw 11.1.1.1 0.0.0.0 area 1
R1(config-router)#Netw 11.1.2.1 0.0.0.0 area 1
R1(config-router)#Netw 11.1.3.1 0.0.0.0 area 1
R1(config-router)#Netw 1.1.1.1 0.0.0.0 area 1
R1(config-router)#Netw 12.1.1.1 0.0.0.0 area 1
```

On R2:

```
R2(config)#Int range lo0 - 1
R1(config-if-range)#IP ospf network point-to-point

R2(config)#Router ospf 1
R2(config-router)#router-id 0.0.0.2
R2(config-router)#Netw 12.1.1.2 0.0.0.0 area 1
R2(config-router)#Netw 2.2.2.2 0.0.0.0 area 1
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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 [110/782] via 12.1.1.1, 00:01:20, Serial1/1
      11.0.0.0/24 is subnetted, 4 subnets
O        11.1.0.0 [110/782] via 12.1.1.1, 00:01:20, Serial1/1
O        11.1.1.0 [110/782] via 12.1.1.1, 00:01:20, Serial1/1
O        11.1.2.0 [110/782] via 12.1.1.1, 00:01:20, Serial1/1
O        11.1.3.0 [110/782] via 12.1.1.1, 00:01:20, Serial1/1
```

Task 3

Configure OSPF Area 2 based on the following chart; the loopback interfaces should be advertised with their correct mask.

Router	Interface	Area	Router ID
R4	Lo1 F0/0	2 2	0.0.0.4
R5	Lo0 F0/0	2 2	0.0.0.5

On R4:

```
R4(config)#Int range Lo0 - 1
R4(config-if-range)#IP ospf network point-to-point

R4(config)#Router ospf 1
R4(config-router)#router-id 0.0.0.4
R4(config-router)#Netw 44.4.4.4 0.0.0.0 area 2
R4(config-router)#Netw 45.1.1.4 0.0.0.0 area 2
```

On R5:

```
R5(config)#Int lo0
R5(config-if)#ip ospf network point-to-point

R5(config)#Router ospf 1
```


© https://t.me/learningnets

```
R5(config-router)#router-id 0.0.0.5
R5(config-router)#Netw 5.5.5.5 0.0.0.0 area 2
R5(config-router)#Netw 45.1.1.5 0.0.0.0 area 2
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

To verify the configuration:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/2] via 45.1.1.4, 00:00:24, FastEthernet0/0
```

Task 4

Configure OSPF Area 0 based on the following chart; the loopback interfaces should be advertised with their correct mask.

Router	Interface	Area	Router ID
R2	Lo1 F0/0	0 0	0.0.0.2
R3	Lo0 Lo1 F0/0 S1/4	0 0 0 0	0.0.0.3
R4	Lo0 S1/3	0 0	0.0.0.4

On R2:

```
R2(config)#Router ospf 1
R2(config-router)#Netw 22.2.2.2 0.0.0.0 area 0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R2(config-router)#Netw 23.1.1.2 0.0.0.0 area 0
```

On R3:

```
R3(config)#Int range Lo0 - 1
R3(config-if-range)#IP ospf network point-to-point

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 33.3.3.3 0.0.0.0 area 0
R3(config-router)#Netw 34.1.1.3 0.0.0.0 area 0
R3(config-router)#Netw 23.1.1.3 0.0.0.0 area 0
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on FastEthernet0/0 from LOADING to
FULL, Loading Done
```

On R4:

```
R4(config)#Router ospf 1
R4(config-router)#Netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#Netw 34.1.1.4 0.0.0.0 area 0
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial1/3 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

      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/782] via 12.1.1.2, 00:21:47, Serial1/2
      3.0.0.0/24 is subnetted, 1 subnets
O IA     3.3.3.0 [110/783] via 12.1.1.2, 00:02:03, Serial1/2
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/1564] via 12.1.1.2, 00:00:52, Serial1/2
      5.0.0.0/24 is subnetted, 1 subnets
O IA     5.5.5.0 [110/1565] via 12.1.1.2, 00:00:52, Serial1/2
```


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

```
      22.0.0.0/24 is subnetted, 1 subnets
O IA     22.2.2.0 [110/782] via 12.1.1.2, 00:04:19, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/782] via 12.1.1.2, 00:04:19, Serial1/2
      33.0.0.0/24 is subnetted, 1 subnets
O IA     33.3.3.0 [110/783] via 12.1.1.2, 00:02:03, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1563] via 12.1.1.2, 00:01:53, Serial1/2
      44.0.0.0/24 is subnetted, 1 subnets
O IA     44.4.4.0 [110/1564] via 12.1.1.2, 00:00:52, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/1564] via 12.1.1.2, 00:00:52, Serial1/2
```

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
O IA     1.1.1.0 [110/1565] via 45.1.1.4, 00:01:54, FastEthernet0/0
      2.0.0.0/24 is subnetted, 1 subnets
O IA     2.2.2.0 [110/784] via 45.1.1.4, 00:01:54, FastEthernet0/0
      3.0.0.0/24 is subnetted, 1 subnets
O IA     3.3.3.0 [110/783] via 45.1.1.4, 00:01:54, FastEthernet0/0
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/2] via 45.1.1.4, 00:01:54, FastEthernet0/0
      11.0.0.0/24 is subnetted, 4 subnets
O IA     11.1.0.0 [110/1565] via 45.1.1.4, 00:01:54, FastEthernet0/0
O IA     11.1.1.0 [110/1565] via 45.1.1.4, 00:01:54, FastEthernet0/0
O IA     11.1.2.0 [110/1565] via 45.1.1.4, 00:01:54, FastEthernet0/0
O IA     11.1.3.0 [110/1565] via 45.1.1.4, 00:01:54, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/1564] via 45.1.1.4, 00:01:54, FastEthernet0/0
      22.0.0.0/24 is subnetted, 1 subnets
O IA     22.2.2.0 [110/784] via 45.1.1.4, 00:01:54, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/783] via 45.1.1.4, 00:01:54, FastEthernet0/0
      33.0.0.0/24 is subnetted, 1 subnets
O IA     33.3.3.0 [110/783] via 45.1.1.4, 00:01:54, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 45.1.1.4, 00:01:54, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/2] via 45.1.1.4, 00:09:21, FastEthernet0/0
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

Task 5

Configure area 1 such that it does not receive LSA types 4 or LSA 5s, but the routers in this area should still maintain inter-area routes in their routing table. These routers should have reachability to the existing and future external routes redistributed into this routing domain.

We can configure OSPF Stub area to accomplish this task.

Important points to understand about a STUB area:

- A Stub area cannot be a transit area for Virtual links; GRE tunnels should be used instead.
- A Stub area cannot have an ASBR.
- The back bone area cannot be configured as a Stub area.
- Every router and the ABR of that area should be configured with "area xx stub" command to match the area stub flag.
- External routes are NOT allowed in a Stub area, but the routers in the Stub area can connect to the External routes via the default route that is injected by their ABR.
- A stub area can NOT have LSA type 4s.
- By default, the cost of the default route injected into this area by the ABR is 1; this can be verified using the "Show ip ospf", and/or "Show ip route". The cost of the default route can be changed using "area xx default-cost cc", where xx is the area number, and cc is the desired cost

On R1:

```
R1(config)#router ospf 1
R1(config-router)#area 1 stub
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Serial1/2 from FULL to DOWN,
Neighbor Down: Adjacency forced to reset
```

NOTE: The adjacency with R2 (0.0.0.2) transitioned from "Full" to "Down" because the area stub flag no longer matches.

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

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 stub
```

You should see the following console message stating that the adjacency is in FULL state:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R1:

```
R1#Show ip ospf | Inc area

 Supports area transit capability
 Number of areas in this router is 1. 0 normal 1 stub 0 nssa
 Number of areas transit capable is 0
        Number of interfaces in this area is 6
        It is a stub area

R1#Show ip route ospf | B Gate
Gateway of last resort is 12.1.1.2 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/782] via 12.1.1.2, 00:08:47, Serial1/2
      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/782] via 12.1.1.2, 00:08:47, Serial1/2
      3.0.0.0/24 is subnetted, 1 subnets
O IA     3.3.3.0 [110/783] via 12.1.1.2, 00:08:47, Serial1/2
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/1564] via 12.1.1.2, 00:08:47, Serial1/2
      5.0.0.0/24 is subnetted, 1 subnets
O IA     5.5.5.0 [110/1565] via 12.1.1.2, 00:08:47, Serial1/2
      22.0.0.0/24 is subnetted, 1 subnets
O IA     22.2.2.0 [110/782] via 12.1.1.2, 00:08:47, Serial1/2
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/782] via 12.1.1.2, 00:08:47, Serial1/2
      33.0.0.0/24 is subnetted, 1 subnets
O IA     33.3.3.0 [110/783] via 12.1.1.2, 00:08:47, Serial1/2
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/1563] via 12.1.1.2, 00:08:47, Serial1/2
      44.0.0.0/24 is subnetted, 1 subnets
O IA     44.4.4.0 [110/1564] via 12.1.1.2, 00:08:47, Serial1/2
      45.0.0.0/24 is subnetted, 1 subnets
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
O IA     45.1.1.0 [110/1564] via 12.1.1.2, 00:08:47, Serial1/2
```

The output of the above show command reveals that the cost of the default route from R1's perspective is 782.

Why 782?
The default cost of the default route which is 1 PLUS the cost of the Serial link which is 781. Let's verify:

On R2:

```
R2#Show ip ospf | Inc cost

          generates stub default route with cost 1

R2#Show ip ospf interface S1/1 | Inc Cost

  Process ID 1, Router ID 0.0.0.1, Network Type POINT_TO_POINT, Cost: 781
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

Therefore, if 1 is added to 781, you should see the resulting cost of 782:

On R1:

```
R1#Sh ip ospf route | Inc 0.0.0.0

*>  0.0.0.0/0, Inter, cost 782, area 1
```

The routers in the other areas will see all the routes from area 1:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
O IA     1.1.1.0 [110/1565] via 45.1.1.4, 00:16:58, FastEthernet0/0
      2.0.0.0/24 is subnetted, 1 subnets
O IA     2.2.2.0 [110/784] via 45.1.1.4, 00:27:29, FastEthernet0/0
      3.0.0.0/24 is subnetted, 1 subnets
O IA     3.3.3.0 [110/783] via 45.1.1.4, 00:27:29, FastEthernet0/0
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/2] via 45.1.1.4, 00:27:29, FastEthernet0/0
      11.0.0.0/24 is subnetted, 4 subnets
O IA     11.1.0.0 [110/1565] via 45.1.1.4, 00:16:58, FastEthernet0/0
```


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

```
O IA     11.1.1.0 [110/1565] via 45.1.1.4, 00:16:58, FastEthernet0/0
O IA     11.1.2.0 [110/1565] via 45.1.1.4, 00:16:58, FastEthernet0/0
O IA     11.1.3.0 [110/1565] via 45.1.1.4, 00:16:58, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/1564] via 45.1.1.4, 00:27:29, FastEthernet0/0
      22.0.0.0/24 is subnetted, 1 subnets
O IA     22.2.2.0 [110/784] via 45.1.1.4, 00:27:29, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/783] via 45.1.1.4, 00:27:29, FastEthernet0/0
      33.0.0.0/24 is subnetted, 1 subnets
O IA     33.3.3.0 [110/783] via 45.1.1.4, 00:27:29, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 45.1.1.4, 00:27:29, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/2] via 45.1.1.4, 00:34:56, FastEthernet0/0
```

Task 6

Area 2 should not receive LSAs types 3 (Coming from other areas), 4 or 5. These routers should have reachability to the existing and future external routes redistributed into this routing domain.

On R5:

You can reduce the size of the routing table further by configuring an area as totally stubby. Since all the "IA" and "E" (Inter-area and External) routes can be reached via a default route which is injected by the ABR of this area, there is no reason to maintain these routes.

```
R5(config)#router ospf 1
R5(config-router)#area 2 stub
```

You should see the following console message stating that the adjacency transitioned from FULL state to DOWN state, this happened because the Area Stub Flag no longer matches, remember that if the Area Stub Flag does NOT match, the adjacency can NOT be established. Therefore, the above command must be configured on all the routers within area 2.

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.4 on FastEthernet0/0 from FULL to
DOWN, Neighbor Down: Adjacency forced to reset
```

On R4:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R4(config)#router ospf 1
R4(config-router)#Area 2 stub No-summary
```

The above command must only be configured on the ABR of this area.

To verify the configuration:

On R5:

```
R5#Show ip route ospf | B Gate
Gateway of last resort is 45.1.1.4 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/2] via 45.1.1.4, 00:00:14, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/2] via 45.1.1.4, 00:00:14, FastEthernet0/0
```

Important points to understand about a Totally Stub area:

- **A Totally Stub area cannot be a transit area for Virtual links; GRE tunnels should be used instead.**
- **A Totally Stub area cannot have an ASBR.**
- **The back bone area cannot be configured as a Totally Stub area.**
- **The routers within the area should be configured with "Area xx stub" command, whereas, the ABR of the Totally Stub area MUST be configured with "Area xx Stub No-summary".**
- **External routes are NOT allowed in a Totally Stub area, but the routers in the Totally Stub area can connect to the External routes via the default route that is injected in the area by their ABR.**
- **By default, the cost of the default route injected into this area by the ABR is 1; this can be verified using the "Show ip ospf" on the ABR, and/or "Show ip route" on the other routers in the area. The cost of the default route can be changed using "Area xx default-cost cc", where xx is the area number, and cc is the desired cost**
- **The routers in a Totally Stub area don't get inter-area routes, but they have NLRI to the inter-area routes via the default route that is injected by the ABR**

Let's see why the default route is the ONLY inter-area route injected into this area:

```
R4#Show ip ospf da summary 0.0.0.0

            OSPF Router with ID (0.0.0.4) (Process ID 1)

                Summary Net Link States (Area 2)
```


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

```
LS age: 331
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 0.0.0.0 (summary Network Number)
Advertising Router: 0.0.0.4
LS Seq Number: 80000001
Checksum: 0x99A0
Length: 28
Network Mask: /0
      MTID: 0         Metric: 1
```

We can see that the default route is injected by the ABR, in this case R4 (0.0.0.4), and it was generated within area 2, therefore, it is wrong to say that the default route is the ONLY inter-area route injected into totally stubby area.
The inter-area routes that have originated in the other areas are the only routes that are blocked and since the default route is originated in area 2, it is not blocked.

Task 7

Redistribute Lo0 – 4 interfaces of R1 into this routing domain. After the redistribution, area 1 should only receive and propagate LSA types 1, 2, 3 and 7. This area should not have the ability to connect to any external routes redistributed else where within this routing domain.

Since a stub area can NOT have external routes, area 1 must be converted into an NSSA.

Let's remove the following commands from R1:

On R1:

```
R1(config)#router ospf 1
R1(config-router)#No Netw 1.1.1.1 0.0.0.0 area 1
R1(config-router)#No Netw 11.1.0.1 0.0.0.0 area 1
R1(config-router)#No Netw 11.1.1.1 0.0.0.0 area 1
R1(config-router)#No Netw 11.1.2.1 0.0.0.0 area 1
R1(config-router)#No Netw 11.1.3.1 0.0.0.0 area 1

R1(config-router)#No area 1 stub
R1(config-router)#area 1 nssa
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.2 on Serial1/2 from DOWN to DOWN,
Neighbor Down: Adjacency forced to reset
```

```
R1(config)#route-map TST permit 10
R1(config-route-map)#match interface Lo0 Lo1 Lo2 Lo3 Lo4

R1(config)#router ospf 1
R1(config-router)#redistribute connected route-map TST subnets
```

On R2:

```
R2(config)#router ospf 1
R2(config-router)#No area 1 stub
R2(config-router)#area 1 nssa
```

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 1, Nbr 0.0.0.1 on Serial1/1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R2:

NOTE: R1 does NOT have a default route, because the ABR (R2) did NOT inject one:

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

     1.0.0.0/24 is subnetted, 1 subnets
O N2    1.1.1.0 [110/20] via 12.1.1.1, 00:01:53, Serial1/1
     3.0.0.0/24 is subnetted, 1 subnets
O       3.3.3.0 [110/2] via 23.1.1.3, 00:02:03, FastEthernet0/0
     4.0.0.0/24 is subnetted, 1 subnets
O       4.4.4.0 [110/783] via 23.1.1.3, 00:02:03, FastEthernet0/0
     5.0.0.0/24 is subnetted, 1 subnets
O IA    5.5.5.0 [110/784] via 23.1.1.3, 00:02:03, FastEthernet0/0
     11.0.0.0/24 is subnetted, 4 subnets
O N2    11.1.0.0 [110/20] via 12.1.1.1, 00:01:53, Serial1/1
O N2    11.1.1.0 [110/20] via 12.1.1.1, 00:01:53, Serial1/1
O N2    11.1.2.0 [110/20] via 12.1.1.1, 00:01:53, Serial1/1
O N2    11.1.3.0 [110/20] via 12.1.1.1, 00:01:53, Serial1/1
```


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

```
      33.0.0.0/24 is subnetted, 1 subnets
O        33.3.3.0 [110/2] via 23.1.1.3, 00:02:03, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/782] via 23.1.1.3, 00:02:03, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O IA     44.4.4.0 [110/783] via 23.1.1.3, 00:02:03, FastEthernet0/0
      45.0.0.0/24 is subnetted, 1 subnets
O IA     45.1.1.0 [110/783] via 23.1.1.3, 00:02:03, FastEthernet0/0
```

On R4:

NOTE: R2 received the external routes from area 1 as "N2" and generated an "E2" for every "N2".

```
R4#Show ip route ospf | B Gate
Gateway of last resort is not set

      1.0.0.0/24 is subnetted, 1 subnets
O E2     1.1.1.0 [110/20] via 34.1.1.3, 00:03:29, Serial1/3
      2.0.0.0/24 is subnetted, 1 subnets
O IA     2.2.2.0 [110/783] via 34.1.1.3, 00:22:15, Serial1/3
      3.0.0.0/24 is subnetted, 1 subnets
O        3.3.3.0 [110/782] via 34.1.1.3, 00:22:15, Serial1/3
      5.0.0.0/24 is subnetted, 1 subnets
O        5.5.5.0 [110/2] via 45.1.1.5, 00:22:03, FastEthernet0/0
      11.0.0.0/24 is subnetted, 4 subnets
O E2     11.1.0.0 [110/20] via 34.1.1.3, 00:03:29, Serial1/3
O E2     11.1.1.0 [110/20] via 34.1.1.3, 00:03:29, Serial1/3
O E2     11.1.2.0 [110/20] via 34.1.1.3, 00:03:29, Serial1/3
O E2     11.1.3.0 [110/20] via 34.1.1.3, 00:03:29, Serial1/3
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/1563] via 34.1.1.3, 00:22:15, Serial1/3
      22.0.0.0/24 is subnetted, 1 subnets
O        22.2.2.0 [110/783] via 34.1.1.3, 00:22:15, Serial1/3
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/782] via 34.1.1.3, 00:22:15, Serial1/3
      33.0.0.0/24 is subnetted, 1 subnets
O        33.3.3.0 [110/782] via 34.1.1.3, 00:22:15, Serial1/3
```

In any Stub area the ABR receives the Type-7 LSAs and converts them into Type-5. The following can confirm this fact:

```
R2#Show ip ospf database nssa-external | Inc Link State

                Type-7 AS External Link States (Area 1)
  Link State ID: 1.1.1.0 (External Network Number )
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
  Link State ID: 11.1.0.0 (External Network Number )
  Link State ID: 11.1.1.0 (External Network Number )
  Link State ID: 11.1.2.0 (External Network Number )
  Link State ID: 11.1.3.0 (External Network Number )

R2#Show ip ospf database external | Inc Link State

                Type-5 AS External Link States
  Link State ID: 1.1.1.0 (External Network Number )
  Link State ID: 11.1.0.0 (External Network Number )
  Link State ID: 11.1.1.0 (External Network Number )
  Link State ID: 11.1.2.0 (External Network Number )
  Link State ID: 11.1.3.0 (External Network Number )
```

Task 8

Configure the following Loopback interfaces on R5:

Lo1 – 55.5.1.5 /24
Lo2 – 55.5.2.5 /24
Lo3 – 55.5.3.5 /24
Lo4 – 55.5.4.5 /24

On R5:

```
R5(config-if)#int Lo1
R5(config-if)#IP address 55.5.1.5 255.255.255.0

R5(config-if)#int Lo2
R5(config-if)#IP address 55.5.2.5 255.255.255.0

R5(config-if)#int Lo3
R5(config-if)#IP address 55.5.3.5 255.255.255.0

R5(config-if)#int Lo4
R5(config-if)#IP address 55.5.4.5 255.255.255.0
```

To verify the configuration:

On R5:

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

```
R5#Show ip int br | exc unass

Interface              IP-Address      OK? Method Status                Protocol
FastEthernet0/0        45.1.1.5        YES manual up                    up
Loopback0              5.5.5.5         YES manual up                    up
Loopback1              55.5.1.5        YES manual up                    up
Loopback2              55.5.2.5        YES manual up                    up
Loopback3              55.5.3.5        YES manual up                    up
Loopback4              55.5.4.5        YES manual up                    up
```

Task 9

Redistribute Lo1 – 4 interfaces on R5. Configure R5 such that the routers in this area only maintain and propagate LSA types 1, 2, 3, 7 and a default route.

Since area 2 is a Totally stubby area, it cannot have external routes unless it's converted into a NSSA, since the task states that the routers of this area should also get a default route, the ABR of this area (R4) is configured to inject a default route using the "Default-information-originate" keyword:

On R5:

```
R5(config-if)#route-map TST permit 10
R5(config-route-map)#match interface Lo1 Lo2 Lo3 Lo4

R5(config)#router ospf 1
R5(config-router)#No area 2 stub
R5(config-router)#area 2 nssa
R5(config-router)#redistribute connected subnets route-map TST
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#No area 2 stub
R4(config-router)#area 2 nssa default-information-originate
```

The "Default-information-originate" keyword at the end of the "area 2 nssa" command will inject a default route into the area. This default route will be injected as "N2" and NOT "IA".

To verify the configuration:

On R5:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
R5#Show ip route ospf | B Gate
Gateway of last resort is 45.1.1.4 to network 0.0.0.0

O*N2  0.0.0.0/0 [110/1] via 45.1.1.4, 00:02:09, FastEthernet0/0
      2.0.0.0/24 is subnetted, 1 subnets
O IA     2.2.2.0 [110/784] via 45.1.1.4, 00:02:09, FastEthernet0/0
      3.0.0.0/24 is subnetted, 1 subnets
O IA     3.3.3.0 [110/783] via 45.1.1.4, 00:02:09, FastEthernet0/0
      4.0.0.0/24 is subnetted, 1 subnets
O IA     4.4.4.0 [110/2] via 45.1.1.4, 00:02:09, FastEthernet0/0
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/1564] via 45.1.1.4, 00:02:09, FastEthernet0/0
      22.0.0.0/24 is subnetted, 1 subnets
O IA     22.2.2.0 [110/784] via 45.1.1.4, 00:02:09, FastEthernet0/0
      23.0.0.0/24 is subnetted, 1 subnets
O IA     23.1.1.0 [110/783] via 45.1.1.4, 00:02:09, FastEthernet0/0
      33.0.0.0/24 is subnetted, 1 subnets
O IA     33.3.3.0 [110/783] via 45.1.1.4, 00:02:09, FastEthernet0/0
      34.0.0.0/24 is subnetted, 1 subnets
O IA     34.1.1.0 [110/782] via 45.1.1.4, 00:02:09, FastEthernet0/0
      44.0.0.0/24 is subnetted, 1 subnets
O        44.4.4.0 [110/2] via 45.1.1.4, 00:02:09, FastEthernet0/0
```

Task 10

Area 1 should be changed such that it receives and propagates LSA types 1, 2, 7 plus a default route. This area should NOT maintain Inter-area routes, but must have the ability to connect to these routes.

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 nssa no-summary
```

The "no-summary" keyword blocks the summary LSAs which are the inter-area routes. Whenever the inter-area routes are blocked using the "No-Summary" keyword, the ABR will always inject a default route.

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 12.1.1.2 to network 0.0.0.0

O*IA  0.0.0.0/0 [110/782] via 12.1.1.2, 00:03:14, Serial1/2
      2.0.0.0/24 is subnetted, 1 subnets
O        2.2.2.0 [110/782] via 12.1.1.2, 00:26:23, Serial1/2
```

Task 11

The default route that was injected into area 1 should have a cost of 50.

To see the cost of the default route that was injected:

On R1:

```
R1#Show ip route ospf | Inc 0.0.0.0/0

O*IA  0.0.0.0/0 [110/782] via 12.1.1.2, 00:04:17, Serial1/2

R1#Sh ip ospf route | B Area
    Area 1

    Intra-area Route List
*>  2.2.2.0/24, Intra, cost 782, area 1
      via 12.1.1.2, Serial1/2
*   12.1.1.0/24, Intra, cost 781, area 1, Connected
      via 12.1.1.1, Serial1/2

    Intra-area Router Path List
i 0.0.0.2 [781] via 12.1.1.2, Serial1/2, ABR/ASBR, Area 1, SPF 17

    Inter-area Route List
*>  0.0.0.0/0, Inter, cost 782, area 1
      via 12.1.1.2, Serial1/2

R1#Show ip ospf inter s1/2 | Inc Cost

  Process ID 1, Router ID 0.0.0.1, Network Type POINT_TO_POINT, Cost: 781
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
```

We can see that the cost of the S1/2 interface is 781, and the total cost of the default route injected is 782, therefore, the cost of the defaulf route is 1. Let's change it to 50 and verify:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

On R2:

```
R2(config)#router ospf 1
R2(config-router)#area 1 default-cost 50
```

By default, the cost of the default route that is injected into a given area is 1; this can be verified by "Show ip route" command; the output of the above show command reveals that the cost of the injected default route is 782; this is the cost of the link to the ABR (The cost of the S1/2 interface) plus 1 (The default cost of the default route which is 1).

The default cost of the injected default route can be changed using the "Area xx default-cost cc", where cc is the new cost replacing the default value.

To verify the configuration:

On R1:

```
R1#Sh ip ospf route | B Inter-area

    Inter-area Route List
*>  0.0.0.0/0, Inter, cost 831, area 1
      via 12.1.1.2, Serial1/2
```

If the OSPF cost of the S1/2 interface is subtracted from 831, we should see the cost of the default route injected into this area:

831 – 781 = 50

```
R1#Show ip route ospf | Inc 0.0.0.0/0

O*IA  0.0.0.0/0 [110/831] via 12.1.1.2, 00:03:18, Serial1/2
```

The database is the best place to go to to find the exact value:

```
R1#Show ip ospf da summary

            OSPF Router with ID (0.0.0.1) (Process ID 1)

                Summary Net Link States (Area 1)

  Routing Bit Set on this LSA in topology Base with MTID 0

  LS age: 242
  Options: (No TOS-capability, DC, Upward)
```


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

```
LS Type: Summary Links(Network)
Link State ID: 0.0.0.0 (summary Network Number)
Advertising Router: 0.0.0.2
LS Seq Number: 80000002
Checksum: 0x17EA
Length: 28
Network Mask: /0
      MTID: 0         Metric: 50
```

Task 12

Erase the startup config and reload the routers before proceeding to the next lab.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

CCIE Foundation

5.0

www.MicronicsTraining.com

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

Redistribution

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

Lab 1 – Basics of Redistribution

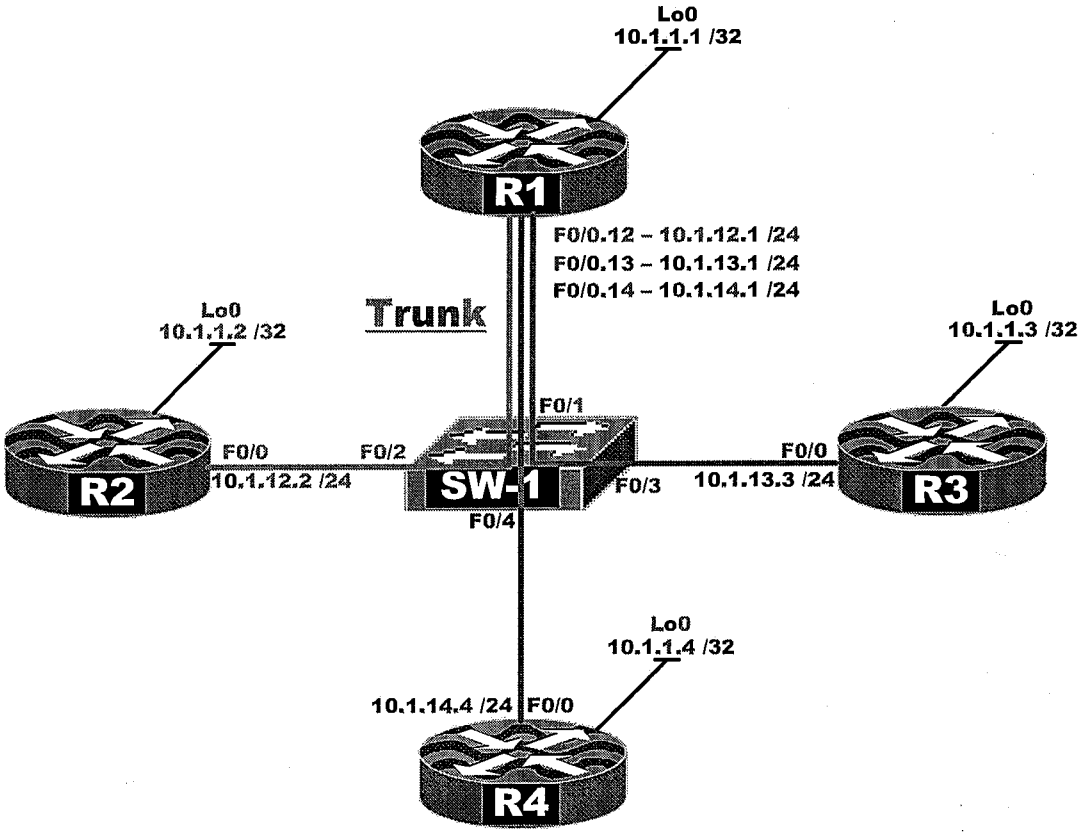

Task 1

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
SW1(config-if)#No shut
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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
R3(config-if)#IP address 10.1.13.3 255.255.255.0
R3(config-if)#No shut
```


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

```
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
```

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.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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
O IA     10.1.1.3/32 [110/2] via 10.1.13.3, 00:01:33, FastEthernet0/0.13
```


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

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:

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 networks.

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

      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
```


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

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:

```
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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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.

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:

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

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

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

     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:

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

```
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
```

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:

© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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:

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
```


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

```
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

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

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

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.

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

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.

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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

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:

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

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

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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

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
```


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

```
R     192.168.23.0/24 [120/1] via 10.1.12.2, 00:00:24, 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:

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
```


© 2014 Narbik Kocharians. All rights reserved

https://t.me/learningnets

```
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:

```
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
```


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

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

https://t.me/learningnets