

CCNP

CISCO CERTIFIED NETWORK PROFESSIONAL

LAB MANUAL

VER 2.0

Page 1 of 315

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com

All contents are copyright @ 2007-2010

All rights reserved.

ACKNOWLEDGEMENT

We can write a 1000 page book, but we can't find enough words to describe the credit Mr. Siddiq Ahmed deserves for what is good about this book. Your ability to guide us made the toughest things easy in developing this book. We still refuse to write a book unless you develop it!!

Behind the scenes at Netmetrics is a vast array of talented people – all of whom made it possible to develop the book. People, who provided the right equipment in the right time, people who maintained a friendly atmosphere to finish the work in time. Many thanks to you all for the hard and good work.

Thanks to Nizam, Anwar, Afsar, Awadh, Sreenivas and Mr. Prasad for finding better ways to describe how the technical pieces fit together, for fixing errors.

Thanks to Mr. Abdur Rahman for giving us the time we needed to work on this book and by making sure the job was fun.

Finally, no acknowledgement section could be complete without acknowledging ourselves, who took all the responsibility and dedication in completing the book.

Finally, to GOD, who gives us strength when things are tough, and peace beyond belief – THANK YOU!!!

PAPER 1

Routing

BUILDING SCALABLE CISCO INTERNETWORKS


BSCI (642–901)

Module 1 – EIGRP

EIGRP LABS INDEX

- 1. CONFIGURING BASIC EIGRP**
- 2. CONFIGURING IP DEFAULT-NETWORK COMMAND**
- 3. CONFIGURE ROUTE SUMMARIZATION**
- 4. LOAD BALANCING ACROSS EQUAL COST PATH**
- 5. LOAD BALANCING ACROSS UNEQUAL COST PATH**
- 6. CONFIGURE EIGRP AUTHENTICATION (MD5)**
- 7. CONFIGURE EIGRP STUB**
- 8. EIGRP REDISTRIBUTION WITH RIPv2**
- 9. EIGRP REDISTRIBUTION WITH OSPF**
- 10. CONFIGURE EIGRP WITH REDISTRIBUTE CONNECTED.**
- 11. CONFIGURE EIGRP AND IGRP**

Lab 1 – Basic EIGRP Configuration



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure EIGRP on 2 routers in AS 100. Disable Auto-summary.

R1	R2
Router eigrp 100	Router eigrp 100
Network 1.0.0.0	Network 1.0.0.0
Network 10.0.0.0	Network 20.0.0.0
No auto-summary	No auto-summary

Verification :

R1#show ip route


C 1.0.0.0/8 is directly connected, Serial0/2/0
D 20.0.0.0/8 [90/2195456] via 1.1.1.2, 00:43:52, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0

R1#show ip eigrp neighbors

IP-EIGRP neighbors for process 100

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq Num
0	1.1.1.2	Se0/2/0	13 00:45:08	355	2130	0	106

Lab 2 – Configuring ip default-network Command



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S0	2.2.2.1	255.0.0.0
E 0	30.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0/2	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure EIGRP according to the above scenario. Configure R1 (S0, E0), R2 (S1, S0, E0) in EIGRP AS 100 and R3 (S0/2) in EIGRP AS 200. Do not advertise network 30.0.0.0 in EIGRP process. R1 wants to send packets to network 30.0.0.0. Use the Ip default-network command to accomplish this task. Also disable auto-summary.

R1 Router eigrp 100 Network 10.0.0.0 Network 1.0.0.0 No auto-summary	R2 Router eigrp 100 Network 1.0.0.0 Network 20.0.0.0 Network 2.0.0.0 No auto-summary Ip route 30.0.0.0 255.0.0.0 2.2.2.2 Ip default-network 2.0.0.0
R3 Router eigrp 100 Network 2.0.0.0 No auto-summary.	

Verification :

R1#show ip route

Gateway of last resort is 1.1.1.2 to network 2.0.0.0

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
D* 2.0.0.0/8 [90/2681856] via 1.1.1.2, 00:00:14, Serial0/2/0
D 20.0.0.0/8 [90/2195456] via 1.1.1.2, 00:04:43, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```


The output displays network 2.0.0.0 as a D* route in the routing table as this is candidate default-route established in R1 to reach network 30.0.0.0.

Note: When we ping from R1 to 30.1.1.1 network

R1 # ping 30.1.1.1

Result: 100% success

Lab 3 – Route Summarization with EIGRP



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S0	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0/2	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 1	172.168.0.1	255.255.255.0
Loopback 2	172.168.1.1	255.255.255.0
Loopback 3	172.168.2.1	255.255.255.0
Loopback 4	172.168.3.1	255.255.255.0
Loopback 5	172.168.4.1	255.255.255.0
Loopback 6	172.168.5.1	255.255.255.0
Loopback 7	172.168.6.1	255.255.255.0
Loopback 8	172.168.7.1	255.255.255.0

Lab Objective:

Task 1

Configure the following Loopback Interfaces on R3 and advertise them under EIGRP:

Loopback 1: 172.168.0.1/24
Loopback 2: 172.168.1.1/24
Loopback 3: 172.168.2.1/24
Loopback 4: 172.168.3.1/24
Loopback 5: 172.168.4.1/24
Loopback 6: 172.168.5.1/24
Loopback 7: 172.168.6.1/24
Loopback 8: 172.168.7.1/24

R3

```
Interface loopback 1
Ip address 172.168.0.1 255.255.255.0
Interface loopback 2
Ip address 172.168.1.1 255.255.255.0
Interface loopback 3
Ip address 172.168.2.1 255.255.255.0
Interface loopback 4
Ip address 172.168.3.1 255.255.255.0
Interface loopback 5
Ip address 172.168.4.1 255.255.255.0
Interface loopback 6
Ip address 172.168.5.1 255.255.255.0
Interface loopback 7
Ip address 172.168.6.1 255.255.255.0
Interface loopback 8
Ip address 172.168.7.1 255.255.255.0

Router eigrp 100
Network 2.0.0.0
Network 30.0.0.0
Network 172.168.1.0 0.0.0.255
Network 172.168.2.0 0.0.0.255
Network 172.168.3.0 0.0.0.255
Network 172.168.4.0 0.0.0.255
Network 172.168.5.0 0.0.0.255
Network 172.168.6.0 0.0.0.255
Network 172.168.7.0 0.0.0.255
Network 172.168.0.0 0.0.0.255
No auto-summary
```

Task 2

Configure EIGRP on R1 and R2. Advertise the directly connected networks in EIGRP in AS 100. Disable auto-summary. Also configure route summarization so that only one summary route is advertised to R1.

R1	R2
<pre>Router eigrp 100 Network 10.0.0.0 Network 1.0.0.0 No auto-summary</pre>	<pre>Router eigrp 100 Network 1.0.0.0 Network 20.0.0.0 Network 2.0.0.0 No auto-summary Int s0 Ip summary-address eigrp 100 172.168.0.0 255.255.248.0</pre>

Verification:

R1#show ip route


```
C 1.0.0.0/8 is directly connected, Serial0/2/0
D 2.0.0.0/8 [90/2681856] via 1.1.1.2, 00:00:02, Serial0/2/0
D 20.0.0.0/8 [90/2195456] via 1.1.1.2, 00:00:02, Serial0/2/0
D 172.168.0.0/21 is subnetted, 1 subnets
D 172.168.0.0 [90/2809856] via 1.1.1.2, 00:00:02, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
D 30.0.0.0/8 [90/2707456] via 1.1.1.2, 00:00:02, Serial0/2/0
```

With route summarization on R2 a summary route is created pointing to null 0

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
D 172.168.0.0/16 is variably subnetted, 9 subnets, 2 masks
D 172.168.4.0/24 [90/2297856] via 2.2.2.2, 00:07:13, Serial1
D 172.168.5.0/24 [90/2297856] via 2.2.2.2, 00:07:08, Serial1
D 172.168.6.0/24 [90/2297856] via 2.2.2.2, 00:07:04, Serial1
D 172.168.7.0/24 [90/2297856] via 2.2.2.2, 00:06:56, Serial1
D 172.168.0.0/24 [90/2297856] via 2.2.2.2, 00:06:49, Serial1
D 172.168.0.0/21 is a summary, 00:01:24, Null0
D 172.168.1.0/24 [90/2297856] via 2.2.2.2, 00:07:33, Serial1
D 172.168.2.0/24 [90/2297856] via 2.2.2.2, 00:07:25, Serial1
D 172.168.3.0/24 [90/2297856] via 2.2.2.2, 00:07:18, Serial1
D 10.0.0.0/8 [90/2172416] via 1.1.1.1, 00:01:30, Serial0
D 30.0.0.0/8 [90/2195456] via 2.2.2.2, 00:08:03, Serial1
```

Lab 4 – Load balancing across Equal Cost Path



Interface IP Address Configuration

R2

Interface	IP Address	Subnet Mask
S 0	2.2.2.1	255.0.0.0
S 1	1.1.1.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure EIGRP AS 100 as per the above scenario and verify load balancing using the traceroute command on R3 :

R2	R3
Router eigrp 100 Network 20.0.0.0	Router eigrp 100 Network 1.0.0.0

Network 1.0.0.0 Network 2.0.0.0 No auto-summary	Network 30.0.0.0 Network 2.0.0.0 No auto-summary
---	--

Verification:

R3#show ip route

```
C 1.0.0.0/8 is directly connected, Serial1
C 2.0.0.0/8 is directly connected, Serial0
D 20.0.0.0/8 [90/2195456] via 1.1.1.1, 00:07:42, Serial1
  [90/2195456] via 2.2.2.1, 00:07:42, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```

First Traceroute packet going via 1.1.1.1

R3#traceroute 20.1.1.1

Type escape sequence to abort.
Tracing the route to 20.1.1.1

```
1 1.1.1.1 32 msec
  2.2.2.1 20 msec *
```

Second Traceroute packet going via 2.2.2.1

R3#traceroute 20.1.1.1

Type escape sequence to abort.
Tracing the route to 20.1.1.1

```
1 2.2.2.1 20 msec
  1.1.1.1 28 msec *
```

Lab 5 – Load balancing across Unequal Cost Path

(Scenario Based On Lab 4)
Interface IP Address Configuration

R2

Interface	IP Address	Subnet Mask
S 0	2.2.2.1	255.0.0.0
S 1	1.1.1.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure EIGRP AS 100 as per the above scenario. Make the links unequal cost paths using the bandwidth command in interface mode and verify load balancing.

Use the variance command to gain load balancing

R1 Router eigrp 100 Network 20.0.0.0 Network 1.0.0.0 Network 2.0.0.0 No auto-summary	R2 Router eigrp 100 Network 1.0.0.0 Network 30.0.0.0 Network 2.0.0.0 Variance 2 No auto-summary Interface S 0 Bandwidth 800
--	--

The variance multiplier set in the variance command when multiplied by the successor FD, must be greater than the feasible successor FD. Thus the feasible successors whose FD is less than the above calculated value are installed in the routing table.

Verification:

With out the variance command:

R2#sh ip eigrp topology

```
P 1.0.0.0/8, 1 successors, FD is 2169856
  via Connected, Serial1
P 2.0.0.0/8, 1 successors, FD is 3712000
  via Connected, Serial0
  via 1.1.1.2 (2681856/2169856), Serial1
P 20.0.0.0/8, 1 successors, FD is 281600
  via Connected, Ethernet0
P 30.0.0.0/8, 1 successors, FD is 2195456
  via 1.1.1.2 (2195456/281600), Serial1
  via 2.2.2.2 (3737600/281600), Serial0
```

The output displays 2 routes installed in the topology table with 2 different costs.

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial1
C 2.0.0.0/8 is directly connected, Serial0
C 20.0.0.0/8 is directly connected, Ethernet0
D 30.0.0.0/8 [90/2195456] via 1.1.1.2, 00:01:42, Serial1
```


With the variance command:

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial1
C 2.0.0.0/8 is directly connected, Serial0
C 20.0.0.0/8 is directly connected, Ethernet0
D 30.0.0.0/8 [90/2195456] via 1.1.1.2, 00:00:04, Serial1
  [90/3737600] via 2.2.2.2, 00:00:04, Serial0
```

The output displays 2 routes installed in the routing table.

Lab 6 – EIGRP Authentication



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure MD5 authentication for the links. Use cisco123 as the key-string with a key-id of 1.

R1	R2
<pre> Int S1 Ip authentication mode eigrp 100 md5 Ip authentication key-chain eigrp 100 chain1 Key chain chain1 Key 1 Key-string cisco123 </pre>	<pre> Int S 0 Ip authentication mode eigrp 100 md5 Ip authentication key-chain eigrp 100 chain 2 Key chain chain 2 Key 1 Key-string cisco123 </pre>

Verification

With EIGRP Authentication:

R2#show ip eigrp neighbors

IP-EIGRP neighbors for process 100

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq Num
0	2.2.2.2	Se1	14 00:00:24	40	240	0	2

Verify authentication by using debug EIGRP packets

R2#debug eigrp packets

*Mar 1 02:52:50.895: EIGRP: Sending HELLO on Ethernet0

*Mar 1 02:52:50.899: AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iibQ un/rely 0/0

*Mar 1 02:52:53.219: EIGRP: received packet with MD5 authentication, key id = 1

*Mar 1 02:52:53.223: EIGRP: Received HELLO on Serial1 nbr 2.2.2.2

*Mar 1 02:52:53.223: AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iibQ un/rely 0/0 p

With authentication mismatch:

R2#show ip eigrp neighbors

IP-EIGRP neighbors for process 100

-----NIL-----

R2#debug eigrp packets

*Mar 1 02:58:05.895: EIGRP: Sending HELLO on Serial1

*Mar 1 02:58:05.895: AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iibQ un/rely 0/0


*Mar 1 02:58:06.347: EIGRP: Sending HELLO on Ethernet0

*Mar 1 02:58:06.351: AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iibQ un/rely 0/0

*Mar 1 02:58:08.471: EIGRP: pkt key id = 1, authentication mismatch

*Mar 1 02:58:08.475: EIGRP: Serial1: ignored packet from 2.2.2.2, opcode = 5 (invalid authentication)

Lab 7 – Configuring EIGRP STUB



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	2.2.2.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.2	255.0.0.0
S1	3.3.3.1	255.0.0.0
E 0/0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1/0	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	172.168.0.1	255.255.255.0
Loopback 1	172.168.1.1	255.255.255.0
Loopback 2	172.168.2.1	255.255.255.0
Loopback 3	172.168.3.1	255.255.255.0

Lab Objective:

Task 1

Configure EIGRP AS 100 as per the above scenario on R1, R2 and R3. Disable auto-summary. Only one summary route must be advertised to R2 and R1

R1 Router eigrp 100 Network 2.0.0.0 Network 10.0.0.0 No auto-summary	R2 Router eigrp 100 Network 2.0.0.0 Network 3.0.0.0 Network 20.0.0.0 No auto-summary
R3 Router eigrp 100 Network 3.0.0.0 Network 30.0.0.0 Network 172.168.0.0 No auto-summary Interface s 1/0 Ip summary-address eigrp 100 172.168.0.0 255.255.252.0	

Verification:

Without configuring stub in R3:

R2#show ip route

- C 2.0.0.0/8 is directly connected, Serial1/0
- C 3.0.0.0/8 is directly connected, Serial1/1
- C 20.0.0.0/8 is directly connected, Ethernet0/0
- 172.168.0.0/22 is subnetted, 1 subnets
- D 172.168.0.0 [90/20640000] via 3.3.3.2, Serial1/1
- D 10.0.0.0/8 [90/20537600] via 2.2.2.1, Serial1/0
- D 30.0.0.0/8 [90/20537600] via 3.3.3.2, Serial1/1

The output displays directly connected routes, summary route and Eigrp routes.

Task 2 :

Configure Eigrp Stub on R3, preventing R3 to send any routes to R2, but R2 receives routes from R1.

R3

```
Router eigrp 100
Eigrp stub receive-only
```

Verification:

R2#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0
C 3.0.0.0/8 is directly connected, Serial1/1
C 20.0.0.0/8 is directly connected, Ethernet0/0
D 10.0.0.0/8 [90/20537600] via 2.2.2.1, Serial1/0
```

The output displays only network 10.0.0.0 (eigrp route) coming from R1 but no eigrp routes from R3.

Task 3 :

Configure Eigrp Stub on R3, allowing R3 to send only connected routes to R2, but R2 receives any routes from R1.

R3

```
Router eigrp 100
Eigrp stub connected
```

Verification:

R2#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0
C 3.0.0.0/8 is directly connected, Serial1/1
C 20.0.0.0/8 is directly connected, Ethernet0/0
 172.168.0.0/24 is subnetted, 4 subnets
D 172.168.0.0 [90/20640000] via 3.3.3.2, Serial1/1
D 172.168.1.0 [90/20640000] via 3.3.3.2, Serial1/1
D 172.168.2.0 [90/20640000] via 3.3.3.2, Serial1/1
D 172.168.3.0 [90/20640000] via 3.3.3.2, Serial1/1
D 10.0.0.0/8 [90/20537600] via 2.2.2.1, Serial1/0
```

D 30.0.0.0/8 [90/20537600] via 3.3.3.2, Serial1/1

The output displays only connected eigrp routes from R3 to R2, but receives all routes from R1.

Task 4 :

Configure Eigrp Stub on R3, allowing only summary routes from R3 to R2, but R2 receives any routes from R1.

R3

```
Router eigrp 100
Eigrp stub summary
```

Verification:

R2#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0
C 3.0.0.0/8 is directly connected, Serial1/1
C 20.0.0.0/8 is directly connected, Ethernet0/0
  172.168.0.0/22 is subnetted, 1 subnets
D   172.168.0.0 [90/20640000] via 3.3.3.2, Serial1/1
D  10.0.0.0/8 [90/20537600] via 2.2.2.1, Serial1/0
```

The output displays only summary route from R3, and also all routes from R1.

Task 5 :

Configure Eigrp Stub on R3, allowing connected and summary routes from R3 to R2, but R2 receives any routes from R1.

R3

```
Router eigrp 100
Eigrp stub
```

Verification:

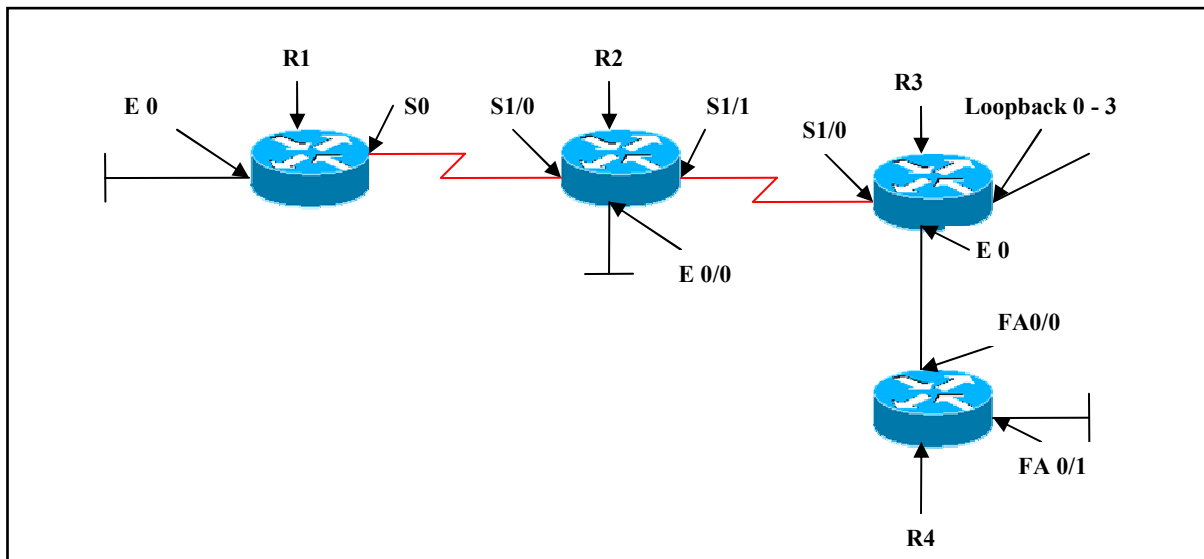
R2#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0
```

- C 3.0.0.0/8 is directly connected, Serial1/1
- C 20.0.0.0/8 is directly connected, Ethernet0/0
- 172.168.0.0/22 is subnetted, 1 subnets
- D 172.168.0.0 [90/20640000] via 3.3.3.2, Serial1/1
- D 10.0.0.0/8 [90/20537600] via 2.2.2.1, Serial1/0
- D 30.0.0.0/8 [90/20537600] via 3.3.3.2, Serial1/1

The output displays both connected and summary routes from R3, as the command `eigrp stub defaults to "eigrp stub connected summary"`.

Task 6 :



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	2.2.2.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.2	255.0.0.0
S1/1	3.3.3.1	255.0.0.0
E 0/0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1/0	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	172.168.0.1	255.255.255.0
Loopback 1	172.168.1.1	255.255.255.0
Loopback 2	172.168.2.1	255.255.255.0
Loopback 3	172.168.3.1	255.255.255.0

R4

Interface	IP Address	Subnet Mask
Fa 0/0	30.1.1.2	255.0.0.0
Fa 0/1	40.1.1.1	255.0.0.0

Lab Objective:

Configure EIGRP in AS 100 on R1, R2, R3. Advertise only interface fa0/0 on R4 in EIGRP AS 100. Configure static route in R3 to reach network 40.0.0.0 via 30.1.1.1. Redistribute the static route in EIGRP AS 100.

R3

```
Ip route 40.0.0.0 255.0.0.0 30.1.1.2
```

```
Router eigrp 100
```

```
Redistribute static metric 10 10 10 10 10
```

```
Eigrp stub static
```

Verification:

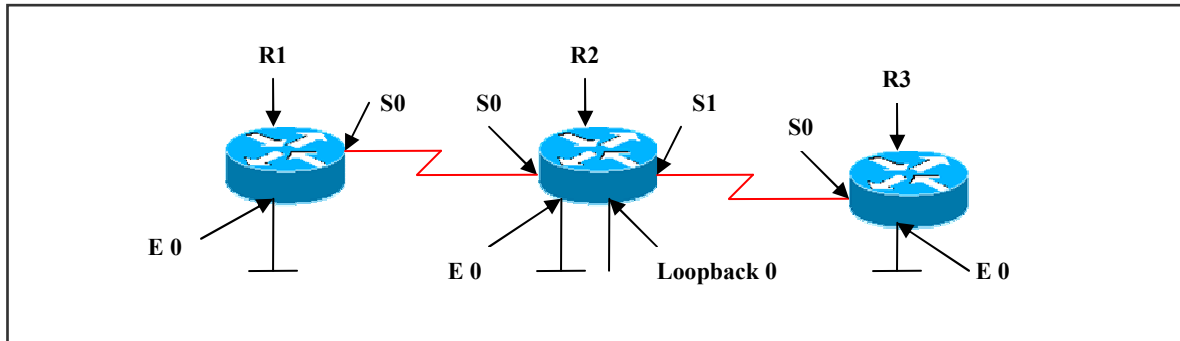
```
R2#show ip route
```


- C 2.0.0.0/8 is directly connected, Serial1/0
- C 3.0.0.0/8 is directly connected, Serial1/1
- C 20.0.0.0/8 is directly connected, Ethernet0/0
- D EX 40.0.0.0/8 [170/256514560] via 3.3.3.2, Serial1/1**
- D 10.0.0.0/8 [90/20537600] via 2.2.2.1, Serial1/0

The output displays only directly connected of R1, R2 and redistributed static route from R3, but blocking connected routes and summary routes from R3.

The output also displays the redistributed route as an external EIGRP route with AD 170

Lab 8– Redistribute EIGRP with RIPv2



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0
Loopback 0	40.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure EIGRP AS 100 on R1 (S0, E0), R2 (S0, E0) and RIPv2 on R2 (S1, Loopback 0) and R3 (S0, E0) as per the above scenario. Mutually redistribute both protocols.

R1 Router eigrp 100 Network 1.0.0.0 Network 10.0.0.0 No auto-summary	R3 Router rip Version 2 Network 2.0.0.0 Network 30.0.0.0 No auto-summary
---	--

R2

```

Router eigrp 100
Network 1.0.0.0
Network 20.0.0.0
No auto-summary
Redistribute rip metric 10 10 10 10 10

Router rip
Version 2
Network 2.0.0.0
Network 40.0.0.0
No auto-summary
Redistribute eigrp 100 metric 10

```

Verification :

R1#show ip route

```

C   1.0.0.0/8 is directly connected, Serial0/2/0
D EX 2.0.0.0/8 [170/256514560] via 1.1.1.2, 00:01:24, Serial0/2/0
D   20.0.0.0/8 [90/2195456] via 1.1.1.2, 00:12:18, Serial0/2/0
D EX 40.0.0.0/8 [170/256514560] via 1.1.1.2, 00:01:24, Serial0/2/0
C   10.0.0.0/8 is directly connected, FastEthernet0/0
D EX 30.0.0.0/8 [170/256514560] via 1.1.1.2, 00:01:24, Serial0/2/0

```

R2#show ip route

```

C   1.0.0.0/8 is directly connected, Serial0
C   2.0.0.0/8 is directly connected, Serial1
C   20.0.0.0/8 is directly connected, Ethernet0
C   40.0.0.0/8 is directly connected, Loopback0
D   10.0.0.0/8 [90/2172416] via 1.1.1.1, 00:14:29, Serial0
R   30.0.0.0/8 [120/1] via 2.2.2.2, 00:00:15, Serial1


```

R3#show ip route

```
R 1.0.0.0/8 [120/10] via 2.2.2.1, 00:00:23, Serial0  
C 2.0.0.0/8 is directly connected, Serial0  
R 20.0.0.0/8 [120/10] via 2.2.2.1, 00:00:23, Serial0  
R 40.0.0.0/8 [120/1] via 2.2.2.1, 00:00:23, Serial0  
R 10.0.0.0/8 [120/10] via 2.2.2.1, 00:00:23, Serial0  
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays that RIP routes are advertised in R1 EIGRP AS 100 as 'D EX' routes. EIGRP routes are advertised in RIP as 'R' routes.

Lab 9 – Redistributing EIGRP with OSPF



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1:

Configure EIGRP AS 100 on R1 (S0, E0), R2 (S0) and OSPF area 0 on R2 (S1, E0), R3 (S0) and OSPF area 1 on R3 (E0) as per the above scenario. Mutually redistribute both protocols.

R1	R3 Router ospf 1
-----------	----------------------------

Router eigrp 100 Network 1.0.0.0 Network 10.0.0.0 No auto-summary	Network 2.2.2.2 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 1
--	---

R2

```
Router eigrp 100
Network 1.0.0.0
No auto-summary
Redistribute ospf 1 metric 10 10 10 10 10

Router ospf 1
Network 2.2.2.1 0.0.0.0 area 0
Network 20.0.0.0 0.255.255.255 area 0
Redistribute eigrp 100 metric 10 subnets
```

Verification :

R1#show ip route

```
1.0.0.0/8 is directly connected, Serial0/2/0
D EX 2.0.0.0/8 [170/256514560] via 1.1.1.2, 00:00:57, Serial0/2/0
D EX 20.0.0.0/8 [170/256514560] via 1.1.1.2, 00:00:57, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
D EX 30.0.0.0/8 [170/256514560] via 1.1.1.2, 00:00:57, Serial0/2/0
```


R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
D 10.0.0.0/8 [90/2172416] via 1.1.1.1, 00:11:11, Serial0
O IA 30.0.0.0/8 [110/74] via 2.2.2.2, 00:00:42, Serial1
```

R3#show ip route

```
O E2 1.0.0.0/8 [110/10] via 2.2.2.1, 00:01:05, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O 20.0.0.0/8 [110/74] via 2.2.2.1, 00:01:05, Serial0
O E2 10.0.0.0/8 [110/10] via 2.2.2.1, 00:01:05, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```

Lab 10 – Configuring EIGRP with Redistribute Connected



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

Task 1 :

Configure EIGRP AS 100 and do not advertise network 10.0.0.0 and network 20.0.0.0 and redistribute network 10.0.0.0 and 20.0.0.0 into EIGRP.

R1	R2
<pre>Router eigrp 100 Network 1.0.0.0 No auto-summary Redistribute connected metric 10 10 10 10 10</pre>	<pre>Router eigrp 100 Network 1.0.0.0 No auto-summary Redistribute connected metric 10 10 10 10 10</pre>

Verification :

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
D EX 20.0.0.0/8 [170/256514560] via 1.1.1.2, 00:00:40, Serial0/2/0
```

C 10.0.0.0/8 is directly connected, FastEthernet0/0

R2#show ip route

C 1.0.0.0/8 is directly connected, Serial0

C 20.0.0.0/8 is directly connected, Ethernet0

D EX 10.0.0.0/8 [170/256514560] via 1.1.1.1, 00:00:33, Serial0


The output displays 'D EX' routes for both networks 10.0.0.0 and 20.0.0.0 in the routing tables.

Module 2 – OSPF

OSPF LAB INDEX

- 1. CONFIGURING OSPF IN SINGLE AREA**
- 2. CONFIGURING OSPF IN MULTIPLE AREA**
- 3. CONFIGURING ABR AND ASBR**
- 4. CONFIGURE STUB**
- 5. CONFIGURE TOTAL STUB**
- 6. CONFIGURE NSSA**
- 7. CONFIGURE NSSA TOTAL STUB**
- 8. OSPF ROUTE SUMMARIZATION**
- 9. OSPF VIRTUAL LINK**
- 10. CONFIGURING OSPF AUTHENTICATION**
- 11. OSPF ON BROADCAST MULTIACCESS**
- 12. OSPF OVER FRAME-RELAY POINT-TO-POINT (SUB-INTERFACE)**
- 13. OSPF OVER FRAME-RELAY POINT-TO-MULTIPOINT (PHYSICAL INTERFACE)**

Lab 1 – Configuring OSPF in a Single Area



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Configure the Interface IP addresses based on the above table

Task 1

Configure OSPF in Area 0. Advertise all networks on all routers.

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 0	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area 0 Network 2.2.2.1 0.0.0.0 area 0 Network 20.0.0.0 0.255.255.255 area 0
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 0	

Verification :

R1 # show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O 2.0.0.0/8 [110/128] via 1.1.1.2, 00:03:58, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:03:58, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O 30.0.0.0/8 [110/138] via 1.1.1.2, 00:03:58, Serial0
```

OSPF routes are displayed as “O” routes in the routing table.

R1 # show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
30.1.1.1	0	FULL/ -	00:00:32	2.2.2.2	Serial1
10.1.1.1	0	FULL/ -	00:00:33	1.1.1.1	Serial0

The symbol indicated by a dash [-] represents that the neighbor is on the serial interface and DR and BDR are not used on point-to-point interfaces.

R1 # show ip ospf

```
Routing Process "OSPF 1" with ID 10.1.1.1
---output omitted---
```

This command displays the OSPF router-id.

Task 2

Configure OSPF in Area 0. Advertise all networks on all routers. Hard Code the Router-id based on the following Loop back ip address:

R1 Loopback 0 6.6.6.6
R2 Loopback 0 7.7.7.7
R3 Loopback 0 8.8.8.8

R1 int loopback 0 ip address 6.6.6.6 255.255.255.255 Router ospf 1 Network 6.6.6.6 0.0.0.0 area 0	R2 int loopback 0 ip address 7.7.7.7 255.255.255.255 Router ospf 1 Network 7.7.7.7 0.0.0.0 area 0
R3 int loopback 0 ip address 8.8.8.8 255.255.255.255 Router ospf 1 Network 8.8.8.8 0.0.0.0 area 0	

Verification:

R1# show ip ospf

Routing Process "OSPF 1" with ID 6.6.6.6
---output omitted---

This output displays that router-id chosen is 6.6.6.6 as it is the loopback address.

Repeat the same on router2 with loopback address as 7.7.7.7 and on router 3 with loopback as 8.8.8.8 and verify using show ip OSPF command

Task 3

Configure OSPF in Area 0. Advertise all networks on all routers. Hard Code the Router-id based on the following :

R1 3.3.3.3
R2 4.4.4.4
R3 5.5.5.5

<p>R1</p> <pre>Router ospf 1 Router-id 3.3.3.3 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 0 Network 6.6.6.6 0.0.0.0 area 0</pre>	<p>R2</p> <pre>Router ospf 1 Router-id 4.4.4.4 Network 1.1.1.2 0.0.0.0 area 0 Network 2.2.2.1 0.0.0.0 area 0 Network 20.0.0.0 0.255.255.255 area 0 Network 7.7.7.7 0.0.0.0 area 0</pre>
<p>R3</p> <pre>Router ospf 1 Router-id 5.5.5.5 Network 2.2.2.2 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 0 Network 8.8.8.8 0.0.0.0 area 0</pre>	

Verification :-

R1#show ip ospf

Routing Process "ospf 1" with ID **3.3.3.3**

This output displays that 3.3.3.3 router-id takes preference over physical and loopback interface.

Lab 2 – Configuring OSPF in Multiple Areas

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure OSPF in Area 0 on R1 (S0), R2 (S0, E0).

Configure OSPF in Area 1 on R1 (E0).

Configure OPSF in Area 2 on R2 (S1), R3 (S0, E0)

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 1	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area0 Network 2.2.2.1 0.0.0.0 area2 Network 20.0.0.0 0.255.255.255 area0
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area2 Network 30.0.0.0 0.255.255.255 area2	

Verification:

R1# show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:07:11, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:07:11, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O IA 30.0.0.0/8 [110/138] via 1.1.1.2, 00:07:11, Serial0/2/0
```

The output displays 'O', 'O IA' routes.

The ABR can be verified by using the following command

R1# show ip ospf border-routers

```
OSPF Process 1 internal Routing Table
Codes: i - Intra-area route, I - Inter-area route
i 20.1.1.1 [64] via 1.1.1.2, Serial0/2/0, ABR, Area 0, SPF 2
```

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
O IA 10.0.0.0/8 [110/65] via 1.1.1.1, 00:11:06, Serial0
O 30.0.0.0/8 [110/74] via 2.2.2.2, 00:11:54, Serial1
```

R2#show ip ospf border-routers

```
OSPF Process 1 internal Routing Table
Codes: i - Intra-area route, I - Inter-area route
i 10.1.1.1 [64] via 1.1.1.1, Serial0, ABR, Area 0, SPF 6
```

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:12:44, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O IA 20.0.0.0/8 [110/74] via 2.2.2.1, 00:12:43, Serial0
O IA 10.0.0.0/8 [110/129] via 2.2.2.1, 00:11:55, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```


Task 2

Configure OSPF as per task 1 and manipulate the Hello-interval time on R1

R1

```
int s0
ip ospf hello-interval 5
```

Verification:

Default hello-interval time:

R1#show ip ospf interface serial 0/2/0

```
Serial0/2/0 is up, line protocol is up
Internet Address 1.1.1.1/8, Area 0
Process ID 1, Router ID 10.1.1.1, Network Type POINT_TO_POINT, Cost: 64
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
20.1.1.1	0	FULL/-	00:00:35	1.1.1.2	Serial0/2/0

Verifying ospf neighbors after manipulating the hello-interval time in R1.

R1#show ip ospf neighbor

-----Nil-----

There will be no neighbor relationship because of hello-interval mismatch.

This can be verified by using 'debug ip ospf events' command, where the output displays a mismatch hello parameter statement.

R1#debug ip ospf events

```
*May 28 09:20:31.403: OSPF: Rcv hello from 20.1.1.1 area 0 from Serial0/2/0 1.1.1.2
```

```
*May 28 09:20:31.403: OSPF: Mismatched hello parameters from 1.1.1.2
```

```
*May 28 09:20:31.403: OSPF: Dead R 40 C 20, Hello R 10 C 5
```

The output displays a mismatch hello parameter statement.

Lab 3 – Configuring ABR and ASBR

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure OSPF in Area 0 on R1 (S0/2/0, E0), R2 (S0)
 Configure OSPF in Area 1 on R2 (S1), R3 (S0, E 0).
 Configure EIGRP AS 100 on R2 (E0) and redistribute into OSPF.

<p>R1</p> <pre>Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 0</pre> <p>R3</p> <pre>Router ospf 1 Network 2.2.2.2 0.0.0.0 area 1 Network 30.0.0.0 0.255.255.255 area 1</pre>	<p>R2</p> <pre>Router ospf 1 Network 2.2.2.1 0.0.0.0 area 1 Network 1.1.1.2 0.0.0.0 area 0</pre> <p>Router eigrp 100 Network 20.0.0.0 No auto-summary</p> <p>Router ospf 1 Redistribute eigrp 100 metric 10 subnets</p>
---	---

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:12:21, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O IA 30.0.0.0/8 [110/138] via 1.1.1.2, 00:11:13, Serial0/2/0
```

The output displays 'O' and 'O IA' routes.

The output also shows that network 20.0.0.0 is missing in the routing table.

As EIGRP is a NON-OSPF routing protocol, we need to redistribute EIGRP into OSPF

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:00:04, Serial0/2/0
O E2 20.0.0.0/8 [110/10] via 1.1.1.2, 00:00:04, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O IA 30.0.0.0/8 [110/138] via 1.1.1.2, 00:00:04, Serial0/2/0
```

Note: If we want OE1 routes then the redistribute command should be configured using **metric-type**

R2

```
Router ospf 1
Redistribute eigrp 100 metric-type 1 metric
10 subnets
```

.R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:01:43, Serial0/2/0
O E1 20.0.0.0/8 [110/74] via 1.1.1.2, 00:00:12, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O IA 30.0.0.0/8 [110/138] via 1.1.1.2, 00:01:43, Serial0/2/0
```

To verify which router is ABR / ASBR :-

R1 # show ip ospf border-routers

OSPF Process 1 internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

```
i 20.1.1.1 [64] via 1.1.1.2, Serial0/2/0, ABR/ASBR, Area 0, SPF
```

Lab 4 – Configure OSPF Stub Area

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure OSPF in Area 0 on R1 (S0/2/0, E0), R2 (S0)

Configure OSPF in Area 1 on R2 (S1), R3 (S0, E 0).

<p>R1</p> <pre>Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 1</pre>	<p>R2</p> <pre>Router ospf 1 Network 1.1.1.2 0.0.0.0 area0 Network 2.2.2.1 0.0.0.0 area2</pre>
<p>R3</p> <pre>Router ospf 1 Network 2.2.2.2 0.0.0.0 area2 Network 30.0.0.0 0.255.255.255 area2</pre>	

Task 2 : Configure EIGRP AS 100 on R2 (E0) and redistribute into OSPF.

R2

```
Router eigrp100
Network 20.0.0.0
No auto-summary
```

```
Router ospf 1
Redistribute eigrp 100 metric 10 subnets
```

Verification :

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:01:08, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O E2 20.0.0.0/8 [110/10] via 2.2.2.1, 00:00:03, Serial0
O IA 10.0.0.0/8 [110/129] via 2.2.2.1, 00:01:08, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays inter-area routes (O IA) and OSPF external type 2 (O E2).

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	243	0x8000000A	0x00B788	2
30.1.1.1	30.1.1.1	243	0x80000008	0x0034CD	3

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
1.0.0.0	20.1.1.1	277	0x80000004	0x002FB2
10.0.0.0	20.1.1.1	277	0x80000004	0x00C314

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
20.0.0.0	20.1.1.1	172	0x80000007	0x00A8D0	0

The output displays Type-5 external link-states.

Task 3

Configure OSPF Area 1 as Stub.

R2	R3
Router ospf 1 Area 1 stub	Router ospf 1 Area 1 stub

After configuring stub, verify the routing table on R3

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.1.1	0	FULL/ -	00:00:30	1.1.1.1	Serial0
30.1.1.1	0	FULL/ -	00:00:38	2.2.2.2	Serial1

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:00:03, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O IA 10.0.0.0/8 [110/129] via 2.2.2.1, 00:00:03, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
O*IA 0.0.0.0/0 [110/65] via 2.2.2.1, 00:00:03, Serial0
```

The output displays default route and inter-area routes, both designated with (OIA) in the routing table.

Default route is denoted as (O* IA).

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)
Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	543	0x8000000C	0x00CB76	2
30.1.1.1	30.1.1.1	543	0x8000000A	0x004EB3	3

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	20.1.1.1	552	0x80000001	0x00E73F
1.0.0.0	20.1.1.1	552	0x80000005	0x004B97
10.0.0.0	20.1.1.1	552	0x80000005	0x00DFF8

The output does not display the 'Type 5 External LSA.

Note: If stub is not configured on both routers OSPF neighborship will not establish. It can be verified by the following commands.

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.1.1	0	FULL/ -	00:00:35	1.1.1.1	Serial0
30.1.1.1	0	DOWN/ -	-	2.2.2.2	Serial1

R2#debug ip ospf events

```
Mar 1 03:12:42.491: OSPF: Rcv hello from 30.1.1.1 area 1 from Serial1 2.2.2.2
*Mar 1 03:12:42.491: OSPF: Hello from 2.2.2.2 with mismatched Stub/Transit area
option bit
```

The output displays mismatched Stub/Transit area option bit .

Lab 5 – Configuring Totally Stub Area

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure OSPF in Area 0 on R1 (S0/2/0, E0), R2 (S0)
 Configure OSPF in Area 1 on R2 (S1), R3 (S0, E 0).

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 1	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area0 Network 2.2.2.1 0.0.0.0 area2
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area2 Network 30.0.0.0 0.255.255.255 area2	

Task 2 : Configure EIGRP AS 100 on R2 (E0) and redistribute into OSPF.

R2

```
Router eigrp100
Network 20.0.0.0
No auto-summary
```

```
Router ospf 1
Redistribute eigrp 100 metric 10 subnets
```

Verification :

Verify the routing table on R3:

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:01:08, Serial0
C    2.0.0.0/8 is directly connected, Serial0
O E2 20.0.0.0/8 [110/10] via 2.2.2.1, 00:00:03, Serial0
O IA 10.0.0.0/8 [110/129] via 2.2.2.1, 00:01:08, Serial0
C   30.0.0.0/8 is directly connected, Ethernet0
```

The output displays inter-area (O IA) and external type 2 (O E2) routes.

The OSPF database on R3 can be verified using the following command :

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	243	0x8000000A	0x00B788	2
30.1.1.1	30.1.1.1	243	0x80000008	0x0034CD	3

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
1.0.0.0	20.1.1.1	277	0x80000004	0x002FB2
10.0.0.0	20.1.1.1	277	0x80000004	0x00C314

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
20.0.0.0	20.1.1.1	172	0x80000007	0x00A8D0	0

The output displays summary net link states and type-5 AS external link-states.

Now, to block both the summary net link-states and type-5 external link-states, configure Area 1 as total stub .

Task 3

Configure R2 and R3 as total stub .

R2	R3
Router ospf 1 Area 1 stub no-summary	Router ospf 1 Area 1 stub no-summary

Verifying the routing table on R3

R3#show ip route

```
C 2.0.0.0/8 is directly connected, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
O*IA 0.0.0.0/0 [110/65] via 2.2.2.1, 00:00:30, Serial0
```

Inter-area and external routes are not visible in the routing table, but they are accessible via the inter-area default route (O * IA).

Verify the OSPF database

R3#show ip ospf database

```
OSPF Router with ID (30.1.1.1) (Process ID)
-----Output has been omitted for brevity-----
```

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	20.1.1.1	125	0x80000003	0x00E341

No Type-5 External LSA and Summary Net Link Type 3, but you can see a default route.

Lab 6 – Configuring NSSA

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure OSPF in Area 0 on R1 (S0/2/0), R2 (S0).

Configure OSPF in Area 1 on R2 (S1), R3 (S0).

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area0 Network 2.2.2.1 0.0.0.0 area2
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area2	

Task 2: Configure EIGRP AS 100 on R2 (E0) and redistribute into OSPF.

R2

```
Router eigrp100
Network 20.0.0.0
No auto-summary
```

```
Router ospf 1
Redistribute eigrp 100 metric 10 subnets
```

Task 3: Configure RIPv2 on R1 (E0), R3 (E0) and redistribute into OSPF.

R1

```
Router rip
Net 10.0.0.0
No auto-summary
Version 2
```

```
Router ospf 1
Redistribute rip metric 10 subnets
```

R3

```
Router rip
Net 30.0.0.0
No auto-summary
Version 2
```

```
Router ospf 1
Redistribute rip metric 10 subnets
```

Verification :

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:00:22, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O E2 20.0.0.0/8 [110/10] via 2.2.2.1, 00:00:22, Serial0
O E2 10.0.0.0/8 [110/10] via 2.2.2.1, 00:00:22, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays inter-area (O IA), external type2 (O E2) routes.

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	141	0x80000013	0x00A591	2
30.1.1.1	30.1.1.1	141	0x80000010	0x00939C	2

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
1.0.0.0	20.1.1.1	149	0x80000001	0x0035AF

Summary ASB Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.1.1	20.1.1.1	149	0x80000001	0x009047

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.0.0	10.1.1.1	328	0x80000001	0x009102	0
20.0.0.0	20.1.1.1	1830	0x80000008	0x00A6D1	0

The OSPF database displays summary net link states, type-5 external net link states.

Now, configure NSSA on R2 & R3, where R3 acts as NSSA ASBR that generates type-7 LSA and R2 acts as NSSA ABR that converts the type-7 LSA into type-5 LSA, when it leaves the NSSA area.

Task 4

Configure R2 and R3 as NSSA .

R2	R3
Router ospf 1 Area 1 nssa default-information-originate	Router ospf 1 Area 1 nssa default-information-originate

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:00:06, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O N2 20.0.0.0/8 [110/10] via 2.2.2.1, 00:00:06, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
O*N2 0.0.0.0/0 [110/1] via 2.2.2.1, 00:00:06, Serial0
```

The output displays 'O N2' and 'O* N2' routes in the routing table.

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	428	0x80000015	0x0047E7	2
30.1.1.1	30.1.1.1	428	0x80000012	0x0035F2	2

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
1.0.0.0	20.1.1.1	435	0x80000002	0x00D805

Type-7 AS External Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	20.1.1.1	435	0x80000001	0x0099F9	0
20.0.0.0	20.1.1.1	434	0x80000001	0x00EE87	0
30.0.0.0	30.1.1.1	458	0x80000001	0x00A7B1	0

No Type-5 External Link States but allows Special Type-7 External Link State

R2#show ip route

C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
O E2 10.0.0.0/8 [110/10] via 1.1.1.1, 00:02:33, Serial0
O N2 30.0.0.0/8 [110/10] via 2.2.2.2, 00:02:33, Serial1

R1#show ip route

C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:14:38, Serial0/2/0
O E2 20.0.0.0/8 [110/10] via 1.1.1.2, 00:03:23, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O E2 30.0.0.0/8 [110/10] via 1.1.1.2, 00:03:16, Serial0/2/0

R1#show ip ospf database

OSPF Router with ID (10.1.1.1) (Process ID 1)

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
10.1.1.1	10.1.1.1	145	0x8000000B	0x001250	2
20.1.1.1	20.1.1.1	133	0x8000000B	0x00FF56	2

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
2.0.0.0	20.1.1.1	1429	0x80000005	0x0020BF

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.0.0	10.1.1.1	145	0x80000002	0x008F03	0
20.0.0.0	20.1.1.1	1429	0x80000009	0x00A4D2	0
30.0.0.0	20.1.1.1	1263	0x80000001	0x0096D6	0

Lab 7 – Configure NSSA Total Stub

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1

Configure OSPF in Area 0 on R1 (S0/2/0), R2 (S0).

Configure OSPF in Area 1 on R2 (S1), R3 (S0).

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area0 Network 2.2.2.1 0.0.0.0 area2
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area2	

Task 2: Configure EIGRP AS 100 on R2 (E0) and redistribute into OSPF.

R2

```
Router eigrp100
Network 20.0.0.0
No auto-summary
```

```
Router ospf 1
Redistribute eigrp 100 metric 10 subnets
```

Task 3: Configure RIPv2 on R1 (E0), R3 (E0) and redistribute into OSPF.

R1

```
Router rip
Net 10.0.0.0
No auto-summary
Version 2
```

```
Router ospf 1
Redistribute rip metric 10 subnets
```

R3

```
Router rip
Net 30.0.0.0
No auto-summary
Version 2
```

```
Router ospf 1
Redistribute rip metric 10 subnets
```

Verification:

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:00:22, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O E2 20.0.0.0/8 [110/10] via 2.2.2.1, 00:00:22, Serial0
O E2 10.0.0.0/8 [110/10] via 2.2.2.1, 00:00:22, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays inter-area (O IA), external type2 (O E2) routes.

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	141	0x80000013	0x00A591	2
30.1.1.1	30.1.1.1	141	0x80000010	0x00939C	2

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
1.0.0.0	20.1.1.1	149	0x80000001	0x0035AF

Summary ASB Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
10.1.1.1	20.1.1.1	149	0x80000001	0x009047

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
10.0.0.0	10.1.1.1	328	0x80000001	0x009102	0
20.0.0.0	20.1.1.1	1830	0x80000008	0x00A6D1	0

The OSPF database displays summary net link states, type-5 external net link states.

Task 4 :

Configure R2 and R3 as NSSA Total Stub .

R2	R3
Router ospf 1 Area 1 nssa no-summary	Router ospf 1 Area 1 nssa no-summary

R3#show ip route

```
C 2.0.0.0/8 is directly connected, Serial0
O N2 20.0.0.0/8 [110/10] via 2.2.2.1, 00:00:15, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
O*IA 0.0.0.0/0 [110/65] via 2.2.2.1, 00:00:15, Serial0
```

The output displays O N2 and O* IA routes only.

R3#show ip ospf database

OSPF Router with ID (30.1.1.1) (Process ID 1)

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
20.1.1.1	20.1.1.1	177	0x80000017	0x0043E9	2
30.1.1.1	30.1.1.1	118	0x80000015	0x002FF5	2

Summary Net Link States (Area 1)


Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	20.1.1.1	187	0x80000001	0x006FAF

Type-7 AS External Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
20.0.0.0	20.1.1.1	186	0x80000001	0x00EE87	0
30.0.0.0	30.1.1.1	118	0x80000002	0x00A5B2	0

No Type-5 External Link States, no Type-3 Summary link but allows Special Type-7 External Link State .

Lab 8 – Configure OSPF Route Summarization



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	172.168.0.1	255.255.255.0
Loopback 1	172.168.1.1	255.255.255.0
Loopback 2	172.168.2.1	255.255.255.0
Loopback 3	172.168.3.1	255.255.255.0

Task 1 : Configure Route Summarization at ABR

Configure OSPF in Area 0 on R1 (S0/2/0, E0), R2 (S0).

Configure OSPF in Area 1 on R2 (S1), R3 (S0, E0, Loopback 0 – 3).

Create the following Loopbacks on R3:

Loopback 0 – 172.168.0.1/24
Loopback 1 – 172.168.1.1/24
Loopback 2 – 172.168.2.1/24
Loopback 3 – 172.168.3.1/24

Advertise these newly created loopbacks in OSPF using the network command. Make sure they appear in the routing table using a /24 mask. These routes should be seen as a single summarized route outside of area 1.

R3	R2
<pre>Int loopback 0 Ip add 172.168.0.1 255.255.255. 0 Ip ospf network point-to-point Int loopback 1 Ip add 172.168.1.1 255.255.255. 0 Ip ospf network point-to-point Int loopback 2 Ip add 172.168.2.1 255.255.255. 0 Ip ospf network point-to-point Int loopback 3 Ip add 172.168.3.1 255.255.255. 0 Ip ospf network point-to-point Router ospf 1 Network 172.168.0.0 0.0.255.255 area 1 Network 2.2.2.2 0.0.0.0 area 1 Network 30.0.0.0 0.255.255 area 1</pre>	<pre>Router ospf 1 Area 1 range172.168.0.0 255.255.252.0</pre>

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:17:26, Serial0/2/0
O IA 20.0.0.0/8 [110/74] via 1.1.1.2, 00:17:26, Serial0/2/0
 172.168.0.0/22 is subnetted, 1 subnets
O IA 172.168.0.0 [110/129] via 1.1.1.2, 00:00:11, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```

The output displays a smaller routing table by displaying only one summarized route for the contiguous networks.

Task 2 : Configure Route Summarization At ASBR

(Scenario Based On Task 1)

Configure OSPF on the routers as per the above scenario.
Create the following Loopbacks on R3:

Loopback 0 – 172.168.0.1/24
Loopback 1 – 172.168.1.1/24
Loopback 2– 172.168.2.1/24
Loopback 3 – 172.168.3.1/24

Advertise these newly created loopbacks in EIGRP AS 100 using the network command and redistribute these networks into OSPF Area 1. These routes should be seen as a single summarized route.

R3

```
Int loopback 0
Ip add 172.168.0.1 255.255.255. 0
Ip ospf network point-to-point
```

```
Int loopback 1
Ip add 172.168.1.1 255.255.255. 0
Ip ospf network point-to-point
```

```
Int loopback 2
Ip add 172.168.2.1 255.255.255. 0
Ip ospf network point-to-point
```

```
Int loopback 3
Ip add 172.168.3.1 255.255.255. 0
Ip ospf network point-to-point
```

```
Router ospf 1
Network 2.2.2.2 0.0.0.0 area 1
Network 30.0.0.0 0.255.255 area 1
```

```
Router eigrp 100
Network 172.168.0.0
No auto-summary
```

```
Router ospf 1
Redistribute eigrp 100 metric 10 subnets
Summary-address 172.168.0.0 255.255.252.0
```

Verification:

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
  172.168.0.0/22 is subnetted, 1 subnets
O E2 172.168.0.0 [110/10] via 2.2.2.2, 00:00:45, Serial1
O 10.0.0.0/8 [110/65] via 1.1.1.1, 00:07:28, Serial0
O 30.0.0.0/8 [110/74] via 2.2.2.2, 00:07:28, Serial1
```


The output displays a smaller routing table.

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:12:14, Serial0
C 2.0.0.0/8 is directly connected, Serial0
O 20.0.0.0/8 [110/74] via 2.2.2.1, 00:12:14, Serial0
  172.168.0.0/16 is variably subnetted, 5 subnets, 2 masks
C 172.168.0.0/24 is directly connected, Loopback0
O 172.168.0.0/22 is a summary, 00:03:01, Null0
C 172.168.1.0/24 is directly connected, Loopback1
C 172.168.2.0/24 is directly connected, Loopback2
C 172.168.3.0/24 is directly connected, Loopback3
O IA 10.0.0.0/8 [110/129] via 2.2.2.1, 00:12:14, Serial0
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays a summary route pointing to interface null 0 on R3 routing table. This is automatically generated by default, when manual summarization is configured so as to prevent routing loops.

Lab 9 – Configuring OSPF Virtual Links



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
S1	2.2.2.1	255.0.0.0
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Task 1 :

Configure OSPF in Area 0 on R1 (S0/2/0, E0), R2 (S0, E0).
 Configure OSPF in Area 1 on R2 (S1), R3 (S0).
 Configure OSPF in Area 2 on R3 (E0).

R1	R2
Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0	Router ospf 1 Network 1.1.1.2 0.0.0.0 area0

Network 10.0.0.0 0.255.255.255 area 0	Network 2.2.2.1 0.0.0.0 area1 Network 20.0.0.0 0.255.255.255 area 0
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area1 Network 30.0.0.0 0.255.255.255 area 2	

Verification:

Verifying the routing table on R1 in area0:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:22:43, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:22:43, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```

The output displays net 20.0.0.0 as ‘O’ and net 2.0.0.0 as ‘O IA’, but there is no net 30.0.0.0, as it is not connected to area0.

We need to configure virtual links between R2 & R3 and this area that connects to area0 is called the transit area.

Each router R2 & R3 point at the router-id of the other router.

Task 2:

Configure Virtual Link between R2 and R3:

R2 Router ospf 1 Area 1 virtual-link 30.1.1.1	R3 Router ospf 1 Area 1 virtual-link 20.1.1.1
--	--

Verifying the routing table on R1 :

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:00:00, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:00:00, Serial0/2/0
```

C 10.0.0.0/8 is directly connected, FastEthernet0/0
O IA 30.0.0.0/8 [110/138] via 1.1.1.2, 00:00:00, Serial0/2/0

The output displays network 30.0.0.0 as ‘O’ route because of the virtual link configured, the router1 assumes that net 30.0.0.0 is in the same area0.

R2#show ip ospf virtual-links

Virtual Link OSPF_VL0 to router 30.1.1.1 is up

Run as demand circuit

DoNotAge LSA allowed

Transit area 1, via interface Serial1, Cost of using 64

Transmit Delay is 1 sec, State POINT_TO_POINT,

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

Hello due in 00:00:00

Adjacency State FULL (Hello suppressed)

Index 2/3, retransmission queue length 0, number of retransmission 1

First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)

Last retransmission scan length is 1, maximum is 1

Last retransmission scan time is 0 msec, maximum is 0 msec

The output displays virtual-link to other router and as well, ‘DoNotAge’ option set.

Task 3: Configure Virtual Link when area connecting two backbone areas.

(Scenario Based on Task 1)

Configure OSPF in Area 0 on R1 (E0).

Configure OSPF in Area 1 on R1 (S0/2/0), R2 (S0).

Configure OSPF in Area 2 on R2 (E0, S1), R3 (S0, E0).

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 1 Network 10.0.0.0 0.255.255.255 area 0	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area 1 Network 2.2.2.1 0.0.0.0 area 0 Network 20.0.0.0 0.255.255.255 area 0
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 0	

Verification:

Verify the routing table on R1

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O IA 2.0.0.0/8 [110/128] via 1.1.1.2, 00:00:43, Serial0/2/0
O IA 20.0.0.0/8 [110/74] via 1.1.1.2, 00:00:43, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O IA 30.0.0.0/8 [110/138] via 1.1.1.2, 00:00:34, Serial0/2/0
```

The output displays network 2.0.0.0, 20.0.0.0 and 30.0.0.0 as O IA routes.

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
O 30.0.0.0/8 [110/74] via 2.2.2.2, 00:05:26, Serial 1
O IA 10.0.0.0/8[110/74]via 1.1.1.2, 00:07:43, Serial 0/2/0
```

R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:08:27, Serial1
C 2.0.0.0/8 is directly connected, Serial1
O 20.0.0.0/8 [110/74] via 2.2.2.1, 00:08:27, Serial1
C 30.0.0.0/8 is directly connected, Ethernet0
```

When we check the routing table on R3, the output does not have network 10.0.0.0 in the routing table.

Task 4 :

Configure Virtual Link between R1 and R2 :

R1	R2
Router ospf 1 Area 1 virtual-link 20.1.1.1	Router ospf 1 Area 1 virtual-link 10.1.1.1

After Configuring Virtual Link:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O 2.0.0.0/8 [110/128] via 1.1.1.2, 00:00:45, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:00:45, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
O 30.0.0.0/8 [110/138] via 1.1.1.2, 00:00:45, Serial0/2/0
```

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
C 20.0.0.0/8 is directly connected, Ethernet0
O 10.0.0.0/8 [110/65] via 1.1.1.1, 00:02:10, Serial0
O 30.0.0.0/8 [110/74] via 2.2.2.2, 00:02:10, Serial1
```


R3#show ip route

```
O IA 1.0.0.0/8 [110/128] via 2.2.2.1, 00:05:37, Serial1
C 2.0.0.0/8 is directly connected, Serial1
O 20.0.0.0/8 [110/74] via 2.2.2.1, 00:05:37, Serial1
O 10.0.0.0/8 [110/129] via 2.2.2.1, 00:05:37, Serial1
C 30.0.0.0/8 is directly connected, Ethernet0
```

Now, when we verify the routing table on R1, R2, R3, we see that all O IA routes are advertised as 'O' routes as the routers assume that the networks belong to the same area because of the virtual link.

Also we can see network 10.0.0.0 in R3 routing table as O route.

Lab 10 – Configuring OSPF Authentication



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S0	1.1.1.2	255.0.0.0
E0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

All routers should Authenticate Routing updates using the simple password authentication method. Use a key-string of **cisco123**.

R1

```
Router ospf 1
Network 1.1.1.1 0.0.0.0 area 0
Network 10.0.0.0 0.255.255.255 area 0
```

```
Int s 0/2/0
Ip ospf authentication-key cisco123
Ip ospf authentication
```

R2

```
Router ospf 1
Network 1.1.1.2 0.0.0.0 area 0
Network 20.0.0.0 0.255.255.255 area 0
```

```
Int s0
Ip ospf authentication-key cisco123
Ip ospf authentication
```

Verification :

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
20.1.1.1	0	FULL/ -	00:00:34	1.1.1.2	Serial0/2/0

The output displays neighbor in full state.

If there is a mismatch in the password, there will be no OSPF neighbor relationship established.

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:09:23, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```

Simple authentication on R1 and R2, but different passwords:

R1#debug ip ospf adj

```
*May 29 06:29:03.966: OSPF: Rcv pkt from 1.1.1.2, Serial0/2/0 : Mismatch
Authentication Key - Clear Text
```

Simple authentication on R1, no authentication on R2:

R1#debug ip ospf adj

```
*May 29 06:37:03.982: OSPF: Rcv pkt from 1.1.1.2, Serial0/2/0 : Mismatch
Authentication type. Input packet specified type 0, we use type 1
```

R2#debug ip ospf adj

```
*Mar 1 01:14:09.311: OSPF: Rcv pkt from 1.1.1.1, Serial0 : Mismatch Authentication
type. Input packet specified type 1, we use type 0
```

Task 2

(Scenario based on Task 1)

All routers should Authenticate Routing updates using the most secure authentication method. Use Key 1 with a key-string of **cisco123**. Do not use wide authentication.

R1

```
Router ospf 1
Network 1.1.1.1 0.0.0.0 area 0
Network 10.0.0.0 0.255.255.255 area 0
```

```
Int S0/2/0
Ip ospf message-digest-key 1 md5 cisco123
Ip ospf authentication message-digest
```

R2

```
Router ospf 1
Network 1.1.1.2 0.0.0.0 area 0
Network 20.0.0.0 0.255.255.255 area 0
```

```
Int S0
Ip ospf message-digest-key 1 md5 cisco123
Ip ospf authentication message-digest
```

Verification :

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
20.1.1.1	0	FULL/ -	00:00:34	1.1.1.2	Serial0/2/0

If there is mismatch in key or password, there will not be OSPF neighbor relationship established between the two routers.

R1#show ip route


```
C 1.0.0.0/8 is directly connected, Serial0/2/0
O 20.0.0.0/8 [110/74] via 1.1.1.2, 00:09:23, Serial0/2/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```

MD5 authentication on R1, no authentication on R2:

*May 29 06:50:02.162: OSPF: Send with youngest Key 1

*May 29 06:50:04.054: OSPF: Rcv pkt from 1.1.1.2, Serial0/2/0 : **Mismatch Authentication type**. Input packet specified type 0, **we use type 2**

Lab 11 – OSPF on Broadcast Multiaccess



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Fa0/0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
E0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
E0	30.1.1.1	255.0.0.0

Lab Objective:

Task

Configure OSPF as per the above scenario.

R1 Router ospf 1 Network 10.0.0.0 0.255.255.255 area 0	R2 Router ospf 1 Network 20.0.0.0 0.255.255.255 area 0
R3 Router ospf 1 Network 30.0.0.0 0.255.255.255 area 0	

Verification :

The DR & BDR election can be verified by using the following commands

R2#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.1.2	1	FULL/BDR	00:00:35	10.1.1.1	Ethernet1
10.1.1.3	1	FULL/DR	00:00:30	10.1.1.3	Ethernet1

R3#show ip ospf interface

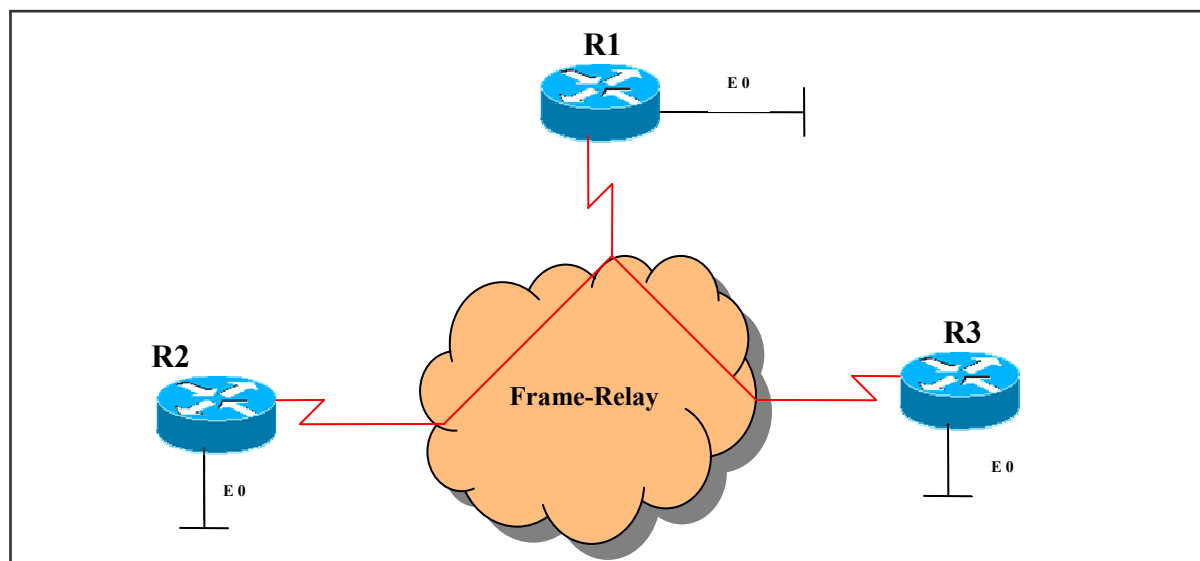
Ethernet0 is up, line protocol is up
Internet Address 10.1.1.3/8, Area 0
Process ID 1, Router ID 10.1.1.3, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 10.1.1.3, Interface address 10.1.1.3
Backup Designated router (ID) 10.1.1.2, Interface address 10.1.1.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
oob-resync timeout 40
Hello due in 00:00:00
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 1
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 2, Adjacent neighbor count is 2
Adjacent with neighbor 10.1.1.2 (Backup Designated Router)
Adjacent with neighbor 10.1.1.1

- Routers notify DR on 224.0.0.6
- DR notifies others on 224.0.0.5.
- Router with highest priority value is the DR
- Router with second highest priority value is BDR
- Default for interface OSPF priority is 1
- In case of a tie, router-id (highest) is preferred
- A router with priority set to 0 cannot become the DR/BDR and is called DR other
- DR election is non-preemptive

Setting priority for DR election

(config-if) # ip ospf priority number
-Default is 1
-Range is 0-255.

Lab 12 – OSPF over Frame-Relay Point-to-Point Subinterfaces



IP addressing and DLCI information Chart

Routers	IP address	Local DLCI	Connecting to:
R1	S 0.1 : 1.1.1.1 / 8	100	R2
	S 0.2 : 2.2.2.1 / 8	200	R3
	E 0 : 10.1.1.1 / 8		
R2	S 0: 1.1.1.2 / 8	300	R1
	E 0 : 20.1.1.1 / 8		
R3	S 0: 2.2.2.2 / 8	400	R1
	E 0 : 30.1.1.1 / 8		

Task 1

Configure the frame-relay cloud in a hub and spoke topology without using frame-relay map statements. These routers should reply to inverse-arp inquiries. Routers should be configured in a point-to-point configuration.

<p>R1</p> <p>Int S0 No ip address No shutdown Encapsulation frame-relay Int serial 0.1 point-to-point Ip add 1.1.1.1 255.0.0.0 Frame-relay interface-dlci 100 No shutdown Int serial 0.2 point-to-point Ip add 2.2.2.1 255.0.0.0 Frame-relay interface-dlci 200 No shutdown</p>	<p>R2</p> <p>Int S0 Ip add 1.1.1.2 255.0.0.0 Encapsulation frame-relay Frame-relay interface-dlci 300 No shutdown</p>
<p>R3</p> <p>Int S0 Ip add 2.2.2.2 255.0.0.0 Encapsulation frame-relay Frame-relay interface-dlci 400</p>	<p>FRS</p> <p>Frame-relay switching</p> <p>Int S0 No ip add Encapsulation frame-relay Frame-relay intf-type dce Clock rate 64000 Frame-relay route 100 interface serial 1 300 Frame-relay route 200 interface serial 2 400 No shutdown</p> <p>Int S1 No ip address Encapsulation frame-relay Frame-relay intf-type dce Clock rate 64000 Frame-relay route 300 interface serial 0 100 No shutdown</p> <p>Int S2 No ip address Encapsulation frame-relay Frame-relay route 400 interface serial 0 200 No shutdown</p>

Verify frame-relay connectivity:

R3 # show frame-relay pvc

The output displays pvc-status = active
When we ping to a network, the rate is 100% successful

Task 2 :

Configure OSPF Over Frame- Relay on R1 , R2 and R3.

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 2.2.2.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 0	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area 0 Network 20.0.0.0 0.255.255.255 area 0
R3 Router ospf 1 Network 2.2.2.2 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 0	

Verify OSPF neighbors:

The output displays no neighbors as the hello-intervals did not match on the routers.

The reason is, as follows:

When R1 - FR was configured on sub-interface, the OSPF default mode is point-to-point (hello 10 seconds)

When R2 & R3 - FR was configured on physical interface, the OSPF default mode is non-broadcast (hello 30 sec)

Therefore, we need to manually change the hello-interval time or change the network type on R2 and R3 interfaces.

R2 Int s0 ip ospf network point-to-point	R3 Int s0 ip ospf network point-to-point
---	---

Now verify OSPF neighbors:

- The output displays neighbors in full state.
- The routing table can also be verified where the output displays all 'O' routes as in the same area 0.

Verification:

Default mode of ospf on a point-to-point frame-relay physical interface:


R2#show ip ospf interface S0

```
Serial0 is up, line protocol is up
Internet Address 1.1.1.2/8, Area 0
Process ID 1, Router ID 20.1.1.1, Network Type NON_BROADCAST, Cost: 64
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 20.1.1.1, Interface address 1.1.1.2
No backup designated router on this network
Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
```

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
O 2.0.0.0/8 [110/128] via 1.1.1.1, 00:00:00, Serial0
C 20.0.0.0/8 is directly connected, Ethernet0
O 10.0.0.0/8 [110/74] via 1.1.1.1, 00:00:00, Serial0
O 30.0.0.0/8 [110/138] via 1.1.1.1, 00:00:00, Serial0
```

Lab 13 – OSPF over Frame-Relay Point-to-Multipoint (Physical Interfaces)



IP addressing and DLCI information Chart

Routers	IP address	Local DLCI	Connecting to:
R1	S 0 : 1.1.1.1 / 8	100	R2 R3
	E 0 : 10.1.1.1 / 8	200	
R2	S 0 : 1.1.1.2 / 8	300	R1
	E 0 : 20.1.1.1 / 8		
R3	S 0 : 1.1.1.3 / 8	400	R1
	E 0 : 30.1.1.1 / 8		

Task 1

Configure the frame-relay cloud in a hub and spoke topology using frame-relay map statements. These routers should NOT reply to inverse-arp inquiries.

<p>R1</p> <p>Int S0 Ip add 1.1.1.1 255.0.0.0 Encapsulation frame-relay Frame-relay map ip 1.1.1.2 100 Frame-relay map ip 1.1.1.3 200 No shutdown</p>	<p>R2</p> <p>Int S0 Ip add 1.1.1.2 255.0.0.0 Encapsulation frame-relay Frame-relay map ip 1.1.1.1 300 No shutdown</p>
<p>R3</p> <p>Int S0 Ip add 1.1.1.3 255.0.0.0 Encapsulation frame-relay Frame-relay map ip 1.1.1.1 400 No shutdown</p>	<p>FRS</p> <p>Frame-relay switching</p> <p>Int S0 No ip add Encapsulation frame-relay Frame-relay intf-type dce Clock rate 64000 Frame-relay route 100 interface serial 2 300 Frame-relay route 200 interface serial 1 400 No shutdown</p> <p>Int S2 No ip address Encapsulation frame-relay Frame-relay intf-type dce Clock rate 64000 Frame-relay route 300 interface serial 0 100 No shutdown</p> <p>Int S1 No ip address Encapsulation frame-relay Frame-relay intf-type dce Frame-relay route 400 interface serial 0 200 No shutdown</p>

Verify frame-relay connectivity :

R3 # show frame-relay pvc

The output displays pvc-status = active
When we ping to a network, the rate is 100% successful

Task 2 :

Configure OSPF Over Frame- Relay on R1 , R2 and R3.

R1 Router ospf 1 Network 1.1.1.1 0.0.0.0 area 0 Network 10.0.0.0 0.255.255.255 area 0	R2 Router ospf 1 Network 1.1.1.2 0.0.0.0 area 0 Network 20.0.0.0 0.255.255.255 area 0
R3 Router ospf 1 Network 1.1.1.3 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 0	

If we don't use broadcast option in the frame-relay map command, the pvc will not allow broadcast through them so that the neighbors should be statically configured.

Configure neighbor statement manually on R1 :

R1 Router ospf 1 Neighbor 1.1.1.2 priority 0 Neighbor 1.1.1.3 priority 0
--

Verify OSPF neighbors:

R1 # show ip ospf neighbors

The output displays neighbor state attempt and then full state when completely established

Attempt state is when neighbors are statically configured.

If we use broadcast option, the neighbors are dynamically detected as with the broadcast option the pvc allows broadcast.

Default mode of ospf on a point-to-multipoint frame-relay physical interface:

R2#show ip ospf interface S0

```
Serial0 is up, line protocol is up
Internet Address 1.1.1.2/8, Area 0
Process ID 1, Router ID 20.1.1.1, Network Type NON_BROADCAST, Cost: 64
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 20.1.1.1, Interface address 1.1.1.2
No backup designated router on this network
Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
```

R2#show ip route

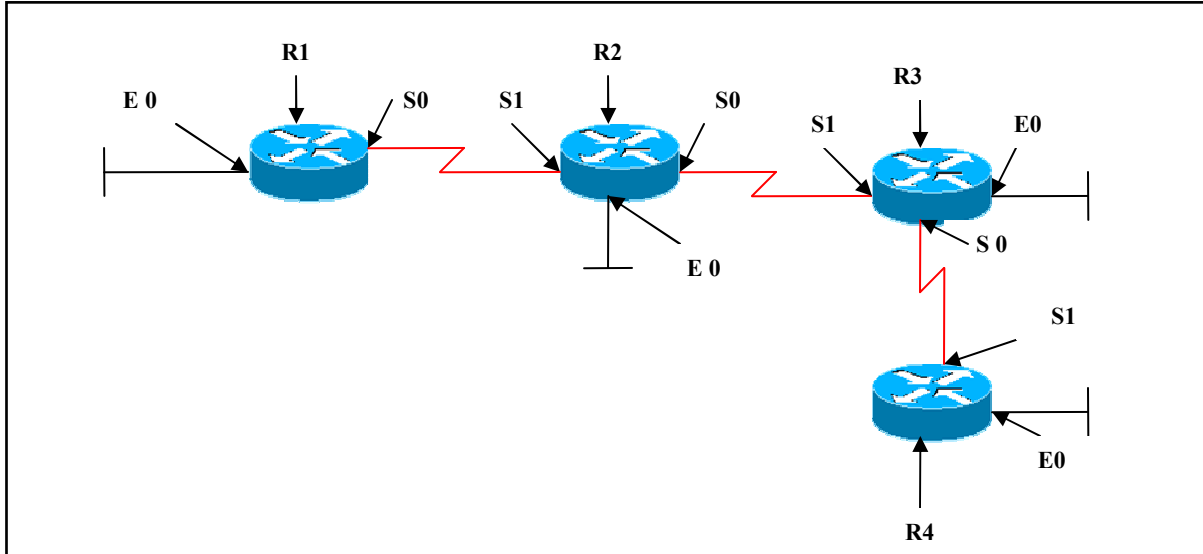
```
C 1.0.0.0/8 is directly connected, Serial0
O 2.0.0.0/8 [110/128] via 1.1.1.1, 00:00:00, Serial0
C 20.0.0.0/8 is directly connected, Ethernet0
O 10.0.0.0/8 [110/74] via 1.1.1.1, 00:00:00, Serial0
O 30.0.0.0/8 [110/138] via 1.1.1.1, 00:00:00, Serial0
```

Module 3 – RIP

RIP LAB INDEX

- 1. CONFIGURE PASSIVE INTERFACE IN RIP**
- 2. CONFIGURE AUTHENTICATION IN RIP**
- 3. MANIPULATING RIP METRICS USING OFFSET-LIST**
- 4. ROUTE FILTERING USING DISTRIBUTE-LIST**
- 5. ROUTE FILTERING USING PREFIX-LIST**

Lab 1 – Configure Passive Interface in RIP



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	192.168.1.2	255.255.255.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	192.168.1.1	255.255.255.0
S 0	172.168.1.1	255.255.255.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	172.168.1.2	255.255.255.0
S 0	172.168.2.1	255.255.255.0
E 0	30.1.1.1	255.0.0.0

R4

Interface	IP Address	Subnet Mask
S 1	172.168.2.2	255.255.255.0
E 0	40.1.1.1.	255.0.0.0

Lab Objective:

Task 1

Configure RIP on all the routers. Do not advertise network 172.168.1.1 in RIP process on R2.

The requirement in the above scenario is to stop RIP broadcasts from R3 being sent to R2. To accomplish this task configure passive-interface on R3 (S1 interface).

R2	R3
Router rip Network 20.0.0.0 Network 192.168.0.0	Router rip Network 172.168.0.0 Passive-interface s1 Network 30.0.0.0

Verification :

R3#debug ip rip

RIP protocol debugging is on

00:42:09: RIP: received v1 update from 172.168.2.2 on Serial0


00:42:09: 40.0.0.0 in 1 hops

00:42:33: RIP: sending v1 update to 255.255.255.255 via Serial0 (172.168.2.1)

00:42:33: subnet 172.168.1.0, metric 1 no updated are sending via s1

The output displays updates received from 172.168.2.2 on S 0 and updates sent via S 0 (172.168.2.1) but does not send updates via S 1 (172.168.1.2).

Lab 2 – Configure Authentication in RIP



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure MD5 Authentication between R1 and R2 using a password of **cisco123**.

R1	R2
<pre> Int s0 Ip rip authentication mode md5 Ip rip authentication key-chain chain1 Key chain chain1 Key 1 Key-string cisco123 </pre>	<pre> Int s1 Ip rip authentication mode md5 Ip rip authentication key-chain chain Key chain chain2 Key 1 Key-string cisco123 </pre>

Verification :

R2#show ip protocol

Routing Protocol is "rip"

Sending updates every 30 seconds, next due in 4 seconds

Invalid after 180 seconds, hold down 180, flushed after 240

Default version control: send version 2, receive version 2

Interface	Send	Recv	Triggered RIP	Key-chain
-----------	------	------	---------------	-----------

Ethernet0	2	2		
-----------	---	---	--	--

Serial1	2	2		chain2
---------	---	---	--	--------

---output omitted---

R2#debug ip rip

01:03:53: RIP: sending v2 update to 224.0.0.9 via Serial1 (1.1.1.2)

01:03:53: RIP: build update entries

01:03:53: 20.0.0.0/8 via 0.0.0.0, metric 1, tag 0

Authentication mismatch in R1 & R2:

R2#debug ip rip events

01:38:27: RIP: sending v2 update to 224.0.0.9 via Ethernet0 (20.1.1.1)

01:38:27: RIP: Update contains 1 routes

01:38:27: RIP: Update queued

01:38:27: RIP: sending v2 update to 224.0.0.9 via Serial1 (1.1.1.2)

01:38:27: RIP: Update contains 1 routes

01:38:27: RIP: Update queued


01:38:27: RIP: Update sent via Ethernet0

01:38:27: RIP: Update sent via Serial1

01:38:28: RIP: ignored v2 packet from 1.1.1.1 (invalid authentication)

If key-identifier and key-string does not match on neighbor end, then there will be an error message stating invalid authentication and updates will not be sent to peers.

Lab 3 – Manipulating RIP Metrics using Offset-List



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S 0	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	2.2.2.2	255.0.0.0
S 0	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1

RIP metric is hop-count.

R1 reaches 30.0.0.0 network via 3.3.3.2, but the requirement is R1 should reach 30.0.0.0 network via R2 .

R3 reaches 10.0.0.0 network via 3.3.3.1, but the requirement is R3 should reach 10.0.0.0 network via R2 .

R1	R3
Access-list 30 permit 30.0.0.0 0.255.255.255	Access-list 10 permit 10.0.0.0 0.255.255.255
Router rip	Router rip
Network 10.0.0.0	Network 30.0.0.0
Network 1.0.0.0	Network 2.0.0.0
Network 3.0.0.0	Network 3.0.0.0
Version 2	Version 2
No auto-summary	No auto-summary
Offset-list 30 in 2 serial 1	Offset-list 10 in 2 serial 0

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
R 2.0.0.0/8 [120/1] via 1.1.1.2, 00:00:19, Serial0
  [120/1] via 3.3.3.2, 00:00:22, Serial1
C 3.0.0.0/8 is directly connected, Serial1
R 20.0.0.0/8 [120/1] via 1.1.1.2, 00:00:19, Serial0
C 10.0.0.0/8 is directly connected, Ethernet0
R 30.0.0.0/8 [120/2] via 1.1.1.2, 00:00:19, Serial0
```

R3#show ip route


```
R 1.0.0.0/8 [120/1] via 3.3.3.1, 00:00:04, Serial0
  [120/1] via 2.2.2.1, 00:00:26, Serial1
C 2.0.0.0/8 is directly connected, Serial1
C 3.0.0.0/8 is directly connected, Serial0
R 20.0.0.0/8 [120/1] via 2.2.2.1, 00:00:27, Serial1
R 10.0.0.0/8 [120/2] via 2.2.2.1, 00:00:27, Serial1
C 30.0.0.0/8 is directly connected, Ethernet0
```

R3#debug ip rip

```
02:08:39: RIP: received v2 update from 3.3.3.1 on Serial0
02:08:39: 1.0.0.0/8 -> 0.0.0.0 in 1 hops
02:08:39: 10.0.0.0/8 -> 0.0.0.0 in 3 hops
02:08:39: 20.0.0.0/8 -> 0.0.0.0 in 2 hops
02:08:39: 30.0.0.0/8 -> 0.0.0.0 in 3 hops
```

The offset-list value '2' is added to the default-metric (1) and applied to incoming routes via S1 in the R1 and S0 in R3. Thus, the routes travel via R2, as that being least in hop count.

Lab 4 – Route filtering using Distribute-List



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S 0	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	192.168.0.1	255.255.255.0
Loopback 1	192.168.1.1	255.255.255.0
Loopback 2	192.168.2.1	255.255.255.0
Loopback 3	192.168.3.1	255.255.255.0
Loopback 4	192.168.4.1	255.255.255.0
Loopback 5	192.168.5.1	255.255.255.0
Loopback 6	192.168.6.1	255.255.255.0
Loopback 7	192.168.7.1	255.255.255.0

Lab Objective:

Task 1

Configure RIP on all the routers as per the scenario . The Requirement is to block networks belonging to 192.168.0.0/22 (192.168.0.0, 192.168.1.0,192.168.2.0, 192.168.3.0) to R1 from R3 using Distribute-List.

R2

```
Access-list 16 deny 192.168.0.0 0.0.3.255  
Access-list 16 permit any
```

```
Router rip  
Distribute-list 16 out serial 1
```

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0  
R 2.0.0.0/8 [120/1] via 1.1.1.2, 00:00:09, Serial0  
R 20.0.0.0/8 [120/1] via 1.1.1.2, 00:00:09, Serial0  
R 192.168.4.0/24 [120/2] via 1.1.1.2, 00:00:09, Serial0  
R 192.168.5.0/24 [120/2] via 1.1.1.2, 00:00:09, Serial0  
C 10.0.0.0/8 is directly connected, Ethernet0  
R 192.168.6.0/24 [120/2] via 1.1.1.2, 00:00:09, Serial0  
R 192.168.7.0/24 [120/2] via 1.1.1.2, 00:00:09, Serial0  
R 30.0.0.0/8 [120/2] via 1.1.1.2, 00:00:09, Serial0
```

The output does not display 192.168.0.0/22 range of networks, but does not block other routes .

Lab 5 – Route Filtering using Prefix-List

(Scenario Based on Lab 4)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1	1.1.1.2	255.0.0.0
S 0	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	192.168.0.1	255.255.255.0
Loopback 1	192.168.1.1	255.255.255.0
Loopback 2	192.168.2.1	255.255.255.0
Loopback 3	192.168.3.1	255.255.255.0
Loopback 4	192.168.4.1	255.255.255.0
Loopback 5	192.168.5.1	255.255.255.0
Loopback 6	192.168.6.1	255.255.255.0
Loopback 7	192.168.7.1	255.255.255.0

Lab Objective:

Task 1

Configure RIP on all the routers as per the scenario . The Requirement is to block networks belonging to 192.168.0.0/22 (192.168.0.0, 192.168.1.0,192.168.2.0, 192.168.3.0) to R1 from R3 using Prefix-List.

R2

```
Ip prefix-list ccnp seq 5 deny 192.168.0.0/22 ge 24 le 24  
Ip prefix-list ccnp seq 10 permit 0.0.0.0/0 le 32
```

```
Router rip  
Distribute-list prefix ccnp out serial 1
```

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0  
R 2.0.0.0/8 [120/1] via 1.1.1.2, 00:00:13, Serial0  
R 20.0.0.0/8 [120/1] via 1.1.1.2, 00:00:13, Serial0  
R 192.168.4.0/24 [120/2] via 1.1.1.2, 00:00:13, Serial0  
R 192.168.5.0/24 [120/2] via 1.1.1.2, 00:00:13, Serial0  
C 10.0.0.0/8 is directly connected, Ethernet0  
R 192.168.6.0/24 [120/2] via 1.1.1.2, 00:00:13, Serial0  
R 192.168.7.0/24 [120/2] via 1.1.1.2, 00:00:13, Serial0  
R 30.0.0.0/8 [120/2] via 1.1.1.2, 00:00:13, Serial0
```

The output does not display 192.168.0.0/22 range of networks, but does not block other routes.

Module 4 – IS-IS

Page 96 of 315

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com


All contents are copyright © 2007-2010

All rights reserved.

IS-IS LAB INDEX

- 1. CONFIGURE IS-IS IN SINGLE AREA**
- 2. CONFIGURE IS-IS IN MULTIPLE AREA**
- 3. CONFIGURE IS-IS ROUTE SUMMARIZATION**

Lab 1– Configure IS-IS in Single Area



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 5/0	2.2.2.1	255.0.0.0
Fa 0/0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.2	255.0.0.0
E 0/0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure IS-IS in single area on both the routers.

R1	R2
Int fa0/0 Ip router isis Isis circuit-type level-1 Int s5/0 Ip router isis Isis circuit-type level-1 Router isis Net 49.0001.0000.0000.0001.00	Int e0/0 Ip router isis Isis circuit-type level-1 Int s1/0 Ip router isis Isis circuit-type level-1 Router isis Net 49.0001.0000.0000.0002.00

Is-type level-1	Is-type level-1
-----------------	-----------------

Verification:

R1#show ip route

```
C 2.0.0.0/8 is directly connected, Serial5/0
i L1 20.0.0.0/8 [115/20] via 2.2.2.2, Serial5/0
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```

The output displays that network 20.0.0.0 is the route from level 1 as indicated by the “i L1” tag.

R2#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0
C 20.0.0.0/8 is directly connected, Ethernet0/0
i L1 10.0.0.0/8 [115/20] via 2.2.2.1, Serial1/0
```

R1#show ip protocol

```
Routing Protocol is "isis"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    None
  Maximum path: 4
  Routing for Networks:
    FastEthernet0/0
    Serial5/0
  Routing Information Sources:
    Gateway         Distance   Last Update
    2.2.2.2          115       00:04:17
  Distance: (default is 115)
```

R1#show isis topology

```
IS-IS paths to level-1 routers
System Id      Metric  Next-Hop      Interface  SNPA
R1             --
0000.0000.0002  10     0000.0000.0002  Se5/0     *HDLC*
```

This command displays the level 1 topology table, which shows the least cost IS – IS paths to the IS’s.

Lab 2 – Configure IS-IS in Multiple Areas

(Scenario Based on Lab 1)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 5/0	2.2.2.1	255.0.0.0
Fa 0/0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.2	255.0.0.0
E 0/0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure IS-IS in multiple areas. Configure R1 (S5/0, FA0/0) in ISIS Area 1 and R2 (S1/0, E0/0) in ISIS Area 2.

R1	R2
Int fa0/0 Ip router isis Isis circuit-type level-1 Int s5/0 Ip router isis Isis circuit-type level-1-2 Router isis Net 49.0001.0000.0000.0001.00 Is-type level-1-2	Int e0/0 Ip router isis Isis circuit-type level-1 Int s1/0 Ip router isis Isis circuit-type level-1-2 Router isis Net 49.0002.0000.0000.0002.00 Is-type level-1-2

Verification:

R1#show ip route


```
C 2.0.0.0/8 is directly connected, Serial5/0  
i L2 20.0.0.0/8 [115/20] via 2.2.2.2, Serial5/0  
C 10.0.0.0/8 is directly connected, FastEthernet0/0
```

The output displays that network 20.0.0.0 is the route from level 2 as indicated by the “i L2” tag.

R2#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0  
C 20.0.0.0/8 is directly connected, Ethernet0/0  
i L2 10.0.0.0/8 [115/20] via 2.2.2.1, Serial1/0
```

Lab 3 – Configure IS-IS Summarization



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 5/0	2.2.2.1	255.0.0.0
Fa 0/0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.2	255.0.0.0
E 0/0	20.1.1.1	255.0.0.0
Loopback 0	192.168.0.1	255.255.255.0
Loopback 1	192.168.1.1	255.255.255.0
Loopback 2	192.168.2.1	255.255.255.0
Loopback 3	192.168.3.1	255.255.255.0

Lab Objective:

Task 1

Configure IS-IS in multiple areas. Configure R1 (S5/0, FA0/0) in ISIS Area 1 and R2 (S1/0, E0/0) in ISIS Area 2.

Configure Loopbacks on R2 and only summarized route should be sent to R1.

R1

```
Int fa0/0
Ip router isis
```

Isis circuit-type level-1

Int s5/0

Ip router isis

Isis circuit-type level-1-2

Router isis

Net 49.0001.0000.0000.0001.00

Is-type level-1-2

R2

Int e0/0

Ip router isis

Isis circuit-type level-1

Int s1/0

Ip router isis

Isis circuit-type level-1-2

Int Loopback 0

Ip address 192.168.0.1 255.255.255.0

Ip router isis

Isis circuit-type level-1

Int Loopback 1

Ip address 192.168.1.1 255.255.255.0

Ip router isis

Isis circuit-type level-1

Int Loopback 2

Ip address 192.168.2.1 255.255.255.0

Ip router isis

Isis circuit-type level-1

Int Loopback 3

Ip address 192.168.3.1 255.255.255.0

Ip router isis

Isis circuit-type level-1

Router isis

Net 49.0002.0000.0000.0002.00

Is-type level-1-2

Summary-address 192.168.0.0 255.255.252.0 level-2

Verification:

R1#show ip route

```
C 2.0.0.0/8 is directly connected, Serial5/0  
i L2 20.0.0.0/8 [115/20] via 2.2.2.2, Serial5/0  
C 10.0.0.0/8 is directly connected, FastEthernet0/0  
i L2 192.168.0.0/22 [115/20] via 2.2.2.2, Serial5/0
```

The output displays a reduced routing table which displays only the summarized route of loopback addresses from level-2 as indicated by “i L2” tag.

Module 5 – BGP

Page 105 of 315

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com

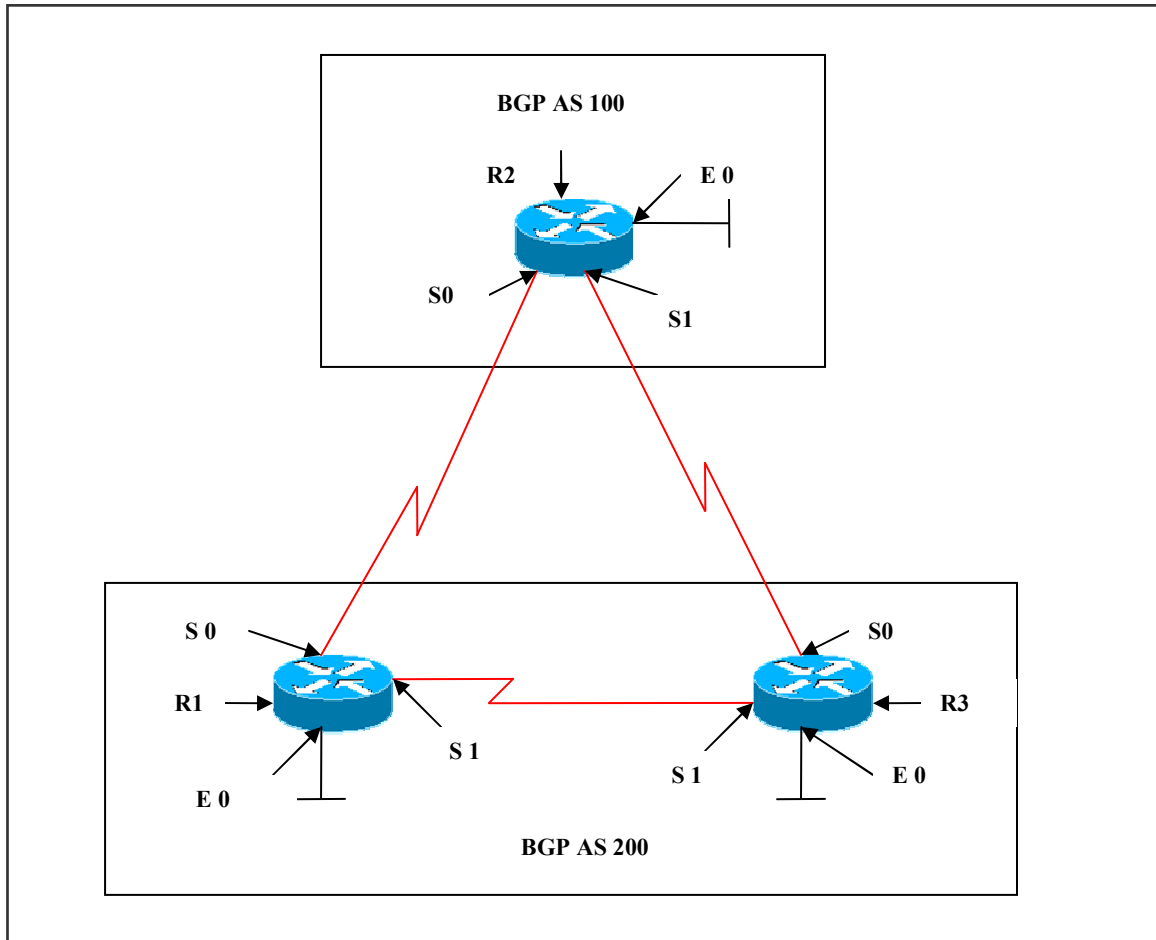
All contents are copyright @ 2007-2010

All rights reserved.

BGP LAB INDEX

- 1. BGP BASIC CONFIGURATION**
- 2. BGP USING LOOPBACK ADDRESS**
- 3. eBGP WITH MULTIHOP COMMAND**
- 4. eBGP WITH MULTIHOP COMMAND (LOAD BALANCING)**
- 5. BGP NEXT-HOP ATTRIBUTE**
- 6. ORIGIN ATTRIBUTE**
- 7. WEIGHT ATTRIBUTE**
- 8. LOCAL PREFERENCE**
- 9. CONFIGURING MED ATTRIBUTE USING DEFAULT-METRIC COMMAND**
- 10. MED ATTRIBUTE**
- 11. COMMUNITY ATTRIBUTE**
- 12. AS-PATH ATTRIBUTE**
- 13. AUTHENTICATION IN BGP**
- 14. CONFIGURING PEER-GROUP**
- 15. ROUTE AGGREGATION IN BGP**
- 16. ROUTE REFLECTOR**
- 17. BGP CONFEDERATION**

Lab 1 – Basic BGP Configuration



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure a BGP neighbor relationship between R1, R2 and R3. R1 should be in AS 200, R2 should be in AS 100 and R3 should be in AS 200.

R1 Router bgp 200 Neighbor 1.1.1.2 remote-as 100 Neighbor 3.3.3.2 remote-as 200 Network 1.0.0.0 Network 3.0.0.0 Network 10.0.0.0 No synchronization	R2 Router bgp 100 Neighbor 1.1.1.1 remote-as 200 Neighbor 2.2.2.2 remote-as 200 Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 No synchronization
R3 Router bgp 200 Neighbor 3.3.3.1 remote-as 200 Neighbor 2.2.2.1 remote-as 100 Network 2.0.0.0 Network 3.0.0.0 Network 30.0.0.0 No synchronization	

Verification:

R1#show ip bgp summary

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	state/PfxRcd
1.1.1.2	4	100	10	12	8	0	0	00:04:54	3
3.3.3.2	4	200	12	11	8	0	0	00:06:57	4

The output displays that BGP neighbors have established a TCP connection.

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
B 2.0.0.0/8 [200/0] via 3.3.3.2, 00:08:31
C 3.0.0.0/8 is directly connected, Serial0/2/1
B 20.0.0.0/8 [20/0] via 1.1.1.2, 00:07:17
C 10.0.0.0/8 is directly connected, FastEthernet0/0
B 30.0.0.0/8 [200/0] via 3.3.3.2, 00:08:31
```

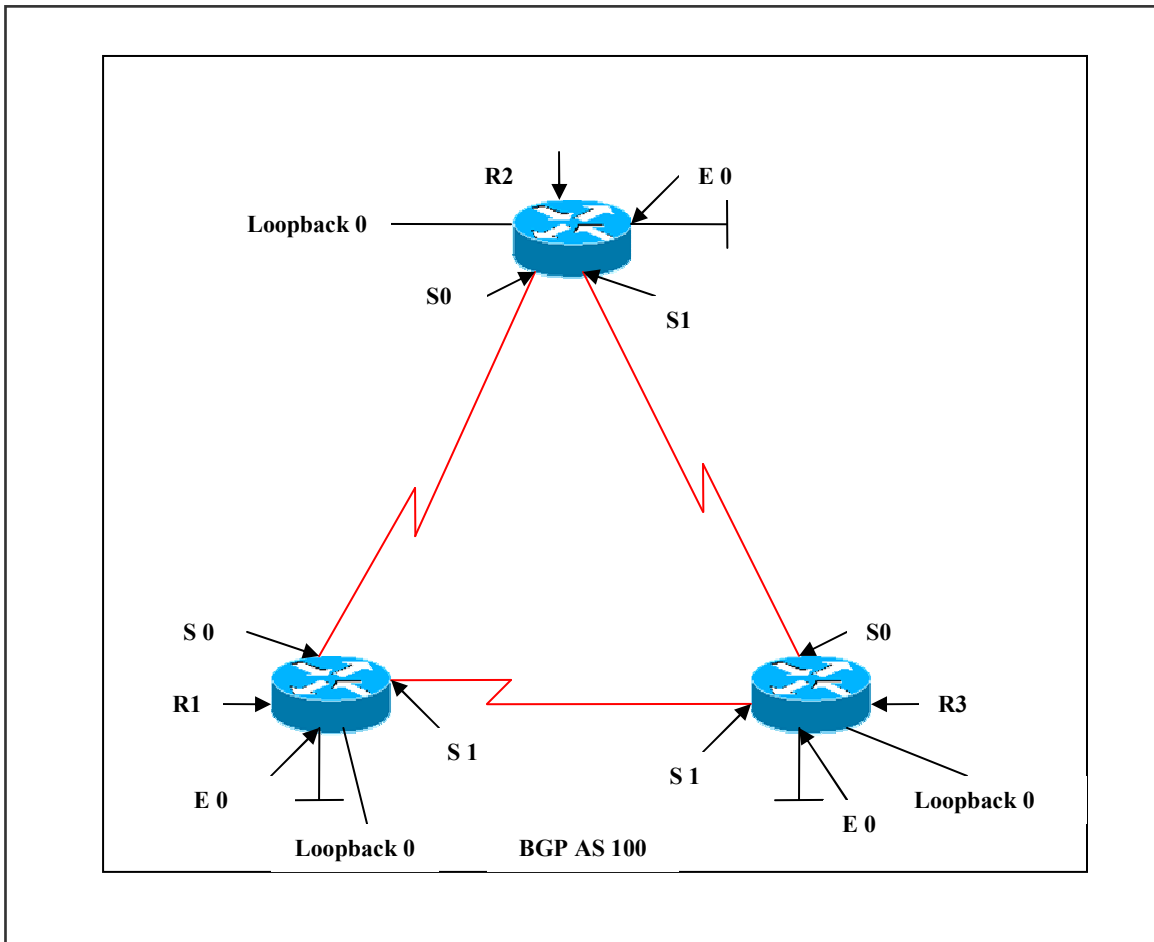
The output states that the BGP routes denoted as 'B' in the routing table.

R1#show ip bgp

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.0.0.0	1.1.1.2	0		0	100 i
*>	0.0.0.0	0		32768	i
* 2.0.0.0	1.1.1.2	0		0	100 i
*>i	3.3.3.2	0	100	0	i
* i3.0.0.0	3.3.3.2	0	100	0	i
*>	0.0.0.0	0		32768	i
*> 10.0.0.0	0.0.0.0	0		32768	i
*> 20.0.0.0	1.1.1.2	0		0	100 i
* i	2.2.2.1	0	100	0	100 i
*>i30.0.0.0	3.3.3.2	0	100	0	i

The output displays the BGP table.

Lab 2 – Connecting BGP using Loopback



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0
Loopback 0	50.50.50.50	255.255.255.255

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0
Loopback 0	75.75.75.75	255.255.255.255

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	100.100.100.100	255.255.255.255

Lab Objective:

Task 1

Configure a BGP neighbor relationship between R1, R2 and R3. All routers should be configured in AS 100. Establish the neighbor relationship based on Loopback 0 addresses. Configure EIGRP as the routing protocol in AS 100. Advertise all loopback networks under EIGRP.

R1	R2
Router eigrp 100 Network 1.0.0.0 Network 3.0.0.0 Network 10.0.0.0 Network 50.0.0.0 No auto-summary Router bgp 100 Neighbor 75.75.75.75 remote-as 100 Neighbor 75.75.75.75 update-source loopback 0 Neighbor 100.100.100.100 remote-as 100 Neighbor 100.100.100.100 update-source loopback 0 No synchronization	Router eigrp 100 Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 Network 75.0.0.0 No auto-summary Router bgp 100 Neighbor 50.50.50.50 remote-as 100 Neighbor 50.50.50.50 update-source loopback 0 Neighbor 100.100.100.100 remote-as 100 Neighbor 100.100.100.100 update-source loopback 0 No synchronization

R3

```
Router eigrp 100
Network 2.0.0.0
Network 3.0.0.0
Network 30.0.0.0
Network 100.0.0.0
No auto-summary
```

```
Router bgp 100
Neighbor 50.50.50.50 remote-as 100
Neighbor 50.50.50.50 update-source loopback 0
Neighbor 75.75.75.75 remote-as 100
Neighbor 75.75.75.75 update-source loopback 0
No synchronization
```

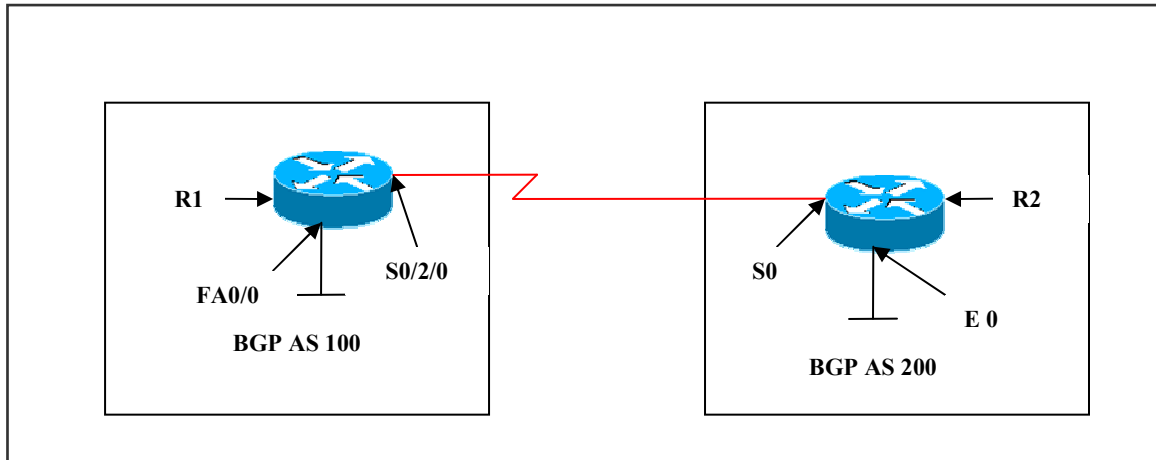
Verification:

R1#show ip bgp summary

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
75.75.75.75		4 100	10	10	1	0	0	00:06:48	0
100.100.100.100		4 100	10	10	1	0	0	00:06:50	0

The output displays that neighbors established a TCP connection between them.

Lab 3 – ebgp-Multihop



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1 :

Configure a BGP neighbor relationship between R1 and R2. R1 should be configured in AS 100 and R2 should be in AS 200. Establish the neighbor relationship between peers that are not directly connected. You are allowed to create a static route on each router to accomplish this task.

R1	R2
Ip route 20.1.1.1 255.0.0.0 1.1.1.2	Ip route 10.1.1.1 255.0.0.0 1.1.1.1
Router bgp 100	Router bgp 200
Network 1.0.0.0	Network 1.0.0.0
Neighbor 20.1.1.1 remote-as 200	Neighbor 10.1.1.1 remote-as 100
Neighbor 20.1.1.1 ebgp-multihop	Neighbor 10.1.1.1 ebgp-multihop
No synchronization	No synchronization

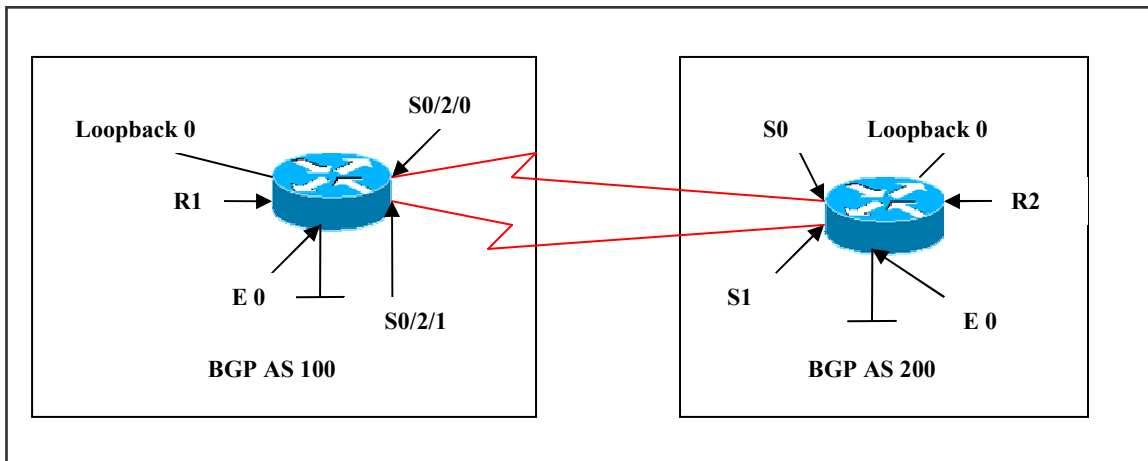
Verification:

R1#show ip bgp summary

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
20.1.1.1 4 200 18 18 1 0 0 00:08:23 0
```

The output displays neighborship as established.

Lab 4 – ebgp-Multihop (Load Balancing)



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0/2/0	1.1.1.1	255.0.0.0
S 0/2/1	2.2.2.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0
Loopback 0	50.50.50.50	255.255.255.255

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0
Loopback 0	75.75.75.75	255.255.255.255

Lab Objective:

Task 1:

Configure a BGP neighbor relationship between R1 and R2. R1 should be configured in AS 100 and R2 should be in AS 200. Establish the neighbor relationship between peers using loopbacks. Create a static route on each router to accomplish this task.

R1	R2
Ip route 75.75.75.75 255.255.255.255 1.1.1.2 Ip route 75.75.75.75 255.255.255.255 2.2.2.2	Ip route 50.50.50.50 255.255.255.255 1.1.1.1 Ip route 50.50.50.50 255.255.255.255 2.2.2.1
Router bgp 100 Neighbor 75.75.75.75 remote-as 200 Neighbor 75.75.75.75 update-source loopback 0 Neighbor 75.75.75.75 ebgp-multihop Network 1.0.0.0 Network 2.0.0.0 No synchronization	Router bgp 200 Neighbor 50.50.50.50 remote-as 100 Neighbor 50.50.50.50 update-source loopback 0 Neighbor 50.50.50.50 ebgp-multihop Network 1.0.0.0 Network 2.0.0.0 No synchronization

Verification:

R1#show ip bgp summary

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
75.75.75.75 4 200 16 16 1 0 0 00:12:01 0
```

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0/2/0
  50.0.0.0/32 is subnetted, 1 subnets
C 50.50.50.50 is directly connected, Loopback0
C 2.0.0.0/8 is directly connected, Serial0/2/1
C 10.0.0.0/8 is directly connected, FastEthernet0/0
  75.0.0.0/32 is subnetted, 1 subnets
S 75.75.75.75 [1/0] via 2.2.2.2
  [1/0] via 1.1.1.2
```

Routing table displays two choices to reach the next hop 75.75.75.75, one via 2.2.2.2 and the other via 1.1.1.2.

The load balancing can be verified by issuing the traceroute command:

R1#traceroute 75.75.75.75

Type escape sequence to abort.
Tracing the route to 75.75.75.75

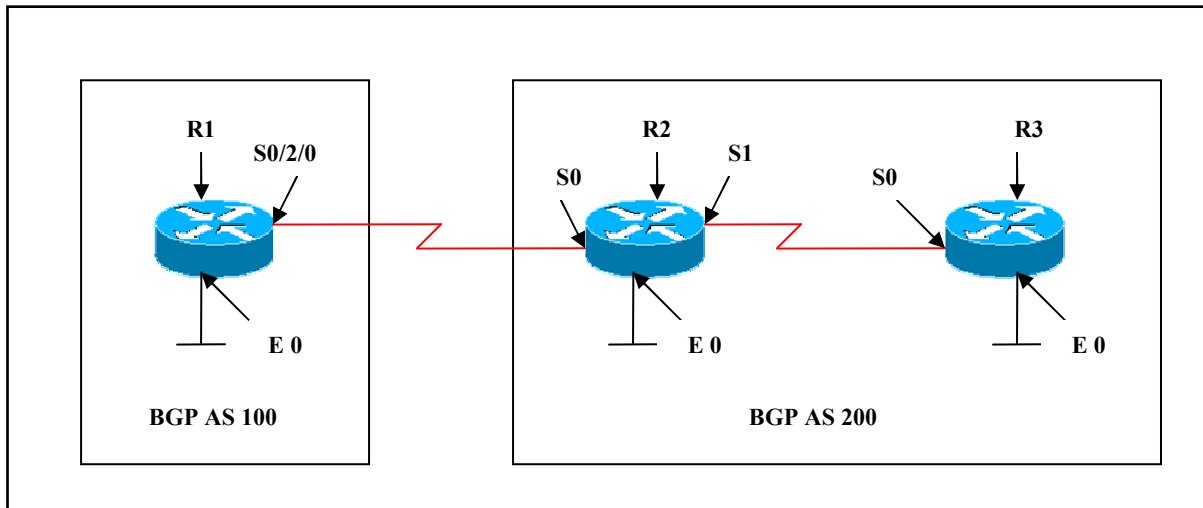
```
1 2.2.2.2 16 msec
  1.1.1.2 16 msec *
```

R1#traceroute 75.75.75.75

Type escape sequence to abort.
Tracing the route to 75.75.75.75

```
1 1.1.1.2 24 msec
  2.2.2.2 16 msec
```

Lab 5 – BGP Next Hop Attribute



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure BGP on the routers as per the above scenario.

R1 Router bgp 100 Neighbor 1.1.1.2 remote-as 200 Network 1.0.0.0 Network 10.0.0.0 No synchronization	R2 Router bgp 200 Neighbor 2.2.2.2 remote-as 200 Neighbor 1.1.1.1 remote-as 100 Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 No synchronization
R3 Router bgp 200 Neighbor 2.2.2.1 remote-as 200 Network 2.0.0.0 Network 30.0.0.0 No synchronization	

R1 advertises network 10.0.0.0 to R2 with the next hop of 1.1.1.1 and R2 advertises network 20.0.0.0 to R1 with the next hop of 1.1.1.2.

For iBGP, the protocol states that the next hop that eBGP advertises, should be carried into iBGP, because of this rule, R2 advertises network 10.0.0.0 to its iBGP peer R3 with a next hop of 1.1.1.1. Therefore for R3, the next hop to reach network 10.0.0.0 is via 1.1.1.1 and not 2.2.2.1.

Make sure that R3 can reach network 10.0.0.0 via IGP, otherwise R3 drops packets with the destination of 10.0.0.0 or advertise via network commands in BGP.

Verification:

R3#show ip route

```
B 1.0.0.0/8 [200/0] via 2.2.2.1, 00:08:24
C 2.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [200/0] via 2.2.2.1, 00:09:11
```

```
B 10.0.0.0/8 [200/0] via 1.1.1.1, 00:08:11
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output on R3 displays that the network 10.0.0.0 can be reached via 1.1.1.1.

Task 2:

Configure BGP such that R2 advertises its updates to iBGP peers via 2.2.2.1 instead of 1.1.1.1. You can use the next-hop-self command to accomplish this task.

R2

```
Router bgp 200
Neighbor 2.2.2.2 remote-as 200
Neighbor 2.2.2.2 next-hop-self
Neighbor 1.1.1.1 remote-as 100
Neighbor 1.1.1.1 next-hop-self
Network 1.0.0.0
Network 2.0.0.0
Network 20.0.0.0
No synchronization
```

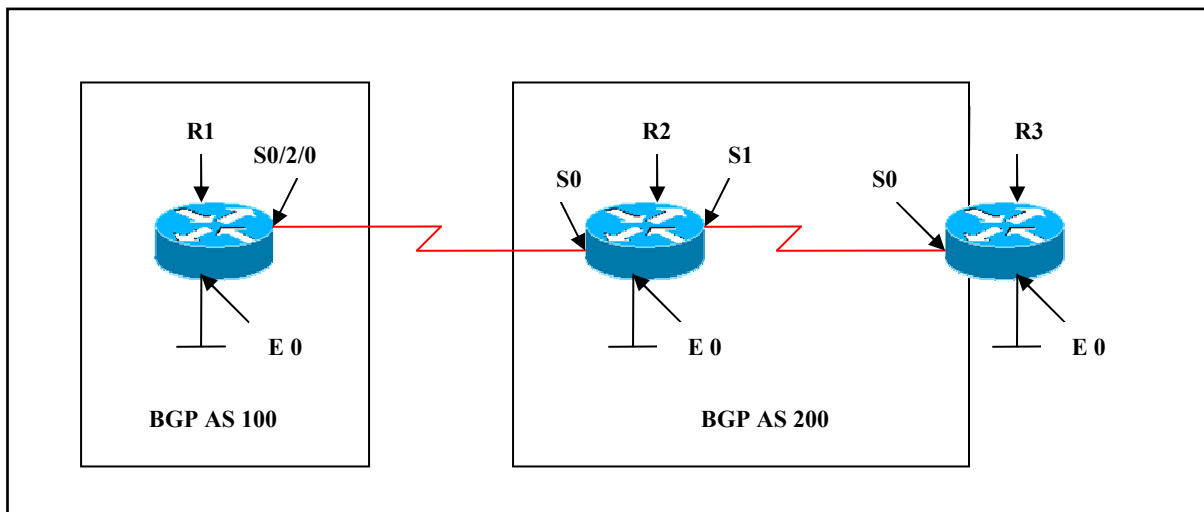
Verification:

R3#show ip route

```
B 1.0.0.0/8 [200/0] via 2.2.2.1, 00:00:05
C 2.0.0.0/8 is directly connected, Serial0
B 10.0.0.0/8 [200/0] via 2.2.2.1, 00:00:05
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays that for R3 to reach network 10.0.0.0 is via 2.2.2.1 because of the next-hop-self command. R2 advertises network 10.0.0.0 via 2.2.2.1 to R3, instead of carrying the next-hop advertised by eBGP.

Lab 6 – Origin Attribute



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0/2/0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure BGP on all the three routers. Do not advertise network 30.0.0.0 on R3 in BGP, instead create static route on R2 to reach 30.0.0.0 via 2.2.2.2 and redistribute this static route into BGP.

R1 Router bgp 100 Neighbor 1.1.1.2 remote-as 200 Network 1.0.0.0 Network 10.0.0.0 No synchronization	R3 Router bgp 200 Neighbor 2.2.2.1 remote-as 200 Network 2.0.0.0 No synchronization
R2 ip route 30.1.1.1 255.0.0.0 2.2.2.2 Router bgp 200 Neighbor 2.2.2.2 remote-as 200 Neighbor 1.1.1.1 remote-as 100 Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 Redistribute static No synchronization	

Verification:

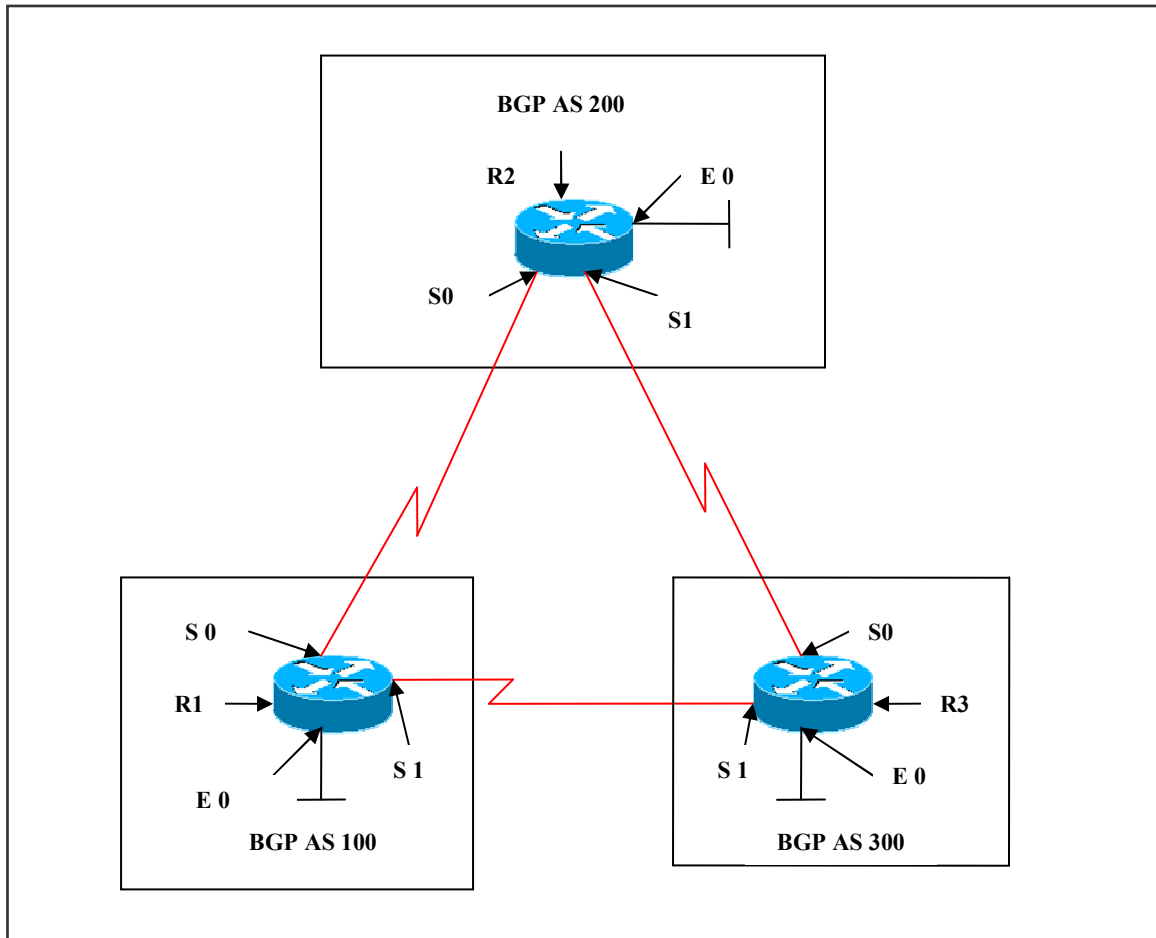
R1#show ip bgp

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.0.0.0	1.1.1.2	0		0	200 i
*>	0.0.0.0	0		32768	i
*> 2.0.0.0	1.1.1.2	0		0	200 i
*> 10.0.0.0	0.0.0.0	0		32768	i
*> 20.0.0.0	1.1.1.2	0		0	200 i
*> 30.1.1.0/24	1.1.1.2	0		0	200 ?

R1 reaches 2.0.0.0 via '200 i' means that the next AS path is 200 and the origin of the route is IGP.

R1 also reaches 30.0.0.0 via '200 ?', means that the next AS is 200 and that the origin is incomplete and is a redistributed static route.

Lab 7 – Setting Cisco Weight Attribute



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure AS 200 such that all traffic destined for network 3.0.0.0 should go through R3. Use the Weight attribute to accomplish this task.

R1 Router bgp 100 Neighbor 1.1.1.2 remote-as 200 Neighbor 3.3.3.2 remote-as 300 Network 1.0.0.0 Network 3.0.0.0 Network 10.0.0.0 No synchronization	R3 Router bgp 300 Neighbor 2.2.2.1 remote-as 200 Neighbor 3.3.3.1 remote-as 100 Network 2.0.0.0 Network 3.0.0.0 Network 30.0.0.0 No synchronization
R2 Router bgp 200 Neighbor 1.1.1.1 remote-as 100 Neighbor 1.1.1.1 weight 500 Neighbor 2.2.2.2 remote-as 300 Neighbor 2.2.2.2 weight 1000 Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 Redistribute static	

No synchronization

Verification:

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
B 3.0.0.0/8 [20/0] via 2.2.2.2, 00:00:04
C 20.0.0.0/8 is directly connected, Ethernet0
B 10.0.0.0/8 [20/0] via 2.2.2.2, 00:00:04
B 30.0.0.0/8 [20/0] via 2.2.2.2, 00:00:04
```

The output displays that R2 has been forced to use R3 as the next-hop to reach network 3.0.0.0

R2#show ip bgp 3.0.0.0

```
BGP routing table entry for 3.0.0.0/8, version 4
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non-peer-group peers:
    1.1.1.1
    100
    1.1.1.1 from 1.1.1.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, weight 500, valid, external
    300
    2.2.2.2 from 2.2.2.2 (30.1.1.1)
      Origin IGP, metric 0, localpref 100, weight 1000, valid, external, best
```

The output displays two paths and shows the path via 2.2.2.2 as the best path chosen because of the highest weight set to that path.

R2#show ip bgp 10.0.0.0

```
BGP routing table entry for 10.0.0.0/8, version 5
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non-peer-group peers:
    1.1.1.1
    100
    1.1.1.1 from 1.1.1.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, weight 500, valid, external
    300 100
    2.2.2.2 from 2.2.2.2 (30.1.1.1)
      Origin IGP, localpref 100, weight 1000, valid, external, best
```

The output displays two paths and shows the path via 2.2.2.2 as the best path chosen because of the highest weight set to that path.

Task 2:

Configure route-map using weight attribute to manipulate the routing information on R2.

```
R2

Access-list 1 permit 30.0.0.0 0.255.255.255
Access-list 2 permit 10.0.0.0 0.255.255.255

Route-map list 1 permit 10
Match ip address 1
Set weight 1000

Route-map list 1 permit 20

Route-map list 2 permit 10
Match ip add 2
Set weight 1000

Route-map list 2 permit 20

Router bgp 200
Neighbor 1.1.1.1 route-map list 1 in
Neighbor 2.2.2.2 route-map list 2 in
```

Verification:

R2#sh ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
B 3.0.0.0/8 [20/0] via 1.1.1.1, 00:00:02
C 20.0.0.0/8 is directly connected, Ethernet0
B 10.0.0.0/8 [20/0] via 2.2.2.2, 00:00:02
B 30.0.0.0/8 [20/0] via 1.1.1.1, 00:00:02
```

The routing table displays that R2 is learning network 10.0.0.0 via 2.2.2.2 and network 30.0.0.0 via 1.1.1.1

R2#show ip bgp 30.0.0.0

```
BGP routing table entry for 30.0.0.0/8, version 7
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    2.2.2.2
    300
    2.2.2.2 from 2.2.2.2 (30.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external
    100 300
    1.1.1.1 from 1.1.1.1 (10.1.1.1)
      Origin IGP, localpref 100, weight 1000, valid, external, best
```

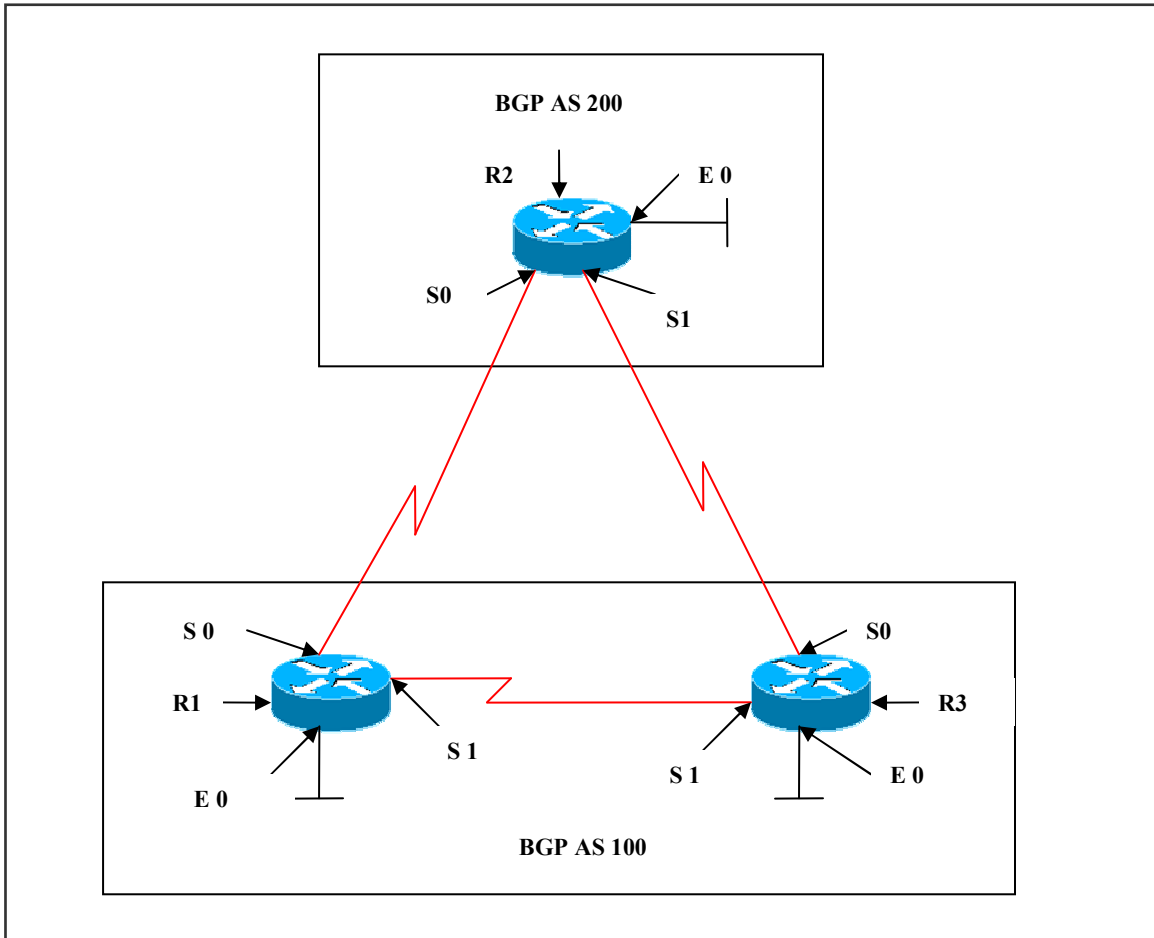
The output displays the best path to reach network 30.0.0.0 from R2 is via 1.1.1.1, because of the highest weight attribute set to that path.

R2#show ip bgp 10.0.0.0

```
BGP routing table entry for 10.0.0.0/8, version 8
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    1.1.1.1
    300 100
    2.2.2.2 from 2.2.2.2 (30.1.1.1)
      Origin IGP, localpref 100, weight 1000, valid, external, best
    100
    1.1.1.1 from 1.1.1.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external
```

The output displays the best path to reach network 10.0.0.0 from R2 is via 2.2.2.2, because of the highest weight attribute set to that path.

Lab 8 – Setting Local Preference



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure AS 100 such that all traffic destined for AS 200 should go through R2. Use Local-Preference Attribute to accomplish this task.

R1 Router bgp 100 Neighbor 1.1.1.2 remote-as 200 Neighbor 3.3.3.2 remote-as 100 Network 1.0.0.0 Network 3.0.0.0 Network 10.0.0.0 No synchronization	R3 Router bgp 100 Neighbor 2.2.2.1 remote-as 200 Neighbor 3.3.3.1 remote-as 100 Bgp default-preference 500 Network 2.0.0.0 Network 3.0.0.0 Network 30.0.0.0 No synchronization
R2 Router bgp 200 Neighbor 1.1.1.1 remote-as 100 Neighbor 2.2.2.2 remote-as 100 Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 No synchronization	

Verification:

R3#show ip route

```
B 1.0.0.0/8 [20/0] via 2.2.2.1, 00:00:01
C 2.0.0.0/8 is directly connected, Serial0
C 3.0.0.0/8 is directly connected, Serial1
B 20.0.0.0/8 [20/0] via 2.2.2.1, 00:00:01
B 10.0.0.0/8 [200/0] via 3.3.3.1, 00:00:01
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output displays that R3 learns network 1.0.0.0 via 2.2.2.1.

R3#show ip bgp 1.0.0.0

```
200
 2.2.2.1 from 2.2.2.1 (20.1.1.1)
   Origin IGP, metric 0, localpref 500, valid, external, best
Local
 3.3.3.1 from 3.3.3.1 (10.1.1.1)
   Origin IGP, metric 0, localpref 100, valid, internal
```

The output displays that path 2.2.2.1 is the best path, because of the highest local preference value over the other path.

Task 2:

Configure route-map using local-preference attribute to manipulate the routing information on R3.

```
R3
Access-list 1 permit 20.0.0.0 0.255.255.255

Route-map list 1 permit 10
Match ip address 1
Set local preference 50

Route-map list 1 permit 20

Router bgp 200
Neighbor 2.2.2.1 remote-as 200
Neighbor 2.2.2.1 route-map list1 in
Neighbor 3.3.3.1 remote-as 100
Network 3.0.0.0
```

```
Network 2.0.0.0
Network 30.0.0.0
No synchronization
```

Verification:

R3#show ip route

```
B 1.0.0.0/8 [200/0] via 3.3.3.1, 00:02:17
C 2.0.0.0/8 is directly connected, Serial0
C 3.0.0.0/8 is directly connected, Serial1
B 20.0.0.0/8 [200/0] via 1.1.1.2, 00:01:41
B 10.0.0.0/8 [200/0] via 3.3.3.1, 00:02:17
C 30.0.0.0/8 is directly connected, Ethernet0
```

R3#show ip bgp 20.0.0.0

```
200
  2.2.2.1 from 2.2.2.1 (20.1.1.1)
    Origin IGP, metric 0, localpref 50, valid, external
200
  1.1.1.2 from 3.3.3.1 (10.1.1.1)
    Origin IGP, metric 0, localpref 100, valid, internal, best
```

The output displays that path 1.1.1.2 is the best path, because of the highest local preference value over the other path.

Lab 10 – Configuring MED

(Scenario Based On Lab 9)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

All ingress (incoming) traffic to AS 200 should use the path through R3 using the MED attribute. Configure the MED on R1 to 100 and Configure the MED on R3 to 50. Lower MED will be preferred.

R1	R3
Access-list 1 permit 10.0.0.0 0.255.255.255	Access-list 1 permit 30.0.0.0 0.255.255.255
Route-map list 1 permit 10 Match ip add 1 Set metric 200	Route-map list 1 permit 10 Match ip add 1 Set metric 50
Route-map list 1 permit 20 Set metric 100	Route-map list 1 permit 20 Set metric 100
Router bgp 100 Neighbor 1.1.1.2 route-map list 1 out	Router bgp 100 Neighbor 2.2.2.1 route-map list 1 out

Verification:

R2#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
C 2.0.0.0/8 is directly connected, Serial1
B 3.0.0.0/8 [20/100] via 1.1.1.1, 00:00:04
C 20.0.0.0/8 is directly connected, Ethernet0
B 10.0.0.0/8 [20/100] via 2.2.2.2, 00:00:04
B 30.0.0.0/8 [20/50] via 2.2.2.2, 00:00:04
```

The output displays that network 10.0.0.0 & 30.0.0.0 are learnt via 2.2.2.2 because of the lowest MED value set to this path.

R2#show ip bgp 10.0.0.0

```
BGP routing table entry for 10.0.0.0/8, version 5
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non-peer-group peers:
    1.1.1.1
    100
    2.2.2.2 from 2.2.2.2 (30.1.1.1)
      Origin IGP, metric 100, localpref 100, valid, external, best
    100
    1.1.1.1 from 1.1.1.1 (10.1.1.1)
      Origin IGP, metric 200, localpref 100, valid, external
```

R2#show ip bgp 30.0.0.0

BGP routing table entry for 30.0.0.0/8, version 7

Paths: (2 available, best #1, table Default-IP-Routing-Table)

Advertised to non peer-group peers:

1.1.1.1

100

2.2.2.2 from 2.2.2.2 (30.1.1.1)

Origin IGP, metric 50, localpref 100, valid, external, best

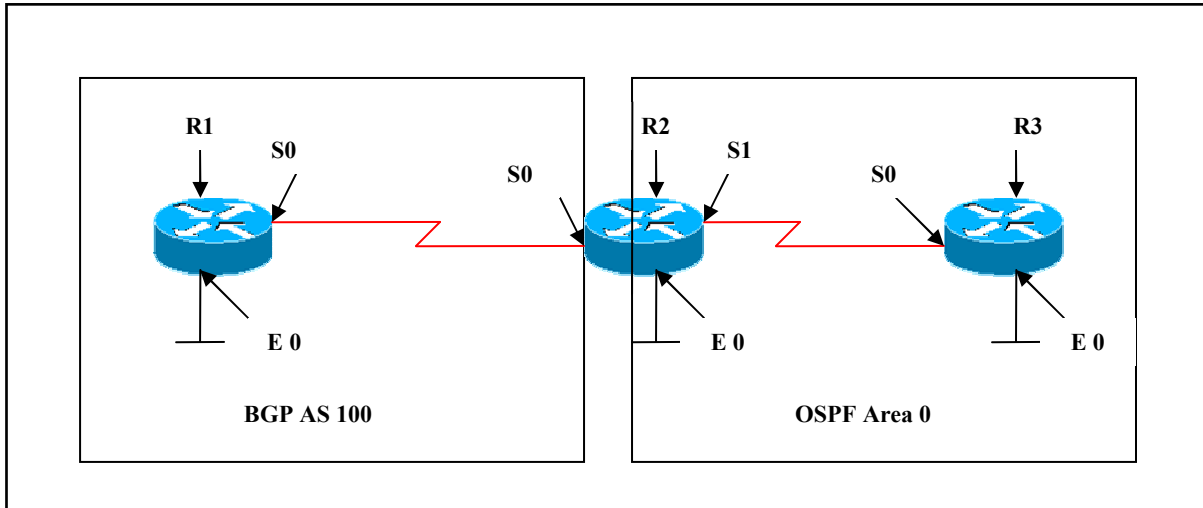
100

1.1.1.1 from 1.1.1.1 (10.1.1.1)

Origin IGP, metric 100, localpref 100, valid, external

The output displays best path 2.2.2.2 with a metric 50 lower than other path.

Lab 11 – Configuring MED using default-metric command



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure BGP and OSPF as per the above scenario. Redistribute OSPF into BGP and verify the metric values displayed in the output by default.

R1 Router bgp 100 Neighbor 1.1.1.2 remote-as 100 Network 1.0.0.0 Network 10.0.0.0 No synchronization	R2 Router ospf 1 Network 2.2.2.1 0.0.0.0 area 0 Network 20.0.0.0 0.255.255.255 area 0 Router bgp 100 Neighbor 1.1.1.1 remote-as 100 Network 1.0.0.0 No synchronization Redistribute ospf 1
R3 Router ospf 1 Network 2.2.2.1 0.0.0.0 area 0 Network 30.0.0.0 0.255.255.255 area 0	

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0  
B 2.0.0.0/8 [200/0] via 1.1.1.2, 00:03:10  
B 20.0.0.0/8 [200/0] via 1.1.1.2, 00:03:10  
C 10.0.0.0/8 is directly connected, Ethernet0  
B 30.0.0.0/8 [200/74] via 2.2.2.2, 00:02:35
```

The output displays routes 2.0.0.0 and 20.0.0.0 with a metric of '0' as they are directly connected to R2 and when passed to R1 travel with a metric of '0'.

Task 2:

Configure BGP and OSPF as per the above scenario. Redistribute OSPF into BGP using a metric value of 5.

R2

```
Router ospf 1
Network 2.2.2.1 0.0.0.0 area 0
Network 20.0.0.0 0.255.255.255 area 0
```

```
Router bgp 100
Neighbor 1.1.1.1 remote-as 100
Network 1.0.0.0
No synchronization
Redistribute ospf 1 metric 5
```

Verification:**R1#show ip route**

```
C 1.0.0.0/8 is directly connected, Serial0
B 2.0.0.0/8 [200/0] via 1.1.1.2, 00:02:39
B 20.0.0.0/8 [200/0] via 1.1.1.2, 00:02:39
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [200/5] via 2.2.2.2, 00:02:07
```

The output displays network 30.0.0.0 changed to metric of 5. But observe that network 2.0.0.0 and 20.0.0.0 still remain with a metric of '0' as they are not displayed as redistributed routes instead they are learn't as connected routes on R2.

Task 3:

Configure BGP and OSPF as per the above scenario. Redistribute OSPF into BGP using a metric value of 5 and also redistribute connected routes with a metric set to 50.

R2

```
Router ospf 1
Network 2.2.2.1 0.0.0.0 area 0
Network 20.0.0.0 0.255.255.255 area 0
```

```
Router bgp 100
Neighbor 1.1.1.1 remote-as 100
Network 1.0.0.0
No synchronization
Redistribute ospf 1 metric 5
Redistribute connected metric 50
```

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
B 2.0.0.0/8 [200/50] via 1.1.1.2, 00:02:39
B 20.0.0.0/8 [200/50] via 1.1.1.2, 00:02:39
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [200/5] via 2.2.2.2, 00:02:07
```

The output displays network 30.0.0.0 changed to metric of '5'. Also network 2.0.0.0 and 20.0.0.0 with a metric of '50'.

Task 4:

Configure BGP and OSPF as per the above scenario. Redistribute OSPF into BGP and use the default-metric command to change the metric.

R2

```
Router ospf 1
Network 2.2.2.1 0.0.0.0 area 0
Network 20.0.0.0 0.255.255.255 area 0

Router bgp 100
Neighbor 1.1.1.1 remote-as 100
Network 1.0.0.0
No synchronization
Redistribute ospf 1
Default-metric 75
```

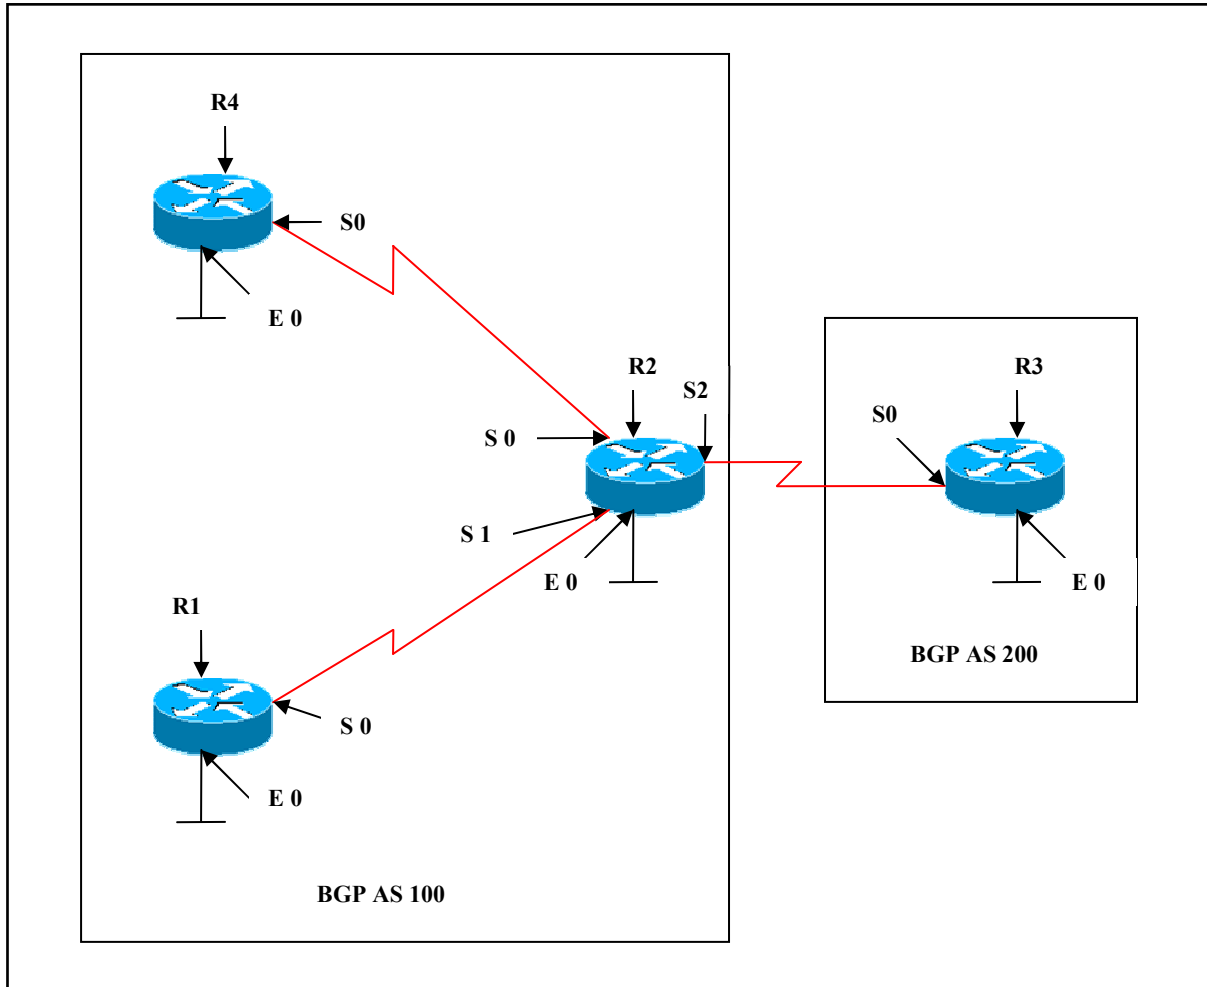
Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
B 2.0.0.0/8 [200/0] via 1.1.1.2, 00:00:33
B 20.0.0.0/8 [200/0] via 1.1.1.2, 00:00:33
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [200/75] via 2.2.2.2, 00:00:00
```

The output displays network 30.0.0.0 with a metric value changed to 75.

Lab 12 – Community Attribute



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	3.3.3.1	255.0.0.0
S 1	1.1.1.2	255.0.0.0
S 2	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

R4

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
E 0	40.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure BGP on all the routers. Configure R1, R2 and R4 in AS 100. Configure R3 in AS 200. Network 10.0.0.0 should not be sent outside AS 100 using no-export community attribute.

R1

```
Access-list 1 permit 10.0.0.0  
0.255.255.255
```

```
Route-map no-exp  
Match ip add 1  
Set community no-export
```

```
Router bgp 100  
Neighbor 1.1.1.2 route-map no-exp out  
Neighbor 1.1.1.2 send-community
```

Verification:

R3#show ip route

```
B 1.0.0.0/8 [20/0] via 2.2.2.1, 00:01:36
C 2.0.0.0/8 is directly connected, Serial0
B 3.0.0.0/8 [20/0] via 2.2.2.1, 00:01:36
B 20.0.0.0/8 [20/0] via 2.2.2.1, 00:01:36
B 40.0.0.0/8 [20/0] via 2.2.2.1, 00:00:13
C 30.0.0.0/8 is directly connected, Ethernet0
```

The output does not display network 10.0.0.0 in the routing table as it is blocked by the no-export community attribute.

R2#show ip bgp 10.0.0.0

```
BGP routing table entry for 10.0.0.0/8, version 13
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to
EBGP peer)
  Not advertised to any peer
  Local
    1.1.1.1 from 1.1.1.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Community: no-export
```

The output displays the community attribute no-export set .

Task 2:

Configure BGP on all the routers. Configure R1, R2 and R4 in AS 100. Configure R3 in AS 200. Network 10.0.0.0 should not be advertised to any peers, internal or external.

R1

```
Access-list 1 permit 10.0.0.0
0.255.255.255

Route-map no-adv
Match ip add 1
Set community no-advertise

Router bgp 100
Neighbor 1.1.1.2 route-map no-adv out
Neighbor 1.1.1.2 send-community
```

Verification:

R3#show ip route

```
B 1.0.0.0/8 [20/0] via 2.2.2.1, 00:01:18
C 2.0.0.0/8 is directly connected, Serial0
B 3.0.0.0/8 [20/0] via 2.2.2.1, 00:01:18
B 20.0.0.0/8 [20/0] via 2.2.2.1, 00:01:18
B 40.0.0.0/8 [20/0] via 2.2.2.1, 00:00:21
C 30.0.0.0/8 is directly connected, Ethernet0
```

R4#show ip route

```
B 1.0.0.0/8 [200/0] via 3.3.3.1, 00:00:21
B 2.0.0.0/8 [200/0] via 3.3.3.1, 00:01:22
C 3.0.0.0/8 is directly connected, Serial1
B 20.0.0.0/8 [200/0] via 3.3.3.1, 00:01:22
C 40.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [200/0] via 2.2.2.2, 00:00:22
```

The output doesn't display network 10.0.0.0 in the routing table of both ibgp and ebgp neighbors.

R2#show ip bgp 10.0.0.0

BGP routing table entry for 10.0.0.0/8, version 18

Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised to any peer)

Not advertised to any peer

Local

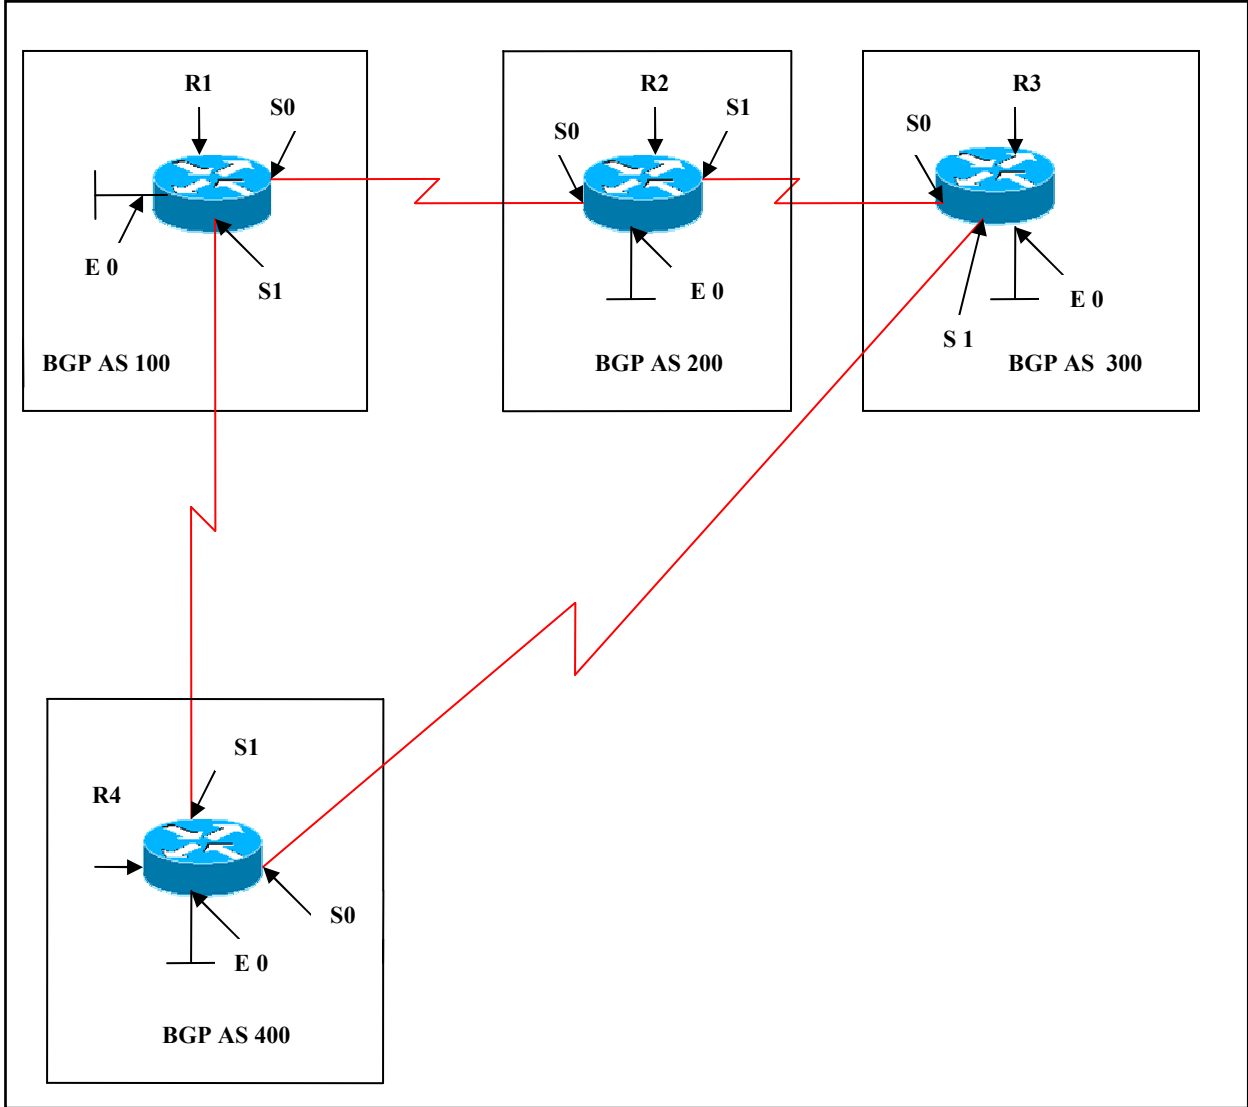
1.1.1.1 from 1.1.1.1 (10.1.1.1)

Origin IGP, metric 0, localpref 100, valid, internal, best

Community: no-advertise

The output displays the community attribute no-advertise set .

Lab 13 – AS-Path Attribute



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
S 1	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	4.4.4.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

R4

Interface	IP Address	Subnet Mask
S 1	3.3.3.2	255.0.0.0
S 0	4.4.4.1	255.0.0.0
E 0	40.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure BGP on all the routers.

R1	R2
Router bgp 100 Neighbor 1.1.1.2 remote-as 200 Neighbor 3.3.3.2 remote-as 400	Router bgp 200 Neighbor 2.2.2.2 remote-as 300 Neighbor 1.1.1.1 remote-as 100

Network 1.0.0.0 Network 3.0.0.0 Network 10.0.0.0 No synchronization	Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 No synchronization
R3 Router bgp 300 Neighbor 2.2.2.1 remote-as 200 Neighbor 4.4.4.1 remote-as 400 Network 2.0.0.0 Network 4.0.0.0 Network 30.0.0.0 No synchronization	R4 Router bgp 400 Neighbor 4.4.4.2 remote-as 300 Neighbor 3.3.3.1 remote-as 100 Network 3.0.0.0 Network 4.0.0.0 Network 40.0.0.0 No synchronization

Verification:

R4#show ip route

```

B 1.0.0.0/8 [20/0] via 3.3.3.1, 00:01:15
B 2.0.0.0/8 [20/0] via 4.4.4.2, 00:01:15
C 3.0.0.0/8 is directly connected, Serial1
C 4.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [20/0] via 3.3.3.1, 00:01:15
C 40.0.0.0/8 is directly connected, Ethernet0
B 10.0.0.0/8 [20/0] via 3.3.3.1, 00:01:15
B 30.0.0.0/8 [20/0] via 4.4.4.2, 00:01:15

```

R4#show ip bgp 30.0.0.0

```

BGP routing table entry for 30.0.0.0/8, version 8
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non-peer-group peers:
    3.3.3.1
    300
    4.4.4.2 from 4.4.4.2 (30.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
    100 200 300
    3.3.3.1 from 3.3.3.1 (3.3.3.1)
      Origin IGP, localpref 100, valid, external

```

The output displays that network 30.0.0.0 is reached via 4.4.4.2 from R4 as it is the shortest path when compared to the other path via 3.3.3.1.

Task 2:

Manipulate the path to reach network 30.0.0.0 on R4. You can use as-path prepend command using route-map to accomplish this task.

R4

```
Access-list 1 permit 30.0.0.0
0.255.255.255

Route-map map 1 permit 10
Match ip add 1
Set as-path prepend 400 400 400 400

Route-map map1 permit 20

Router bgp 400
Neighbor 4.4.4.2 route-map map 1 in
```

Verification:

R4#show ip route

```
B 1.0.0.0/8 [20/0] via 3.3.3.1, 00:00:34
B 2.0.0.0/8 [20/0] via 4.4.4.2, 00:00:34
C 3.0.0.0/8 is directly connected, Serial1
C 4.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [20/0] via 3.3.3.1, 00:00:34
C 40.0.0.0/8 is directly connected, Ethernet0
B 10.0.0.0/8 [20/0] via 3.3.3.1, 00:00:34
B 30.0.0.0/8 [20/0] via 3.3.3.1, 00:00:34
```

R4#show ip bgp 30.0.0.0

```
BGP routing table entry for 30.0.0.0/8, version 8
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non-peer-group peers:
    4.4.4.2
    400 400 400 400 300
```

4.4.4.2 from 4.4.4.2 (30.1.1.1)

Origin IGP, metric 0, localpref 100, valid, external

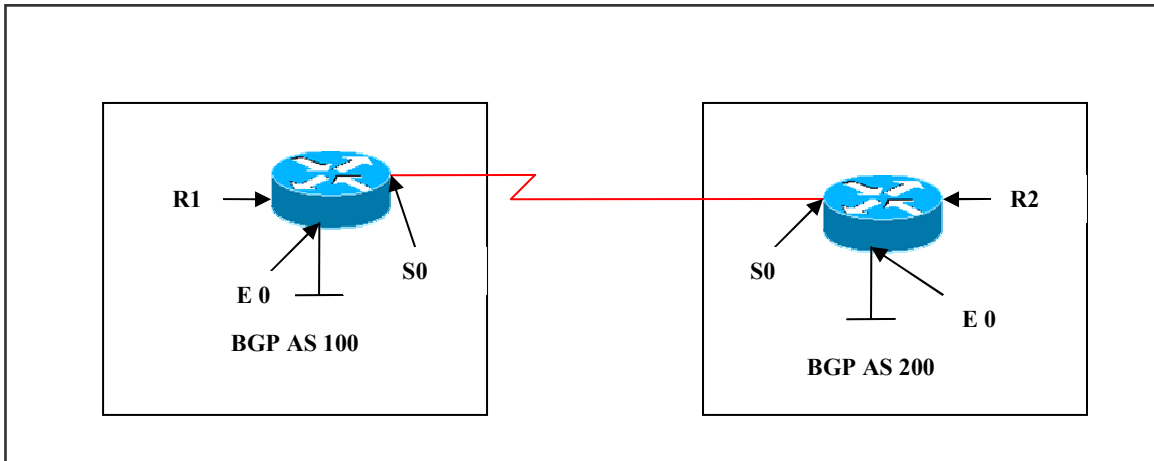
100 200 300

3.3.3.1 from 3.3.3.1 (3.3.3.1)

Origin IGP, localpref 100, valid, external, best

The output displays that network 30.0.0.0 is reached via 3.3.3.1 from R4 as it is the shortest path when compared to the other path via 4.4.4.2.

Lab 14 – BGP Neighbor MD5 Authentication



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure BGP on all the routers. Configure R1 in AS 100 and R2 in AS 200. Configure MD5 Authentication between R1 and R2 using a password of **cisco123**.

R1	R2
Router bgp 100 Neighbor 1.1.1.2 remote-as 200	Router bgp 200 Neighbor 1.1.1.1 remote-as 100

Neighbor 1.1.1.2 password cisco123 Network 1.0.0.0 Network 10.0.0.0 No synchronization	Neighbor 1.1.1.1 password cisco123 Network 1.0.0.0 Network 20.0.0.0 No synchronization
---	---

Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [20/0] via 1.1.1.2, 00:00:04
C 10.0.0.0/8 is directly connected, Ethernet0
```

R1#show ip bgp summary

```
Neighbor  V  AS MsgRcvd MsgSent  TblVer  InQ  OutQ Up/Down  State/PfxRcd
1.1.1.2   4  200   13     11     4    0    0  00:01:43    2
```

Authentication in R2 and no authentication in R1:

R2#debug ip bgp events

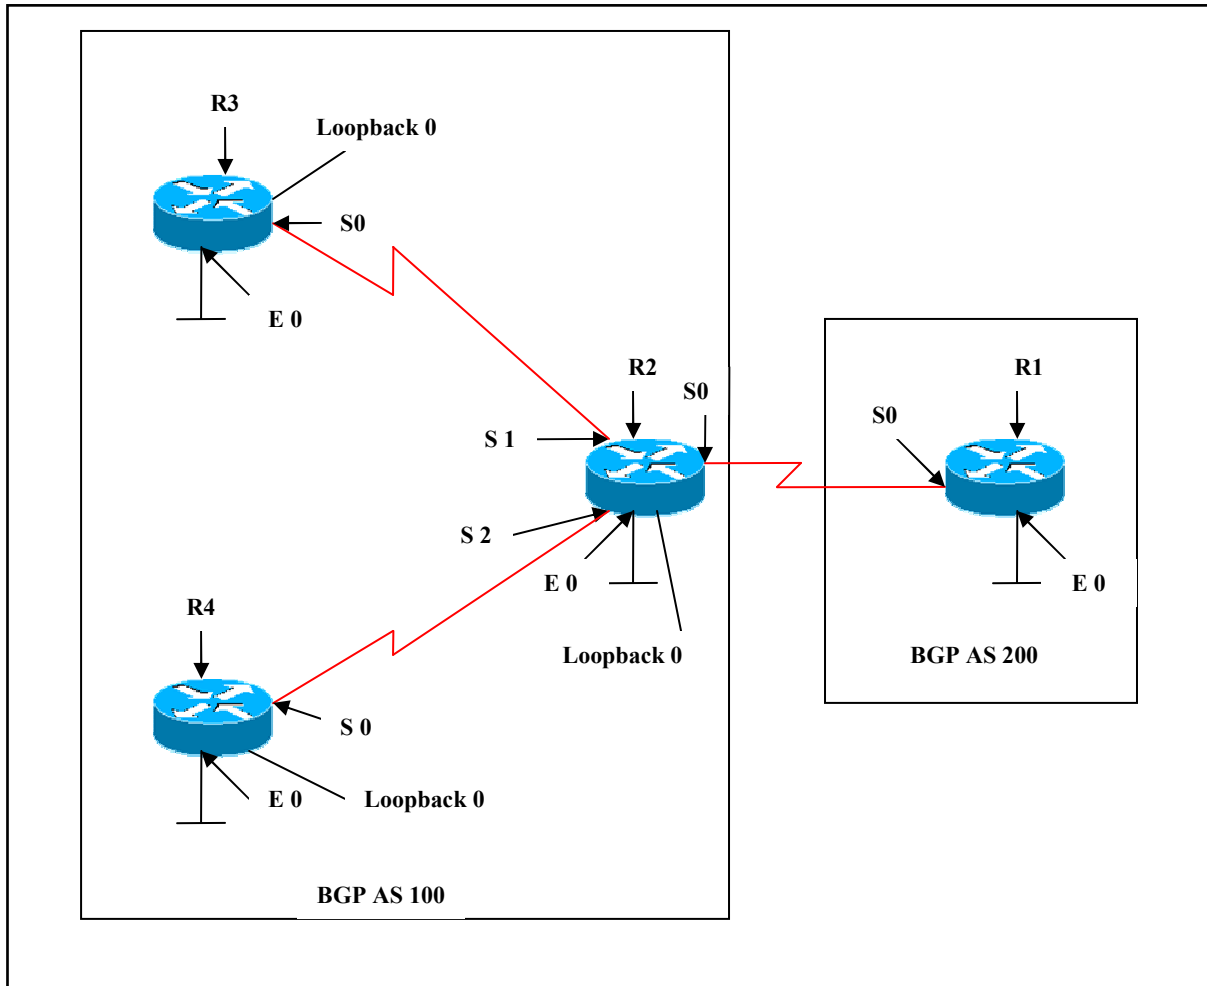
```
04:58:02: %TCP-6-BADAUTH: No MD5 digest from 1.1.1.1:179 to 1.1.1.2:11087
04:58:04: %TCP-6-BADAUTH: No MD5 digest from 1.1.1.1:179 to 1.1.1.2:11087
04:58:04: %TCP-6-BADAUTH: No MD5 digest from 1.1.1.1:179 to 1.1.1.2:11087
```

Authentication mismatch:

R2#debug ip bgp events

```
05:01:09: %TCP-6-BADAUTH: Invalid MD5 digest from 1.1.1.1:11040 to 1.1.1.2:179
05:01:12: %TCP-6-BADAUTH: Invalid MD5 digest from 1.1.1.1:11040 to 1.1.1.2:179
05:01:16: %TCP-6-BADAUTH: Invalid MD5 digest from 1.1.1.1:11040 to 1.1.1.2:179
```

Lab 15 – Configuring Peer-groups



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
S 1	1.1.1.1	255.0.0.0
S 2	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0
Loopback 0	7.7.7.7	255.255.255.255

R3

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 0	8.8.8.8	255.255.255.255

R4

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	40.1.1.1	255.0.0.0
Loopback 0	6.6.6.6	255.255.255.255

Lab Objective:

Task 1:

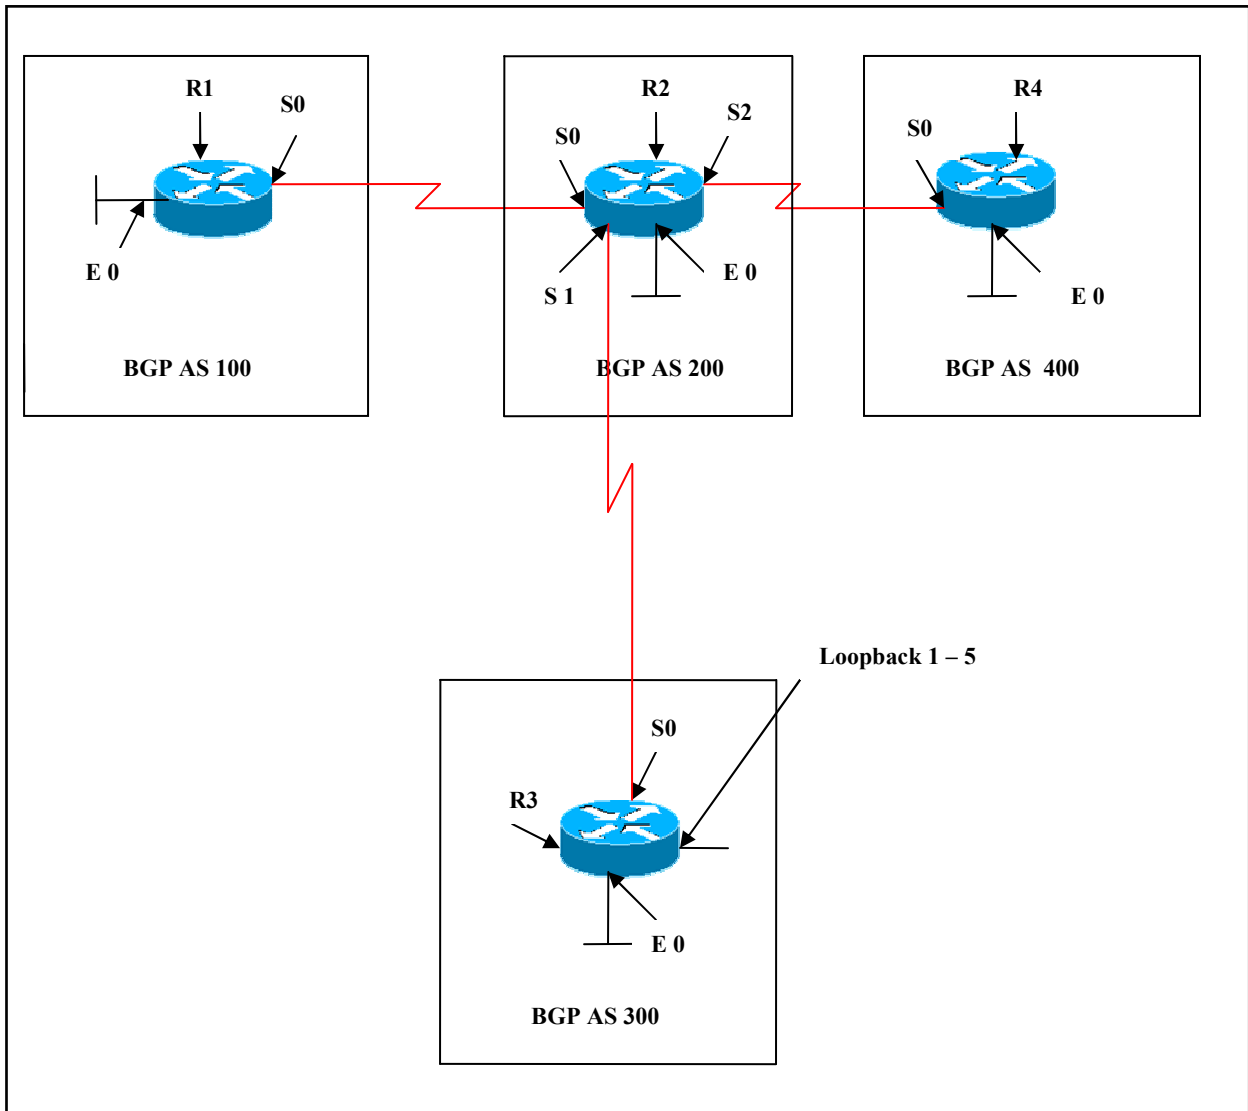
Configure BGP on all the routers. Configure R1 in AS 200 and configure R2, R3 and R4 in AS 100. Configure the loopbacks as per the scenario and advertise in BGP. Also configure route-map blocking network 10.0.0.0 from being advertised to iBGP peers. You can use peer-group to accomplish these tasks.

R1	R2
<pre>ip route 7.7.7.7 255.255.255.255 3.3.3.2 Router bgp 200 Neighbor 7.7.7.7 remote-as 100 Neighbor 7.7.7.7 ebgp-multihop Network 3.0.0.0 Network 10.0.0.0 No synchronization</pre>	<pre>Router bgp 100 Neighbor internal peer-group Neighbor internal remote-AS 100 Neighbor internal update-source loopback 0 Neighbor internal route-map map1 out Neighbor 8.8.8.8 peer-group internal Neighbor 6.6.6.6 peer-group internal</pre>

R3 Router bgp 100 Neighbor internal peer-group Neighbor internal remote-AS 100 Neighbor internal update-source loopback 0 Neighbor 7.7.7.7 peer-group internal Neighbor 6.6.6.6 peer-group internal	R4 Router bgp 100 Neighbor internal peer-group Neighbor internal remote-AS 100 Neighbor internal update-source loopback 0 Neighbor 7.7.7.7 peer-group internal Neighbor 8.8.8.8 peer-group internal
--	--

Configuring BGP using peer-group simplifies configuration reducing the number of statements in the configuration.

Lab 16 – Route Aggregation



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	3.3.3.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
S 1	1.1.1.1	255.0.0.0
S 2	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0
Loopback 1	172.1.0.1	255.255.0.0
Loopback 2	172.2.0.1	255.255.0.0
Loopback 3	172.3.0.1	255.255.0.0
Loopback 4	172.4.0.1	255.255.0.0
Loopback 5	172.5.0.1	255.255.0.0

R4

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	40.1.1.1	255.0.0.0

Lab Objective:

Task 1 :

Configure BGP on all the routers. Create loopbacks on R3 as per the above scenario and advertise them under BGP.

Loopback 1 – 172.1.0.1/16
Loopback 2 – 172.2.0.1/16
Loopback 3 – 172.3.0.1/16
Loopback 4 – 172.4.0.1/16
Loopback 5 – 172.5.0.1/16

R3

```
interface Loopback1
ip address 172.1.0.1 255.255.0.0
```

```
interface Loopback2
ip address 172.2.0.1 255.255.0.0
```

```
interface Loopback3
ip address 172.3.0.1 255.255.0.0
```

```
interface Loopback4
ip address 172.4.0.1 255.255.0.0
```

```
interface Loopback4
ip address 172.5.0.1 255.255.0.0
```

```
Router BGP 300
Network 172.1.0.0
Network 172.2.0.0
Network 172.3.0.0
Network 172.4.0.0
Network 172.5.0.0
```

Task 2

Configure Route Aggregation on R3 such that these routes are summarized as a single route.

R3

```
Router bgp 300
Aggregate-address 172.0.0.0 255.248.0.0
```

Verification :

R1#show ip route

```
B 1.0.0.0/8 [20/0] via 3.3.3.2, 00:15:39
B 2.0.0.0/8 [20/0] via 3.3.3.2, 00:16:38
C 3.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [20/0] via 3.3.3.2, 00:15:39
B 172.1.0.0/16 [20/0] via 3.3.3.2, 00:08:03
B 172.2.0.0/16 [20/0] via 3.3.3.2, 00:08:03
B 172.3.0.0/16 [20/0] via 3.3.3.2, 00:08:03
B 172.4.0.0/16 [20/0] via 3.3.3.2, 00:08:03
B 172.5.0.0/16 [20/0] via 3.3.3.2, 00:07:03
B 40.0.0.0/8 [20/0] via 3.3.3.2, 00:16:38
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [20/0] via 3.3.3.2, 00:08:51
B 172.0.0.0/13 [20/0] via 3.3.3.2, 00:00:34
```

The routing table displays the prefix route (172.0.0.0/13) and also all the specific-routes.

Task 3

Configure Route Aggregation on R3 such that these routes are summarized as a single route. Only the Summary route should be send to R3's neighbors

R3

```
Router bgp 300
Aggregate-address 172.0.0.0 255.248.0.0 summary-only
```

Verification:

R1#show ip route

```
B 1.0.0.0/8 [20/0] via 3.3.3.2, 00:20:36
B 2.0.0.0/8 [20/0] via 3.3.3.2, 00:21:35
C 3.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [20/0] via 3.3.3.2, 00:20:36
B 40.0.0.0/8 [20/0] via 3.3.3.2, 00:21:35
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [20/0] via 3.3.3.2, 00:13:48
B 172.0.0.0/13 [20/0] via 3.3.3.2, 00:00:20
```

The output displays only the prefix route (172.0.0.0/13) and suppresses all the specific routes.

Task 4

Configure Route Aggregation on R3 such that these routes are summarized as a single route. Only the Summary route and the 172.1.0.0, 172.2.0.0 and 172.3.0.0 route should be send to R3's neighbor, blocking 172.4.0.0 and 172.5.0.0 routes.

R3

```
Access-list 1 permit 172.4.0.0 0.0.255.255
Access-list 1 permit 172.5.0.0 0.0.255.255
Access-list 1 deny 0.0.0.0 255.255.255.255
```

```
Route-map map 1 permit 10
Match ip address 1
```

```
Router bgp 200
Aggregate-address 172.0.0.0 255.248.0.0 suppress-map map1
```

By definition of suppress-map, the match criteria set to permit will be suppressed and the rest will be forwarded.

Verification:

R1#show ip route

```
B 1.0.0.0/8 [20/0] via 3.3.3.2, 00:00:32
B 2.0.0.0/8 [20/0] via 3.3.3.2, 00:00:32
C 3.0.0.0/8 is directly connected, Serial0
B 20.0.0.0/8 [20/0] via 3.3.3.2, 00:00:32
B 172.1.0.0/16 [20/0] via 3.3.3.2, 00:00:33
B 172.2.0.0/16 [20/0] via 3.3.3.2, 00:00:33
B 172.3.0.0/16 [20/0] via 3.3.3.2, 00:00:33
B 40.0.0.0/8 [20/0] via 3.3.3.2, 00:00:32
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [20/0] via 3.3.3.2, 00:00:33
B 172.0.0.0/13 [20/0] via 3.3.3.2, 00:00:32
```

The output displays only 172.1.0.0/16, 172.2.0.0/16 & 172.3.0.0/16, thus blocking 172.4.0.0 and 172.5.0.0 routes.

R2#show ip bgp

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.0.0.0	1.1.1.2	0		0	300 i
*>	0.0.0.0	0		32768	i
* 2.0.0.0	2.2.2.2	0		0	400 i
*>	0.0.0.0	0		32768	i
* 3.0.0.0	3.3.3.1	0		0	100 i
*>	0.0.0.0	0		32768	i
*> 10.0.0.0	3.3.3.1	0		0	100 i
*> 20.0.0.0	0.0.0.0	0		32768	i
*> 30.0.0.0	1.1.1.2	0		0	300 i
*> 40.0.0.0	2.2.2.2	0		0	400 i
*> 172.0.0.0/13	0.0.0.0			32768	i
*> 172.1.0.0	1.1.1.2	0		0	300 i
*> 172.2.0.0	1.1.1.2	0		0	300 i
*> 172.3.0.0	1.1.1.2	0		0	300 i
s> 172.4.0.0	1.1.1.2	0		0	300 i
s> 172.5.0.0	1.1.1.2	0		0	300 i

The output displays network 172.4.0.0 and 172.5.0.0 as suppressed routes.

Task 5

Configure Route Aggregation on R3 such that these routes are summarized as a single route. Configure route-map and set the attribute origin to the route-map and implement in BGP process such that the aggregate address appears as incomplete route.

R3

```
Route-map map1  
Set origin incomplete
```


```
Router bgp 300  
Aggregate-address 172.0.0.0 255.248.0.0 attribute-map map1
```

Verification:

```
R1#show ip bgp 172.0.0.0  
BGP routing table entry for 172.0.0.0/13, version 33  
Paths: (1 available, best #1)  
 200, (aggregated by 200 20.1.1.1)  
 3.3.3.2 from 3.3.3.2 (20.1.1.1)  
   Origin incomplete, localpref 100, valid, external, atomic-aggregate, best
```

The output displays network 172.0.0.0 as incomplete route.

Lab 17 – Configuring Route Reflectors



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

Lab Objective:

Task 1:

Configure neighbor relationships between R1 and R2 and another one between R2 and R3. Do not configure a neighbor relationship between R1 and R3. Make sure routes from R1 can get propagated to R3.

R1 Router bgp 100 Neighbor 1.1.1.2 remote-as 100 Network 1.0.0.0 Network 10.0.0.0 No synchronization	R3 Router bgp 100 Neighbor 2.2.2.1 remote-as 100 Network 2.0.0.0 Network 30.0.0.0 No synchronization
R2 Router bgp 100 Neighbor 2.2.2.2 remote-as 100 Neighbor 2.2.2.2 route-reflector-client Neighbor 1.1.1.1 remote-as 100 Neighbor 1.1.1.1 route-reflector-client Network 1.0.0.0 Network 2.0.0.0 Network 20.0.0.0 No synchronization	

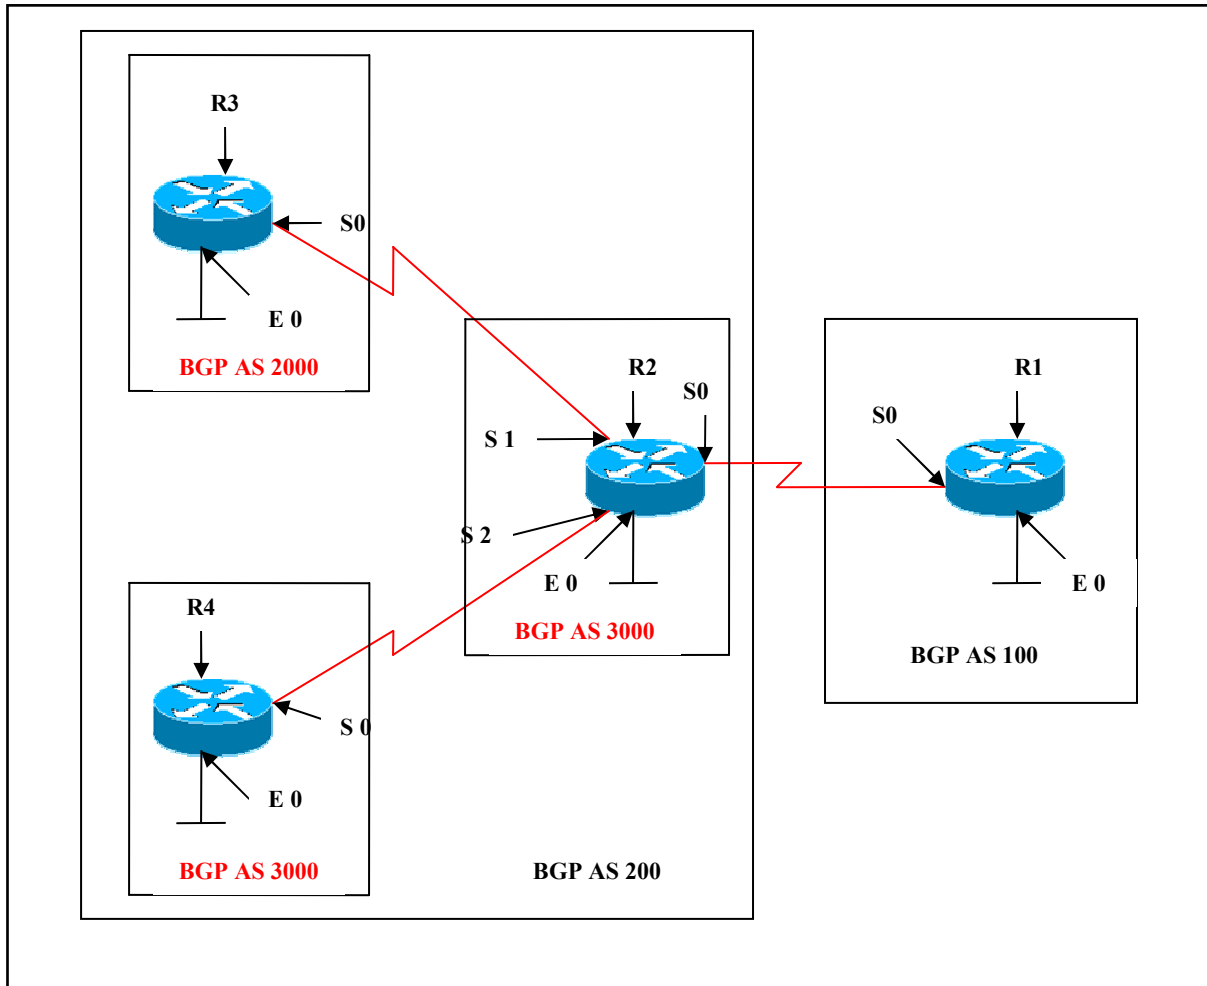
Verification:

R1#show ip route

```
C 1.0.0.0/8 is directly connected, Serial0
B 2.0.0.0/8 [200/0] via 1.1.1.2, 00:00:42
B 20.0.0.0/8 [200/0] via 1.1.1.2, 00:00:42
C 10.0.0.0/8 is directly connected, Ethernet0
B 30.0.0.0/8 [200/0] via 2.2.2.2, 00:00:36
```

If RR was not configured on R2, then the routing table will not display network 30.0.0.0, because of iBGP rule, which states that a BGP speaker will not advertise a route that the BGP speaker learned via another iBGP speaker to a third party iBGP speaker.

Lab 18 – Confederations



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 0	1.1.1.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	1.1.1.2	255.0.0.0
S 1	2.2.2.1	255.0.0.0
S 2	3.3.3.1	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
E 0	30.1.1.1	255.0.0.0

R4

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
E 0	40.1.1.1	255.0.0.0

Lab Objective:

Task 1 :

Configure AS 1000, AS 2000 and AS 3000 are Sub Autonomous Systems of a Larger AS 200 using Confederations. Configure a Neighbor relationship between AS 100 and AS 200 and another Neighbor relationship between AS 1000, AS 2000 and AS 3000.

R1	R2
Router bgp 100 Neighbor 1.1.1.2 remote-as 200 Neighbor 7.7.7.7 ebgp-multihop Network 1.0.0.0 Network 10.0.0.0 No synchronization	Router bgp 1000 Bgp confederation identifier 200 Bgp confederation peers 2000 3000 Neighbor 2.2.2.2 remote-as 2000 Neighbor 3.3.3.2 remote-as 3000 Neighbor 1.1.1.1 remote-as 100 Network 1.0.0.0 Network 2.0.0.0 Network 3.0.0.0 Network 20.0.0.0 No synchronization

R3	R4
Router bgp 2000	Router bgp 3000
Bgp confederation identifier 200	Bgp confederation identifier 200
Bgp confederation peers 1000 3000	Bgp confederation peers 1000 2000
Neighbor 2.2.2.1 remote-as 1000	Neighbor 3.3.3.1 remote-as 1000
Network 2.0.0.0	Network 3.0.0.0
Network 30.0.0.0	Network 40.0.0.0
No synchronization	No synchronization

Verification:

R3#show ip bgp

```

B 1.0.0.0/8 [200/0] via 2.2.2.1, 00:05:48
C 2.0.0.0/8 is directly connected, Serial0
B 3.0.0.0/8 [200/0] via 2.2.2.1, 00:05:48
B 20.0.0.0/8 [200/0] via 2.2.2.1, 00:05:48
B 40.0.0.0/8 [200/0] via 3.3.3.2, 00:05:03
B 10.0.0.0/8 [200/0] via 1.1.1.1, 00:05:03
C 30.0.0.0/8 is directly connected, Ethernet0

```

R3#show ip bgp

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	2.2.2.1	0	100	0	(1000) i
* 2.0.0.0	2.2.2.1	0	100	0	(1000) i
*>	0.0.0.0	0		32768	i
*> 3.0.0.0	2.2.2.1	0	100	0	(1000) i
*> 10.0.0.0	1.1.1.1	0	100	0	(1000) 100 i
*> 20.0.0.0	2.2.2.1	0	100	0	(1000) i
*> 30.0.0.0	0.0.0.0	0		32768	i
*> 40.0.0.0	3.3.3.2	0	100	0	(1000 3000) i

R1#show ip bgp

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.0.0.0	1.1.1.2	0		0	200 i
*>	0.0.0.0	0		32768	i
*> 2.0.0.0	1.1.1.2	0		0	200 i
*> 3.0.0.0	1.1.1.2	0		0	200 i
*> 10.0.0.0	0.0.0.0	0		32768	i
*> 20.0.0.0	1.1.1.2	0		0	200 i
*> 30.0.0.0	1.1.1.2			0	200 i
*> 40.0.0.0	1.1.1.2			0	200 i

The output displays AS-Path 200 whereas on R3 the AS-Path is 1000 2000.

PAPER 2

Switching

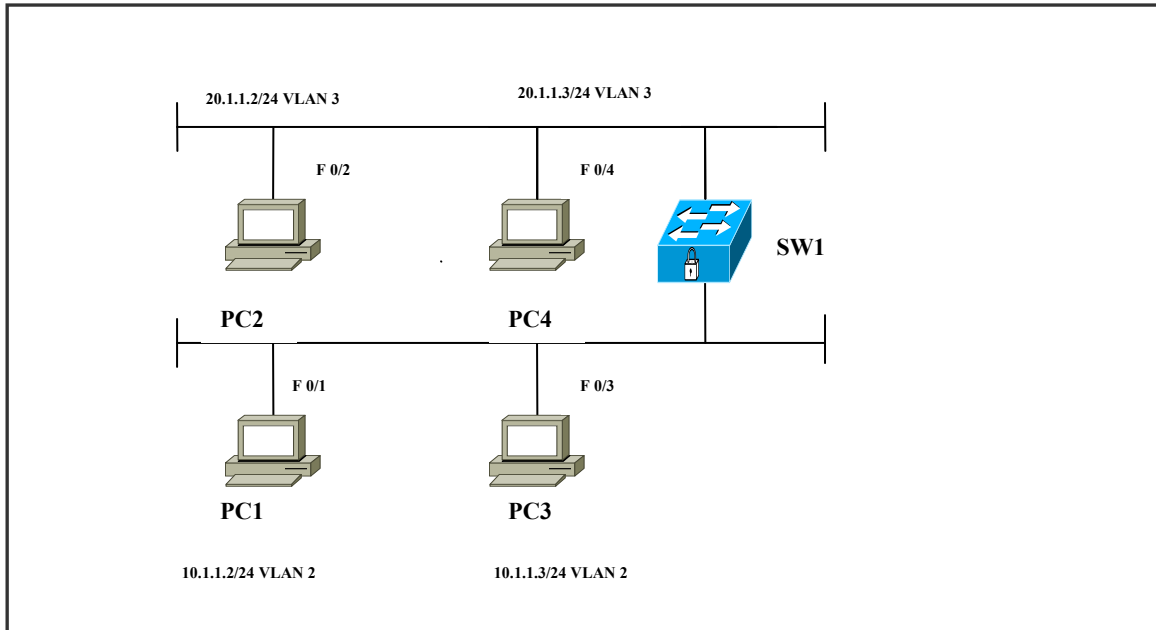
BULDING CISCO MULTILAYER SWITCHED NETWORK

BCMSN (642–812)

SWITCHING LAB INDEX

- 1. IMPLEMENTING VLAN's**
- 2. CONFIGURE TRUNKING**
- 3. DYNAMIC TRUNKING PROTOCOL**
- 4. IMPLEMENTING INTER-VLAN ROUTING**
- 5. PROPAGATING VLAN CONFIGURATION WITH VTP**
- 6. IMPLEMENTING SPANNING TREE PROTOCOL**
- 7. LOAD BALANCING IN STP**
- 8. IMPLEMENTING MSTP**
- 9. CONFIGURE LINK AGGREGATION USING ETHER-CHANNEL**
- 10. CONFIGURE SPAN**
- 11. CONFIGURE LAYER 3 REDUNDANCY WITH HSRP**
- 12. CONFIGURE LAYER 3 REDUNDANCY WITH VRRP**
- 13. CONFIGURE LAYER 3 REDUNDANCY WITH GLBP**

Lab 1 – Implementing VLANs



SW1

Ports	Assigned VLANs	PC
FA 0/1	VLAN 2	PC 1 (10.1.1.2)
FA 0/2	VLAN 3	PC 2 (20.1.1.2)
FA 0/3	VLAN 2	PC 3 (10.1.1.3)
FA 0/4	VLAN 3	PC 4 (20.1.1.3)

Task 1

Create VLAN 2 and VLAN 3 and assign name SALES and FINANCE to each VLAN. Configure ports fa 0/2 –fa 0/4 as access-ports and assign VLAN 2 to ports fa 0/1 and fa0/3. Assign VLAN 3 to ports fa 0/2 and fa 0/4. Configure VLANs using the database mode.

SW1

Vlan database

Vlan 2

vlan 2 name sales

Vlan 3

Vlan 3 name finance

Int fa0/1

Switchport mode access

Switchport access vlan2

Int fa0/2

Switchport mode access

Switchport access vlan3

Int fa0/3

Switchport mode access

Switchport access vlan2

Int fa0/4

Switchport mode access

Switchport access vlan3

Verification:

SW1#show interfaces fastEthernet 0/1 switchport

Name: Fa0/1

Switchport: Enabled

Administrative Mode: static access

Operational Mode: static access

Administrative Trunking Encapsulation: negotiate

Operational Trunking Encapsulation: native

Negotiation of Trunking: Off

Access Mode VLAN: 1 (default)

Trunking Native Mode VLAN: 1 (default)

The output displays mode as access.

SW1#show vlan brief

VLAN Name	Status	Ports
1 default	active	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
2 sales	active	Fa0/1, Fa0/3
3 finance	active	Fa0/2, Fa0/4
1002 fddi-default	act/unsup	
1003 token-ring-default	act/unsup	
1004 fddinet-default	act/unsup	
1005 trnet-default	act/unsup	

The output displays VLAN 2, name SALES assigned to ports fa 0/1 and fa 0/3. Also VLAN 3, name FINANCE assigned to ports fa 0/2 and fa 0/4.

Verifying connectivity between PC 1 and PC 3 (i.e PC's in the same vlan):

From PC 1:

```
C:\>ping 10.1.1.3
```

```
Pinging 10.1.1.3 with 32 bytes of data:
```

```
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255  
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255  
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255  
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255
```

```
Ping statistics for 10.1.1.3:
```

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss)  
Approximate round trip times in milli-seconds:  
Minimum = 1ms, Maximum = 1ms, Average = 1ms
```

Ping is successful.

Verifying connectivity between PC 1 and PC 2 (i.e PC's in different vlan):

From PC 1:

```
C:\>ping 20.1.1.2
Pinging 20.1.1.2 with 32 bytes of data:
Destination host unreachable.
Destination host unreachable.
Destination host unreachable.
Destination host unreachable.

Ping statistics for 20.1.1.2:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Ping not successful.

Therefore we need a router connected to the switch to do inter-vlan communication. Also the link should be configured as trunk.

Task 2

Create VLAN 2 and VLAN 3 and assign name SALES and FINANCE to each VLAN. Configure ports fa 0/2 –fa 0/4 as access-ports and assign VLAN 2 to ports fa 0/1 and fa0/3. Assign VLAN 3 to ports fa 0/2 and fa 0/4. Configure VLANs using the global configuration mode.

SW1

Vlan 2
name sales

Vlan 3
name finance

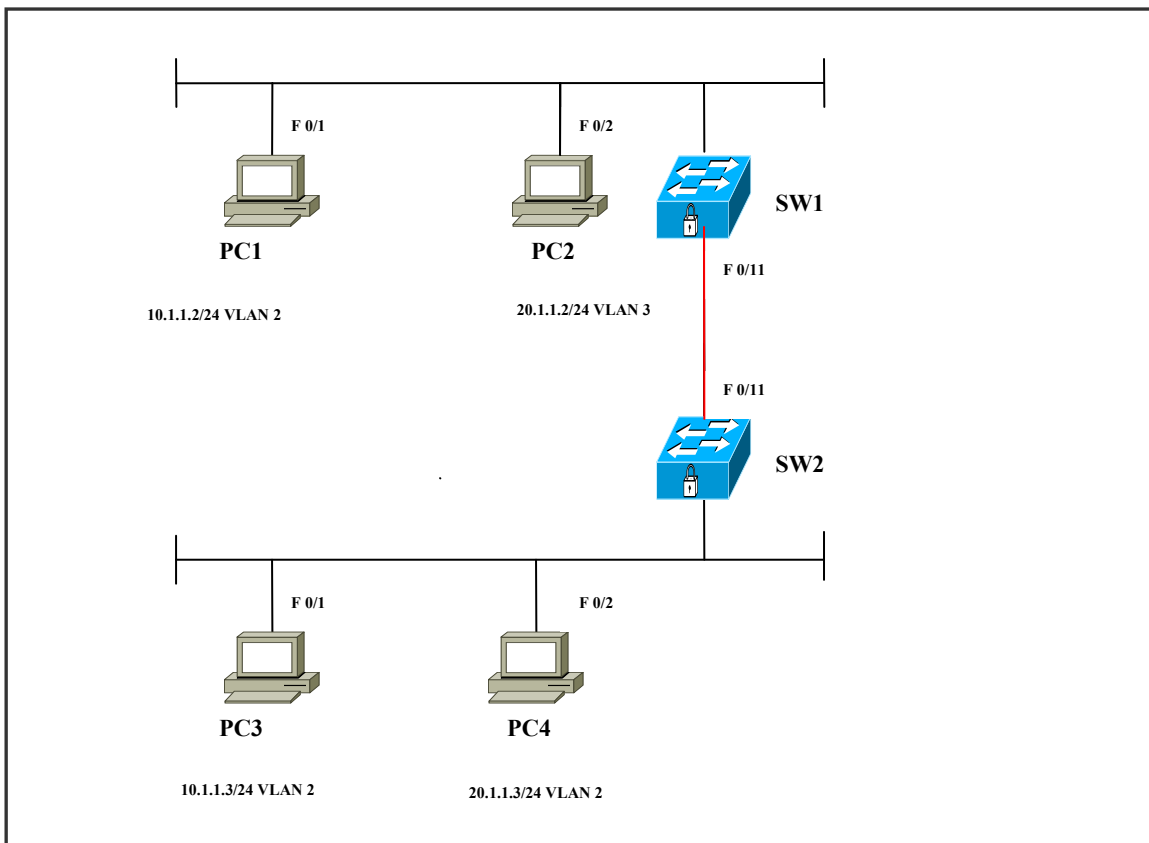
Int fa0/1
Switchport mode access
Switchport access vlan2

Int fa0/2
Switchport mode access
Switchport access vlan3

Int fa0/3
Switchport mode access
Switchport access vlan2

Int fa0/4
Switchport mode access
Switchport access vlan3

Lab 2 – Configure Trunking



SW1

Ports	Assigned VLANs	PC
FA 0/1	VLAN 2	PC 1 (10.1.1.2)
FA 0/2	VLAN 3	PC 2 (20.1.1.2)
FA 0/11 Configured as trunk		

SW2

Ports	Assigned VLANs	PC
FA 0/1	VLAN 2	PC 3 (10.1.1.3)
FA 0/2	VLAN 3	PC 4 (20.1.1.3)
FA 0/11 Configured as trunk		

Task 1

Create VLANs according to the scenario and assign to their respective access-ports.
Configure ISL trunk between SW1 (fa0/11) and SW 2 (fa0/11)

SW1

```
Vlan 2  
name sales
```

```
Vlan 3  
name finance
```

```
Int fa0/1  
Switchport mode access  
Switchport access vlan2
```

```
Int fa0/2  
Switchport mode access  
Switchport access vlan3
```

```
Int fa0/11  
shutdown  
Switchport trunk encapsulation isl  
Switchport mode trunk  
No shutdown
```

SW2

Vlan 2
name sales

Vlan 3
name finance

Int fa0/1
Switchport mode access
Switchport access vlan2

Int fa0/2
Switchport mode access
Switchport access vlan3

Int fa0/11
shutdown
Switchport trunk encapsulation isl
Switchport mode trunk
No shutdown

Verification :

SW1#show interfaces trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/11	on	isl	trunking	1

Port	Vlans allowed on trunk
Fa0/11	1-4094

Port	Vlans allowed and active in management domain
Fa0/11	1-3

Port	Vlans in spanning tree forwarding state and not pruned
Fa0/11	1-3

SW1#show interfaces fastEthernet 0/11 switchport

Name: Fa0/11
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk

Administrative Trunking Encapsulation: isl
Operational Trunking Encapsulation: isl
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)

Verifying connectivity between PC 1 and PC 3(i.e PC's in the same vlan)

From PC 1

```
C:\Documents and Settings\Administrator>ping 10.1.1.3
Pinging 10.1.1.3 with 32 bytes of data:

Reply from 10.1.1.3: bytes=32 time=2ms TTL=255
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255

Ping statistics for 10.1.1.3:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 2ms, Average = 1ms
```

Ping successful.

Verifying connectivity between PC 1 and PC 4 (i.e PC's in different vlan):

From PC 1:

```
C:\Documents and Settings\Administrator>ping 20.1.1.3
Pinging 20.1.1.3 with 32 bytes of data:

Destination host unreachable.
Destination host unreachable.
Destination host unreachable.
Destination host unreachable.

Ping statistics for 20.1.1.3:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Task 2

Create VLANs according to the scenario and assign to their respective access-ports.
Configure 802.1q (dot1q) trunk between SW1 (fa0/11) and SW 2 (fa0/11)

SW1

```
Int fa0/11
shutdown
Switchport trunk encapsulation dot1q
Switchport mode trunk
No shutdown
```

SW2

```
Int fa0/11
shutdown
Switchport trunk encapsulation dot1q
Switchport mode trunk
No shutdown
```

Verification :

SW1#show interfaces trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/11	on	802.1q	trunking	1

Port	Vlans allowed on trunk
Fa0/11	1-4094

Port	Vlans allowed and active in management domain
Fa0/11	1-3

Port	Vlans in spanning tree forwarding state and not pruned
Fa0/11	1-3

SW1#show interfaces fa0/11 switchport

```
Name: Fa0/11
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
```

The output displays mode as trunk and encapsulation used is dot1q, and the default native vlan is vlan1 which is used to carry the untagged frames across.

Verifying connectivity between PC 1 and PC 3 (i.e PC's in the same vlan)

From PC 1

```
C:\Documents and Settings\Administrator>ping 10.1.1.3
Pinging 10.1.1.3 with 32 bytes of data:

Reply from 10.1.1.3: bytes=32 time=2ms TTL=255
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255
Reply from 10.1.1.3: bytes=32 time=1ms TTL=255

Ping statistics for 10.1.1.3:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 2ms, Average = 1ms
```

Ping successful.

Verifying connectivity between PC 1 and PC 4 (i.e PC's in different vlan):

From PC 1:

```
C:\Documents and Settings\Administrator>ping 20.1.1.3
Pinging 20.1.1.3 with 32 bytes of data:

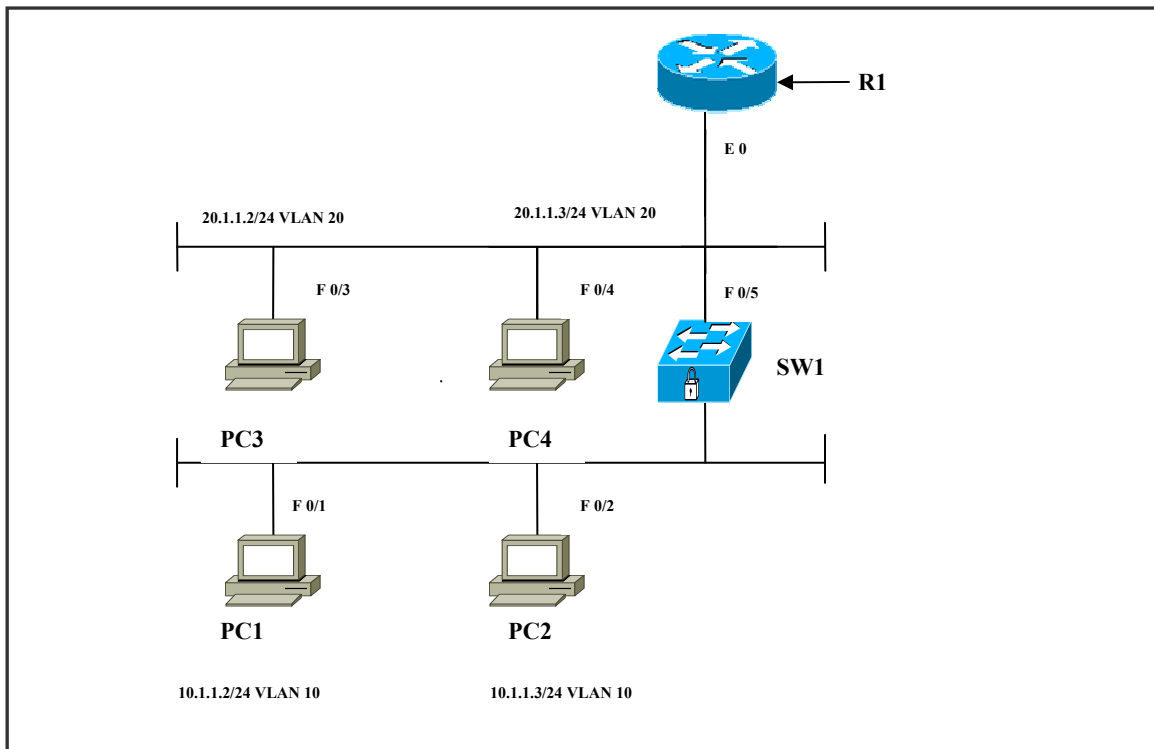
Destination host unreachable.
Destination host unreachable.
Destination host unreachable.
Destination host unreachable.

Ping statistics for 20.1.1.3:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

Ping not successful.

Therefore we need to configure inter-vlan routing.

Lab 3 – Implementing Inter-VLAN Routing



SW1

Ports	Assigned VLANs	PC
FA 0/1	VLAN 10	PC 1 (10.1.1.2)
FA 0/2	VLAN 10	PC 2 (10.1.1.3)
FA 0/3	VLAN 20	PC 3 (20.1.1.2)
FA 0/4	VLAN 20	PC 4 (20.1.1.3)
FA 0/5 Configured as trunk		

R1

Sub-Interfaces	Ip Address	Subnet - Mask
E 0/0.10	10.1.1.1	255.0.0.0
E 0/0.20	20.1.1.1	255.0.0.0
E 0 Configured as trunk		

Task 1

Create VLAN 10 and assign to ports fa 0/1 and fa 0/2. Create VLAN 20 and assign to ports fa 0/3 and fa 0/4. Configure port fa 0/5 as dot1q trunk. Use sub-interfaces on interface e 0 on R1 to accomplish this task.

R1

```
Int e 0  
No ip address
```

```
Int e 0/0.10  
Encapsulation dot1q 10  
Ip address 10.1.1.1 255.0.0.0
```

```
Int e 0/0.20  
Encapsulation dot1q 20  
Ip address 20.1.1.1 255.0.0.0
```

SW1

Vlan 10
Name sales
Vlan 20
Name finance

Int fa 0/1
Switchport mode access
Switchport access vlan 10

Int fa 0/2
Switchport mode access
Switchport access vlan 10

Int fa 0/3
Switchport mode access
Switchport access vlan 20

Int fa 0/4
Switchport mode access
Switchport access vlan 20

Int fa 0/5
Shutdown
Switchport trunk encapsulation dot1q
Switchport mode trunk
Switchport nonegotiate
No shutdown

Verification :

Verify if PC's in VLAN 10 can communicate with PC's in VLAN 20.

From PC 1 (10.1.1.2) :

```
C:\DOCUME~1\ADMINI~1>ping 20.1.1.2
Pinging 20.1.1.2 with 32 bytes of data:
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
Ping statistics for 20.1.1.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 2ms, Maximum = 2ms, Average = 2ms
```

Ping successful which means inter-vlan communication is working properly.

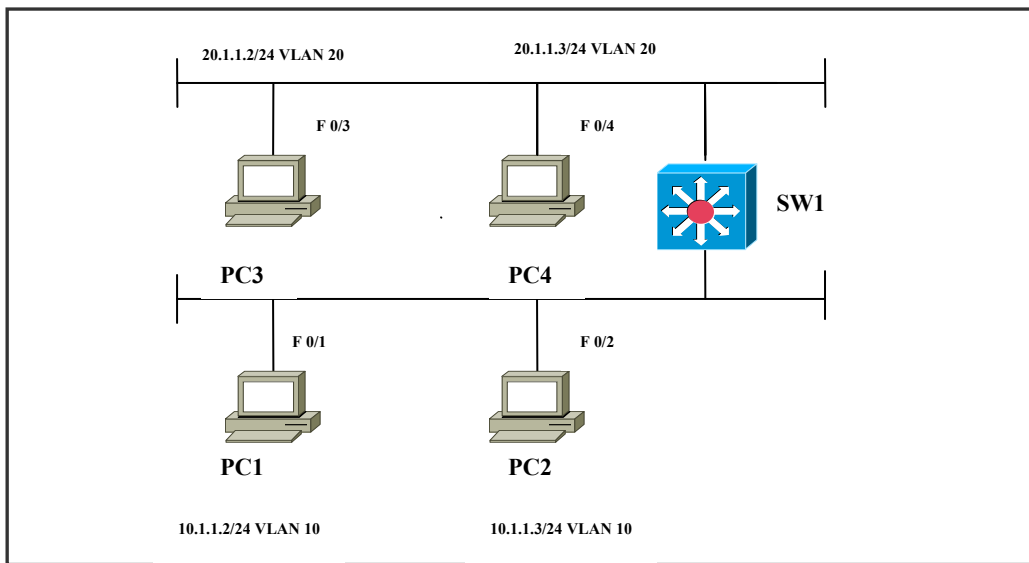
SW1#show int fa0/5 switchport

```
Name: Fa0/5
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
```

The output displays trunk dot1q encapsulation enabled.

Task 2

Implementing inter-vlan communication on a multilayer switch. Create VLAN 10 and assign to ports fa 0/1 and fa 0/2. Create VLAN 20 and assign to ports fa 0/3 and fa 0/4. Configure SVI and assign Ip address.



SW1

Ip routing

Interface vlan 10

Ip address 10.1.1.1 255.0.0.0

No shutdown

Interface vlan 20

Ip address 20.1.1.1 255.0.0.0

No shutdown

Verification :

From PC 1 (10.1.1.2) :

```
C:\>ping 20.1.1.2
```

```
Pinging 20.1.1.2 with 32 bytes of data:
```

```
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
```

```
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
```

```
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
```

```
Reply from 20.1.1.2: bytes=32 time=2ms TTL=254
```

```
Ping statistics for 20.1.1.2:
```


```
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
    approximate round trip times in milli-seconds:
```

```
        Minimum = 2ms, Maximum = 2ms, Average = 2ms
```

Ping 100% successful which means inter-vlan communication is working properly.

Lab 4 – Propagating VLAN Configuration with



Task 1

Configure Switch1 as the VTP Server and the other Switch (SW2) as VTP Client. Use NETMETRICS as the Domain name. Authenticate the relationship using CISCO123 as the password.

Switch1	Switch2
VTP domain NETMETRICS VTP mode server VTP password CISCO123	VTP domain NETMETRICS VTP mode client VTP password CISCO123

Task 2

Create VLANs 2,3,4, and 5 on SW1 (VTP SERVER) and name them as aaa, bbb, ccc, ddd.

Switch1
Vlan 2 Name aaa
Vlan 3 Name bbb
Vlan 4 Name ccc
Vlan 5 Name ddd

Verification:

SW1#show vtp status

```
VTP Version : 2
Configuration Revision : 15
Maximum VLANs supported locally : 1005
Number of existing VLANs : 9
VTP Operating Mode : Server
VTP Domain Name : netmetrics
```

The output displays vtp revision number, configuration revision number, vtp operation mode and vtp domain name.

SW2#show vtp status

```
VTP Version : 2
Configuration Revision : 15
Maximum VLANs supported locally : 1005
Number of existing VLANs : 9
VTP Operating Mode : Client
VTP Domain Name : netmetrics
```


The output displays vtp revision number, configuration revision number, vtp operation mode and vtp domain name.

SW2#show vlan brief

VLAN Name	Status	Ports
1 default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4 Fa0/5, Fa0/6, Fa0/7, Fa0/8 Fa0/9, Fa0/10, 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
2 aaa	active	
3 bbb	active	
4 ccc	active	
5 ddd	active	

The output displays VLANs 2, 3, 4, 5 propagated from vtp server.

Lab 5 – Implementing Spanning Tree Protocol



Task 1

Configure Switch1 as the VTP Server and the other Switch (SW2) as VTP Client. Configure SW1 to be the STP root for VLAN 1. Change the forward delay time such that the port transitions from listening to learning state in just 6 seconds instead of the default of 15 seconds. Configure ports fa 0/9 and fa 0/11 as dot1q trunks on both the switches.

Switch1	Switch2
VTP domain NETMETRICS VTP mode server VTP password CISCO123	VTP domain NETMETRICS VTP mode client VTP password CISCO123
Interface range fa0/9, fa0/11 Switchport trunk encapsulation Switchport mode trunk	Interface range fa0/9, fa0/11 Switchport trunk encapsulation Switchport mode trunk
Spanning-tree vlan 1 root primary Spanning-tree vlan 1 forward-time 6	

Verification:

```
SW1#show spanning-tree vlan 1
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID Priority 24577
```

```
Address 0014.a82f.a680
```

```
This bridge is the root
```

```
Hello Time 2 sec Max Age 20 sec Forward Delay 6 sec
```

Bridge ID Priority 24577 (priority 24576 sys-id-ext 1)
Address 0014.a82f.a680
Hello Time 2 sec Max Age 20 sec Forward Delay 6 sec
Aging Time 300

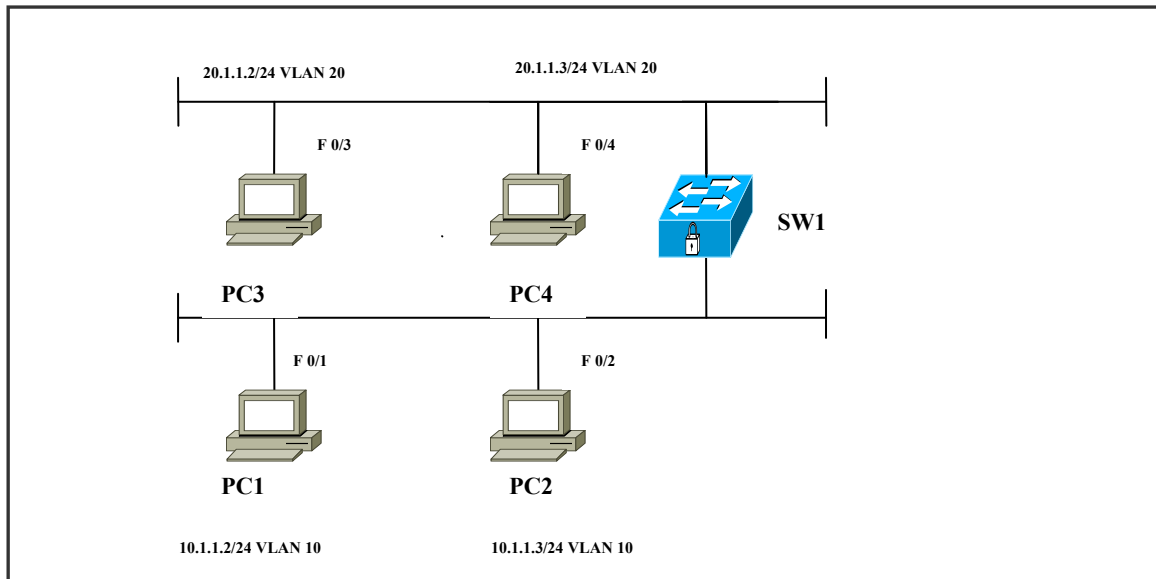
Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/1	Desg	FWD	19	128.1		P2p
Fa0/9	Desg	FWD	19	128.9		P2p
Fa0/11	Desg	FWD	19	128.11		P2p

The output displays that SW1 is the root bridge and forward delay time is 6 seconds.

SW1#debug spanning-tree events

```
05:23:22: STP: VLAN0004 Fa0/9 -> listening
05:23:22: set portid: VLAN0005 Fa0/9: new port id 8009
05:23:22: STP: VLAN0005 Fa0/9 -> listening
05:23:23: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/9,
changed state to up
05:23:28: STP: VLAN0001 Fa0/9 -> learning
05:23:34: STP: VLAN0001 Fa0/9 -> forwarding
05:23:37: STP: VLAN0002 Fa0/9 -> learning
05:23:37: STP: VLAN0005 Fa0/9 -> learning
05:23:52: STP: VLAN0002 Fa0/9 -> forwarding
```

The output displays the transition of ports from listening to learning in just 6 seconds instead of the default of 15 seconds.



Task 2

Configure ports fa0/1 - fa0/3 on SW1 to operate in portfast mode.

Switch1

```
Int range fa0/1 - fa0/3
Spanning-tree portfast
```

Verification:

SW1#show spanning-tree interface fa0/1 detail

```
Port 1 (FastEthernet0/1) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.1.
Designated root has priority 24577, address 0014.a82f.a680
Designated bridge has priority 24577, address 0014.a82f.a680
Designated port id is 128.1, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
The port is in the portfast mode
Link type is point-to-point by default
BPDU: sent 3140, received 0
```

The output displays that port fa0/1 is in portfast mode and also we see that BPDU's are sent

Verify the transition by shutting down interface fa 0/1 and again bringing the interface up.

SW1#debug spanning-tree events

```
05:36:16: set portid: VLAN0001 Fa0/1: new port id 8001
05:36:16: STP: VLAN0001 Fa0/1 -> jump to forwarding from blocking
```

The output displays port fa 0/1 jumps to forwarding state from blocking immediately because of portfast enabled on that port.

Task 3

(Scenario Based on Task 1)

Configure SW1 to quickly switch its root port in the event of an uplink failure. Trunking should be configured between the switches.

Switch1

```
spanning-tree uplinkfast
```

Verification:

SW1#show spanning-tree vlan 1

```
VLAN0001
Spanning tree enabled protocol ieee
Root ID    Priority    32769
  Address   000f.34f4.f080
  Cost      3019
  Port      11 (FastEthernet0/11)
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority    49153 (priority 49152 sys-id-ext 1)
  Address   0014.a82f.a680
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
  Aging Time 300
Uplinkfast enabled
```

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

```
Fa0/9        Altn BLK 3019 128.9  P2p
Fa0/11       Root FWD 3019 128.11 P2p
```

The output displays cost of ports increased by 3000 & priority of the bridge has increased to 49152.

Verify the transition from blocking to forwarding :

- Shutdown the port fa0/9 which is in the forwarding state.

SW1#debug spanning-tree uplinkfast


```
00:47:18: STP FAST: UPLINKFAST: make_forwarding on VLAN0001
FastEthernet0/11 root port id new: 128.11 prev: 128.9
```

```
00:47:18: %SPANTREE_FAST-7-PORT_FWD_UPLINK: VLAN0001 FastEthernet0/11
moved to Forwarding (UplinkFast).
```

```
00:47:18: STP: UFAST: removing prev root port Fa0/9 VLAN0001 port-id 8009
```

The output displays the transition of port fa0/11 from blocking to forwarding in one second.

Task 4



Configure portfast on port fa 0/9 between SW1 and SW2. (By default all ports on the switch are in dynamic desirable mode, they autonegotiate to become trunk. Portfast should not be enabled on trunk as there is a possibility of loops, we need to shutdown the port fa 0/9 and enable portfast on both the switches on port fa 0/9). Enable BPDU guard on port fa 0/9 of SW1 to stop BPDU's on that port.

Switch1	Switch2
Int fa 0/9 Shutdown Spanning-tree portfast Spanning-tree bpdu guard enable	Int fa 0/9 Shutdown Spanning-tree portfast

Now bring the port fa 0/9 on both switches to up. As soon as the BPDU's are being sent on the port . The port enabled with BPDU guard will immediately come into err-disable state .

Verification:

Console messages on SW1 when the bpdu's are received on the bpduguard enabled port fa0/9


```
01:04:30: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/9,
changed state to up
01:04:31: %SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port
FastEthernet0/9 with BPDU Guard enabled. Disabling port.
01:04:31: %PM-4-ERR_DISABLE: bpduguard error detected on Fa0/9, putting Fa0/9 in
err-disable state
01:04:32: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/9,
changed state to downstate.
```

SW1#show interfaces fa0/9

```
FastEthernet0/9 is down, line protocol is down (err-disabled)
Hardware is Fast Ethernet, address is 0014.a82f.a689 (bia 0014.a82f.a689)
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
```

The output displays the port as (err-disabled) state. But the BPDU's are sent out of this port, it doesn't affect that feature.

Task 5



Configure R1 to send BPDUs to SW 1. Enable BPDU guard on port fa0/3 on sw1 to block the access port fa 0/3 on SW 1, if any BPDUs received.

R1	Switch1
Int e0/0 No ip address Bridge-group 1 Bridge 1 protocol ieee Bridge 1 priority 4096	Int fa0/3 Spanning-tree bpduguard enable

Verification:

SW1#show interfaces fa0/3

FastEthernet0/3 is down, line protocol is down (err-disabled)

Hardware is Fast Ethernet, address is 0014.a82f.a683 (bia 0014.a82f.a683)
 MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
 reliability 255/255, txload 1/255, rxload 1/255
 Encapsulation ARPA, loopback not set
 Keepalive set (10 sec)

The output displays the port as (err-disabled) state. But the BPDU's are sent out of this port, it doesn't affect that feature.

SW1#debug spanning-tree events

01:33:12: STP: VLAN0001 Fa0/3 -> listening

01:33:13: %SPANTREE-2-BLOCK_BPDUGUARD: **Received BPDU on port FastEthernet0/3 with BPDU Guard enabled. Disabling port.**

01:33:13: %PM-4-ERR_DISABLE: **bpduguard error detected on Fa0/3, putting Fa0/3 in err-disable state**

The output displays that as soon as BPDU received on port fa0/3, it is disabled because of the BPDU guard enabled on that port.

Task 6

(Scenario Based on Task 5)

Configure portfast on port fa0/3 on sw1. Enable BPDU filter on port fa0/3 on sw1.
 Configure R1 to send BPDUs to port fa0/3 on sw1.

R1 Int e0/0 No shutdown No ip address Bridge-group 1 Bridge 1 protocol ieee Bridge 1 priority 4096	Switch1 Int fa0/3 Switchport mode access Spanning-tree portfast Spanning-tree bpduguard enable
---	---

Verification:

SW1#show spanning-tree interface fa0/3 detail

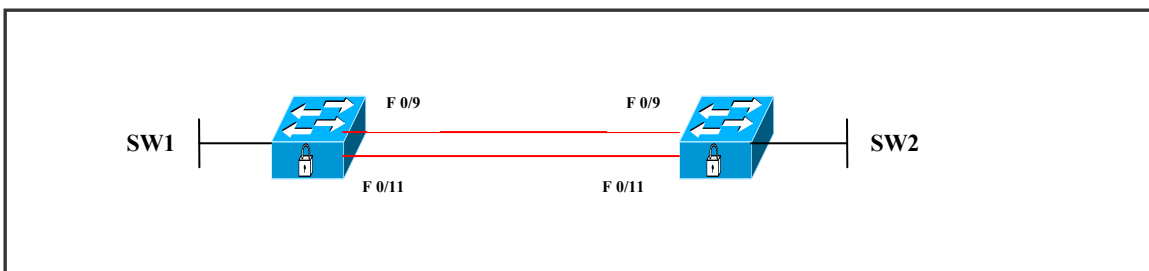
```

Port 3 (FastEthernet0/3) of VLAN0001 is forwarding
  Port path cost 100, Port priority 128, Port Identifier 128.3.
  Designated root has priority 32769, address 000f.34f4.f080
  Designated bridge has priority 32769, address 0014.a82f.a680
  Designated port id is 128.3, designated path cost 19
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is shared by default
  Bpdu filter is enabled
  BPDU: sent 0, received 0

```

The output displays BPDU filter enabled and no BPDU's sent or received

Task 7



Configure SW1 to be the root for VLAN 1. Configure root guard feature on SW1 port fa0/9, fa0/11.

Switch1 Spanning-tree vlan 1 root primary Int range fa0/9, fa0/11 No shutdown

Spanning-tree guard root

Verification:

SW1#show spanning-tree interface fa0/9 detail

Port 9 (FastEthernet0/9) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.9.
Designated root has priority 4097, address 0014.a82f.a680
Designated bridge has priority 4097, address 0014.a82f.a680
Designated port id is 128.9, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 2
Link type is point-to-point by default
Root guard is enabled on the port
BPDU: sent 3671, received 2431

The output displays root guard enabled on port.
Now change the priority in SW2:

Switch2

Spanning-tree vlan 1 priority 4096

As root guard is enabled on SW1, the ports on SW1 change to root inconsistent ports, thus blocking the port when superior BPDUs are received on SW1.

SW1#show spanning-tree interface fa0/9 detail

Port 9 (FastEthernet0/9) of VLAN0001 is broken (Root Inconsistent)
Port path cost 19, Port priority 128, Port Identifier 128.9.
Designated root has priority 32769, address 0014.a82f.a680
Designated bridge has priority 32769, address 0014.a82f.a680
Designated port id is 128.9, designated path cost 0
Timers: message age 1, forward delay 0, hold 0
Number of transitions to forwarding state: 2
Link type is point-to-point by default
Root guard is enabled on the port
BPDU: sent 3991, received 2445

SW1#debug spanning-tree events

```
03:29:15: %SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port
FastEthernet0/9 on VLAN0001.
03:29:15: STP: VLAN0001 Fa0/9 -> blocking
```

The output displays that root guard blocking port fa0/9

SW1#show spanning-tree vlan 1

```
VLAN0001
Spanning tree enabled protocol ieee
Root ID Priority 32769
  Address 0014.a82f.a680
  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 0014.a82f.a680
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
  Aging Time 300
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/9	Desg	BKN*	19	128.9	P2p *ROOT_Inc
Fa0/11	Desg	BKN*	19	128.11	P2p *ROOT_Inc

The output displays that the ports fa0/9 & fa0/11 are in “BKN” state as root-inconsistent type.

SW1#show spanning-tree inconsistentports

Name	Interface	Inconsistency
VLAN0001	FastEthernet0/9	Root Inconsistent
VLAN0001	FastEthernet0/11	Root Inconsistent

The output displays both fa 0/9 and fa 0/11 as inconsistent ports.

Task 8

(Scenario Based on Task 7)

Configure SW1 to the root bridge for vlan1. Configure loop guard on SW2, i.e., on the switch that is not the root bridge. Configure ports fa0/9 & fa0/11 between SW1 and SW2 as trunk ports.

Switch1	Switch2
Int range fa0/9 , fa0/11 Switchport trunk encapsulation dot1q Switchport mode trunk Spanning-tree vlan 1 root primary	Int range fa0/9 , fa0/11 Switchport trunk encapsulation dot1q Switchport mode trunk Spanning-tree guard loop

Verification:

SW2#show spanning-tree interface fastEthernet 0/9 detail

```

Port 9 (FastEthernet0/9) of VLAN0001 is forwarding
  Port path cost 19, Port priority 128, Port Identifier 128.9.
  Designated root has priority 24577, address 0014.a82f.a680
  Designated bridge has priority 24577, address 0014.a82f.a680
  Designated port id is 128.9, designated path cost 0
  Timers: message age 1, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  Loop guard is enabled on the port
  BPDU: sent 6419, received 2212

```

The output displays that loop guard is enabled on the port.

Now filter BPDU's on port fa 0/9 on SW1 :

Switch1
Int fa0/9 Spanning-tree bpdufilter enable

BPDU's will be stopped on SW1 and the port changes to loop inconsistent.

SW2#show spanning-tree interface fa0/9 detail

```

Port 9 (FastEthernet0/9) of VLAN0001 is broken (Loop Inconsistent)
  Port path cost 19, Port priority 128, Port Identifier 128.9.
  Designated root has priority 24577, address 0014.a82f.a680
  Designated bridge has priority 32769, address 000f.34f4.f080
  Designated port id is 128.9, designated path cost 19
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1

```

Link type is point-to-point by default
Loop guard is enabled on the port
BPDU: sent 6420, received 2257

The loop inconsistent state indicates that the port is not receiving any BPDU's or not sending any BPDU's through the port.

SW2#show spanning-tree inconsistentports

Name	Interface	Inconsistency
VLAN0001	FastEthernet0/9	Loop Inconsistent

Number of inconsistent ports (segments) in the system : 1

Lab 6 – Load Balancing in STP

Task 1

(Scenario Based on Lab 5 – Task 7)

Configure VTP to propagate VLAN information. Create 2 VLANs (VLAN 2, VLAN 3) on SW1 (VTP server). Configure SW1 to be the root for VLAN 2 and configure SW2 to be the root for VLAN 3.

Switch1	Switch2
spanning-tree vlan 2 root primary	spanning-tree vlan 3 root primary

Verification :

SW1#show spanning-tree vlan 2

VLAN0002

```
Spanning tree enabled protocol ieee
Root ID  Priority  24578
-----Output Omitted-----
```

```
Bridge ID Priority  24578 (priority 24576 sys-id-ext 2)
-----Output Omitted-----
```

```
Interface    Role Sts Cost    Prio.Nbr Type
-----
Fa0/9        Desg FWD 19    128.9  P2p
Fa0/11       Desg FWD 19    128.11 P2p
```

The output displays that SW1 is root for VLAN 2 i.e both ports fa 0/9 and fa 0/11 are in forwarding state.

SW1#show spanning-tree vlan 3

VLAN0003

```
Spanning tree enabled protocol ieee
Root ID  Priority  24578
-----Output Omitted-----
```

```
Bridge ID Priority  32771 (priority 32768 sys-id-ext 3)
```

-----Output Omitted-----

Interface	Role Sts Cost	Prio.Nbr	Type
Fa0/9	Root FWD 19	128.9	P2p
Fa0/11	Altn BLK 19	128.11	P2p

The output displays that SW1 is not root for VLAN 3 and port fa 0/9 is in forwarding and fa 0/11 is in blocked state as it is not the root bridge. Thus load balancing is achieved.

Task 2

(Scenario Based on Task 1)

Configure dot1q trunk between SW1 and SW2 on ports fa0/9 and fa0/11. Configure VTP on both the switches to propagate VLAN information. Create VLANs 1 to 6 on SW1 (server). Allow VLANs 2, 4, 6 on port fa0/9 on SW1 & SW2. Allow VLANs 1, 3, 5 on port fa0/11 on SW1 & SW2.

Switch1	Switch2
Int fa0/9 Switchport trunk encapsulation dot1q Switchport mode trunk Switchport trunk allowed vlan 2, 4, 6	Int fa0/9 Switchport trunk encapsulation dot1q Switchport mode trunk Switchport trunk allowed vlan 2, 4, 6
Int fa0/11 Switchport trunk encapsulation dot1q Switchport mode trunk Switchport trunk allowed vlan 1, 3, 5	Int fa0/11 Switchport trunk encapsulation dot1q Switchport mode trunk Switchport trunk allowed vlan 1, 3, 5

Verification:

SW1#show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/9	on	802.1q	trunking	1
Fa0/11	on	802.1q	trunking	1

Port Vlans allowed on trunk

Fa0/9	2,4,6
Fa0/11	1,3,5

Port Vlans allowed and active in management domain

```
Fa0/9 2,4,6
Fa0/11 1,3,5
```

Port Vlans in spanning tree forwarding state and not pruned

```
Fa0/9 2,4,6
Fa0/11 1,3,5
```

The output displays that on port fa0/9 only vlans 2, 4, 6 and on port fa0/11 only vlans 1,3,5 are allowed

SW1#show spanning-tree vlan 2

VLAN0002

```
Spanning tree enabled protocol ieee
Root ID Priority 32770
  Address 000f.34f4.f080
  Cost 19
  Port 9 (FastEthernet0/9)
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID Priority 32770 (priority 32768 sys-id-ext 2)
  Address 0014.a82f.a680
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
  Aging Time 300
```

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

```
-----
Fa0/9 Root FWD 19 128.9 P2p
```

The output displays that port fa0/9 is in forwarding state as vlans 2 is configured to allow on port fa0/9.

SW1#show spanning-tree vlan 3

VLAN0003

```
Spanning tree enabled protocol ieee
Root ID Priority 32771
  Address 000f.34f4.f080
  Cost 19
  Port 11 (FastEthernet0/11)
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID Priority 32771 (priority 32768 sys-id-ext 3)
  Address 0014.a82f.a680
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```


Aging Time 300

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

Fa0/11	Root	FWD	19	128.11	P2p	
--------	------	-----	----	--------	-----	--

The output displays that port fa0/11 is in forwarding state as vlan 3 is configured to allow on port fa0/11. Thus load balancing is achieved.

Lab 7 – Implementing MSTP

(Scenario Based on Lab 6 – Task 1)

Task 1

Configure dot1q trunks on ports fa0/9 and fa0/11. Configure VTP to propagate VLAN information. Configure instance 1 MSTP and map VLANs 1, 2, 3. Configure instance 2 MSTP and map VLANs 4, 5, 6. Make SW1 the STP root for instances 1, 2.

Switch1	Switch
Int range fa0/9, fa0/11 Switchport trunk encapsulation dot1q Switchport mode trunk	Int range fa0/9, fa0/11 Switchport trunk encapsulation dot1q Switchport mode trunk
Vtp domain netmet Vtp mode server Vtp password cisco123	Vtp domain netmet Vtp mode client Vtp password cisco123
Vlan 2 Name aaa Vlan 3 Name bbb Vlan 4 Name ccc Vlan5 Name ddd Vlan 6 Name 666	Spanning-tree mode mst Spanning-tree mst configuration Instance 1 vlan 1 – 3 Instance 2 vlan 4 – 6
Spanning-tree mode mst Spanning-tree mst configuration Instance 1 vlan 1 – 3 Instance 2 vlan 4 – 6	
Spanning-tree mst 1 – 2 root primary	

Verification:

SW1#show spanning-tree mst 1

```
##### MST01 vlans mapped: 1-3  
Bridge address 0014.a82f.a680 priority 24577 (24576 sysid 1)  
Root this switch for MST01
```

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/9	Desg	FWD	200000	128.9		P2p
Fa0/11	Desg	FWD	200000	128.11		P2p

The output displays the VLANs mapped to this MST instance 1.

SW1#show spanning-tree mst 2

```
##### MST02 vlans mapped: 4-6  
Bridge address 0014.a82f.a680 priority 24578 (24576 sysid 2)  
Root this switch for MST02
```

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/9	Desg	FWD	200000	128.9		P2p
Fa0/11	Desg	FWD	200000	128.11		P2p

The output displays the VLANs mapped to this MST instance 2.

Task 2
(Scenario Based On Task 1)

Configure MSTP on SW1 & SW2. Make SW1 the STP root for instance 1. Make SW2 the STP root for instance 2. Configure MST instance 1 and map VLANs 1 - 3. Configure MST instance 2 and map VLANs 4 - 6.

Switch1	Switch2
Spanning-tree mode mst Spanning-tree mst configuration Instance 1 vlan 1 – 3 Instance 2 vlan 4 – 6	Spanning-tree mode mst Spanning-tree mst configuration Instance 1 vlan 1 – 3 Instance 2 vlan 4 – 6
Spanning-tree mst 1 root primary	Spanning-tree mst 2 root primary

Verification:

SW1#show spanning-tree mst 1

MST01 vlans mapped: 1-3

Bridge address 0014.a82f.a680 priority 24577 (24576 sysid 1)
Root this switch for MST01

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/9	Desg	FWD	200000	128.9		P2p
Fa0/11	Desg	FWD	200000	128.11		P2p

SW1#show spanning-tree mst 2

MST02 vlans mapped: 4-6

Bridge address 0014.a82f.a680 priority 32770 (32768 sysid 2)
Root address 000f.34f4.f080 priority 24578 (24576 sysid 2)
port Fa0/9 cost 200000 rem hops 19

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/9	Root	FWD	200000	128.9		P2p
Fa0/11	Altn	BLK	200000	128.11		P2p

The output displays that sw1 acts as the root bridge for vlans 1-3 only.

This can be verified from the output that ports fa0/9 and fa0/11 are in forwarding state only for vlans 1-3 whereas one port forwarding and other blocking for vlans 4-6 on the same switch.

SW2#show spanning-tree mst 1

MST01 vlans mapped: 1-3

Bridge address 000f.34f4.f080 priority 32769 (32768 sysid 1)
Root address 0014.a82f.a680 priority 24577 (24576 sysid 1)
port Fa0/9 cost 200000 rem hops 19

Interface	Role	Sts	Cost	Prio.	Nbr	Type
Fa0/9	Root	FWD	200000	128.9		P2p
Fa0/11	Altn	BLK	200000	128.11		P2p

SW2#show spanning-tree mst 2

MST02 vlans mapped: 4-6

Bridge address 000f.34f4.f080 priority 24578 (24576 sysid 2)
Root this switch for MST02


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

Fa0/9	Desg	FWD	200000	128.9		P2p
Fa0/11	Desg	FWD	200000	128.11		P2p

The output displays that sw2 acts as the root bridge for vlans 4-6 only.

This can be verified from the output that ports fa0/9 and fa0/11 are in forwarding state only for vlans 4-6 whereas one port forwarding and other blocking for vlans 1-3 on the same switch.

Lab 8 – Configuring Link Aggregation with EtherChannel



Task 1

Configure L 2 trunk between SW1 & SW2 using default encapsulation on ports fa0/9, fa0/11.

Configure ether channel between SW1 and SW2 on interfaces fa0/9, fa0/11, without using negotiation protocols.

Configure interfaces fa0/9, fa0/11 on SW1 & SW2 in channel group 1 with a mode of “on”.

Switch1	Switch2
Interface port-channel 1	Interface port-channel 1
Int range fa0/9, fa0/11	Int range fa0/9, fa0/11
Channel-group 1 mode on	Channel-group 1 mode on

Verification :

SW2#show etherchannel summary

Flags: D - down P - in port-channel

I - stand-alone s - suspended

H - Hot-standby (LACP only)

R - Layer3 **S - Layer2**

u - unsuitable for bundling

U - in use f - failed to allocate aggregator

d - default port

Number of channel-groups in use: 1

Number of aggregators: 1

```

Group  Port-channel  Protocol  Ports
-----+-----+-----+-----
1      Po1(SU)          -        Fa0/9(P) Fa0/11(P)

```

The output displays port channel created for ports fa0/9, fa0/11 and is denoted as po 1 (su) where, s : layer 2, u : in use, P : in port channel

SW2#show interfaces port-channel 1 switchport

```

Name: Po1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: trunk
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: isl
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)

```

By default all interfaces are in dynamic desirable mode which automatically negotiates to become trunk if not specified.
The output displays default ISL trunking for this port-channel.

SW2#show int trunk

```

Port      Mode      Encapsulation  Status      Native vlan
Po1       desirable n-isl          trunking    1

Port      Vlans allowed on trunk
Po1       1-4094

Port      Vlans allowed and active in management domain
Po1       1

Port      Vlans in spanning tree forwarding state and not pruned
Po1       1

```

The output displays the port-channel interface as trunk instead of individual ports.

SW2#show spanning-tree vlan 1

```

VLAN0001
Spanning tree enabled protocol ieee
Root ID  Priority 32769

```

This bridge is the root
 Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)

```
Interface Role Sts Cost Prio.Nbr Type
-----
Po1 Desg FWD 12 128.65 P2p
```

The output displays the forwarding port as portchannel 1 instead of separate port because of ether channel configured on ports fa0/9, fa0/11, they appear as one bundle.

Task 2
(Scenario Based On Task 1)

Configure L 2 trunk between SW1 & SW2 using dot1q or isl encapsulation on ports fa 0/7 , fa 0/9, fa 0/11.
 Configure ether channel between SW1 and SW2 on interfaces fa 0/7, fa0/9, fa0/11. Both switches SW1 and SW2 should initiate negotiation via PAgP.

Switch1	Switch2
Interface port-channel 1	Interface port-channel 1
Int range fa0/7 , fa0/9 , fa0/11 Channel-group 1 mode desirable	Int range fa0/7 , fa0/9 , fa0/11 Channel-group 1 mode desirable

Verification :

SW1#show etherchannel summary

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

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

```
Group Port-channel Protocol Ports
-----+-----+-----+-----
1 Po1(SU) PAgP Fa0/7(P) Fa0/9(P) Fa0/11(P)
```


The output displays protocol as “PAgP” and Po1 (SU) (port-channel 1) created for ports fa0/7, fa0/9, fa0/11, where s : layer 2, U : in use.

SW1#show interfaces port-channel 1 switchport

Name: Po1

Switchport: Enabled

Administrative Mode: dynamic desirable

Operational Mode: trunk

Administrative Trunking Encapsulation: negotiate

Operational Trunking Encapsulation: isl

Negotiation of Trunking: On

Access Mode VLAN: 1 (default)

Trunking Native Mode VLAN: 1 (default)

The output displays that this interface port-channel 1 has automatically negotiated to become trunk.

SW1#show int trunk

Port	Mode	Encapsulation	Status	Native vlan
Po1	desirable	n-isl	trunking	1

Port	Vlans allowed on trunk
Po1	1-4094

Port	Vlans allowed and active in management domain
Po1	1

Port	Vlans in spanning tree forwarding state and not pruned
Po1	1

The output displays port-channel 1 as trunk instead of individual ports.

SW1#show spanning-tree vlan 1

VLAN0001

Spanning tree enabled protocol ieee

Root ID Priority 32769

Address 000f.34f4.f080

Cost 9

Port 65 (Port-channel1)

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

```

Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
Address 0014.a82f.a680
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 300

```

```

Interface      Role Sts Cost    Prio.Nbr Type
-----
Po1            Root FWD 9      128.65 P2p

```

The output displays the forwarding port as portchannel 1 instead of separate ports.

Task 3
(Scenario Based On Task 1)

Configure L 2 trunk between SW1 & SW2 using dot1q or isl encapsulation on ports fa 0/7 , fa 0/9, fa 0/11.
Configure ether channel between SW1 and SW2 on interfaces fa 0/7, fa0/9, fa0/11. Both switches SW1 and SW2 should initiate negotiation via LAcP.

Switch1	Switch2
Interface port-channel 1	Interface port-channel 1
Int range fa0/7 , fa0/9 , fa0/11	Int range fa0/7 , fa0/9 , fa0/11
Channel-group 1 mode active	Channel-group 1 mode active

Verification :

SW1#show etherchannel summary

```

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

```

```

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

```

```

Group Port-channel Protocol Ports
-----+-----+-----+-----

```

1 Po1(SU) LACP Fa0/7(P) Fa0/9(P) Fa0/11(P)

The output displays protocol as “LACP” and po1 (SU) (port-channel 1) created for ports fa0/7, fa0/9, fa0/11.

Task 4

(Scenario Based On Task 1)

Configure interface port-channel 1 to ports fa0/7, fa0/9, fa0/11.

Configure VLAN 100 and assign to ports fa0/7, fa0/9, fa0/11.

Configure ether-channel between SW1 & SW2 i.e. create channel-group 1 with the mode “on” (without using negotiating protocols).

Switch1	Switch2
Interface range fa0/7, fa0/9, fa0/11 Switchport mode access Switchport access vlan 100 Channel-group 1 mode on	Interface range fa0/7, fa0/9, fa0/11 Switchport mode access Switchport access vlan 100 Channel-group 1 mode on

Verification :

SW1#show vlan brief

```
VLAN Name           Status    Ports
-----
1  default             active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                               Fa0/5, Fa0/6, Fa0/8, Fa0/10
                               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
100 VLAN0100         active    Po1
```

The output displays portchannel 1 in VLAN 100 instead of individual ports.

SW1#show etherchannel summary

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

```
Group Port-channel  Protocol           Ports
-----+-----+-----+-----
```

1 Po1(SU) - Fa0/7(P) Fa0/9(P) Fa0/11(P)

The output displays no protocol and po1 (SU) created for ports fa0/7, fa0/9, fa0/11, where :

P = Port-channel, s = layer 2, U = in use.

Task 5

(Scenario Based On Task 1)

To configure Layer 3 ether-channel, create the port channel logical interface, assign ip address and then put the ethernet interfaces into the port-channel.

Switch1	Switch2
Interface port-channel 1 Ip add 100.0.0.1 255.0.0.0	Interface port-channel 1 Ip add 100.0.0.2 255.0.0.0
Interface range fa0/7, fa0/9, fa0/11 No switchport No ip address Channel-group 1 mode on	Interface range fa0/7, fa0/9, fa0/11 No switchport No ip address Channel-group 1 mode on

Verification :

Test the connectivity of port-channel

Ping from SW1 to SW2

100 % successful

SW1#ping 100.0.0.2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 100.0.0.2, timeout is 2 seconds:

!!!!

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

SW1#show Etherchannel summary

Flags: D - down P - in port-channel

I - stand-alone s - suspended

H - Hot-standby (LACP only)

R - Layer3 S - Layer2

Number of channel-groups in use: 1

Number of aggregators: 1

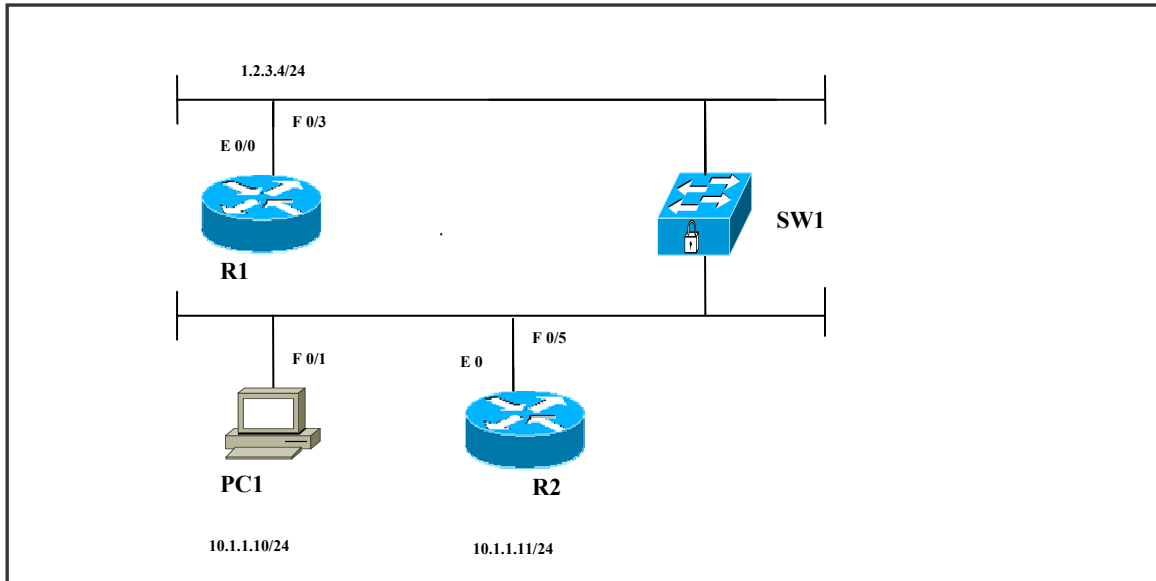
Group Port-channel Protocol Ports

-----+-----+-----+-----

1 Po1(RU) - Fa0/7(P) Fa0/9(P) Fa0/11(P)

The output displays port channel 1 created and denoted as po1 (RU) where : R = layer 3,
U = in use, p = port-channel

Lab 9 – SPAN: Switched Port Analyzer



Task 1

Create VLAN 10 and assign to ports fa0/1 & fa0/5 on SW1.
Configure SW1 to redirect all traffic from VLAN 10 to port fa0/3.
Enable R1 for debug process.

Switch1	Router1
Vlan 10	Int e0/0
Int range fa0/1, fa0/5	Ip add 1.2.3.4 255.0.0.0
Switchport mode access	No shutdown
Switchport access vlan 10	
Monitor session 1 source vlan 10 rx	
Monitor session 1 destination interface fa0/3	
Int fa0/3	
Switchport mode access	

Verification :

R1#debug ip packet

Now ping from PC1:

PC 1 > Ping 255.255.255.255.


IP packet debugging is on

***Mar 1 05:32:11.626: IP: s=10.1.1.11 (Ethernet0/0), d=255.255.255.255, len 100, rcvd 2**

***Mar 1 05:32:11.626: IP: s=1.2.3.4 (local), d=10.1.1.11, len 100, unroutable**

The output displays source ip 10.1.1.11 and destination ip as 255.255.255.255.
The second message displays, source ip 1.2.3.4 and destination ip as 10.1.1.11.
Thus, R1 receives packets sent from R2 even through they are not in the same VLAN.

Task 2



Configure dot1q encapsulation on port fa 0/9 of SW1 to become trunk.
Configure SPAN monitoring on port fa 0/9 of SW2 and also configure dot1q encapsulation of port fa 0/9 of SW2.

Switch1	Switch2
Int fa0/9 Switchport trunk encapsulation dot1q Switchport mode trunk	Int fa0/9 Switchport trunk encapsulation dot1q Switchport mode trunk Monitor session 1 source vlan1 rx Monitor session 1 destination interface fa0/9

Verification :

SW1#show interfaces trunk

Port	Mode	Encapsulation	Status	Native vlan
Fa0/9	on	802.1q	trunking	1

Port	Vlans allowed on trunk
Fa0/9	1-4094

Port	Vlans allowed and active in management domain
Fa0/9	1,10

Port	Vlans in spanning tree forwarding state and not pruned
Fa0/9	1,10

The output displays port fa0/9 as dot1q trunk

SW2#show interfaces trunk

-----Nil-----

The output doesn't display anything as there is no trunk established on port fa0/9 of SW2.

SW2#show int fa0/9

FastEthernet0/9 is up, line protocol is down (monitoring)

Hardware is Fast Ethernet, address is 000f.34f4.f089 (bia 000f.34f4.f089)

MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
reliability 255/255, txload 1/255, rxload 1/255

Encapsulation ARPA, loopback not set

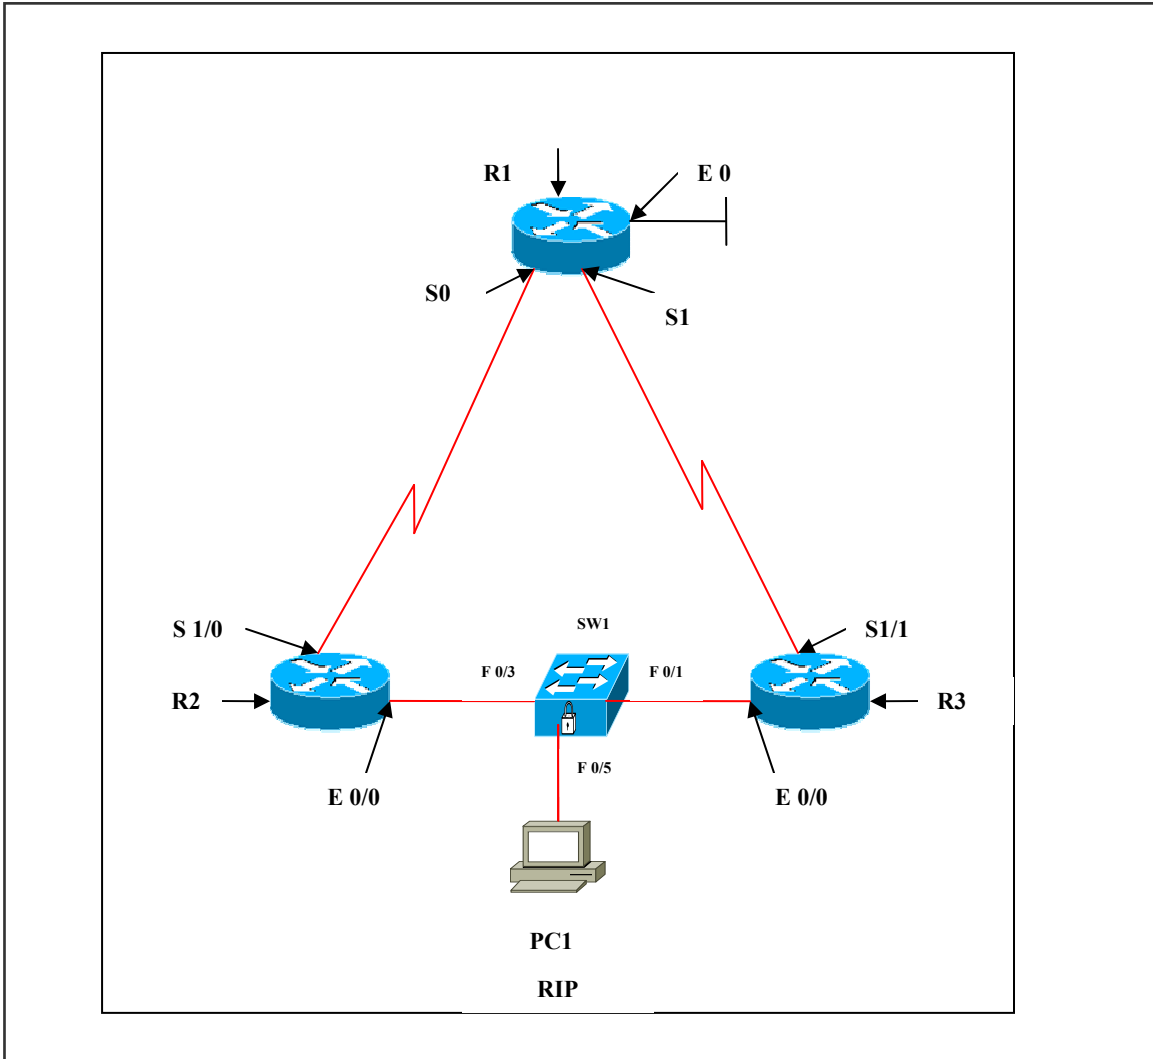
Keepalive set (10 sec)

The output displays line protocol down (monitoring).

To troubleshoot this issue we have to remove the SPAN monitoring for port fa 0/9 on SW2.

NOTE : SPAN monitoring should not be configured on trunk ports.

Lab 10 – Configuring HSRP



Interface IP Address Configuration :

R1

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
S 1	2.2.2.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 1/0	3.3.3.1	255.0.0.0
E 0/0	10.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1/1	2.2.2.1	255.0.0.0
E 0/0	10.1.1.2	255.0.0.0

SW1

Ports	VLAN Assigned	Connected To
FA 0/1	VLAN 1	R3 (10.1.1.2)
FA 0/3	VLAN 1	R2 (10.1.1.1)
FA 0/5	VLAN 1	PC 1 (10.1.1.3)

Task 1

Configure routing protocol (RIP) on R1, R2, and R3.

Configure ports fa0/3, fa0/1 & fa0/5 as access ports on SW1

Configure HSRP group 1 on R2 & R3, using the virtual ip address 10.1.1.10.

R2	R3
Int e0/0 Standby 1 ip 10.1.1.10 Standby 1 preempt Standby 1 priority 200	Int e0/0 Standby 1 ip 10.1.1.10 Standby 1 priority 100

Verification :

R3#show standby brief

P indicates configured to preempt.

```
Interface Grp Prio P State Active Standby Virtual IP
Et0/0 1 100 P Standby 10.1.1.1 local 10.1.1.10
```

The output displays that this is the standby router and the active router is (10.1.1.1), the virtual ip is (10.1.1.10) and this router configured to preempt.

R2#show ip arp

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	10.1.1.10	-	0000.0c07.ac01	ARPA	Ethernet0/0
Internet	10.1.1.2	23	0001.4289.a241	ARPA	Ethernet0/0
Internet	10.1.1.3	19	0008.0216.0d31	ARPA	Ethernet0/0
Internet	10.1.1.1	-	0008.a3d1.b540	ARPA	Ethernet0/0

The output displays ip add 10.1.1.10 (virtual ip) with the MAC address 0000.0c07.ac01, which is well-known HSRP MAC address 01 : hsrp group identifier, 0000.0c : vendor code, 07.ac : HSRP .

Verify the route chosen to reach network 20.1.1.1

From PC1 : tracert 20.1.1.1

From PC1:

```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

```
Tracing route to 20.1.1.1 over a maximum of 30 hops
```

```
 1      2 ms      3 ms      3 ms  10.1.1.1
 2     28 ms     28 ms     28 ms  20.1.1.1
```

Traceroute command displays that the packet reaches 10.1.1.1 (active router) and reaches 20.1.1.1.

Shut down the interface E0/0 on router 2 and traceroute .

From PC1 : tracert 20.1.1.1

From PC1:

```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

```
Tracing route to 20.1.1.1 over a maximum of 30 hops
```

```
 1      1 ms      1 ms      1 ms  10.1.1.2
 2     27 ms     27 ms     28 ms  20.1.1.1
```

Traceroute command displays that the packet reaches 10.1.1.2 (standby router becomes active) and then reaches 20.1.1.1.

Task 2

Configure tracking on R2, so that in case of failure of S 1/0, the priority is decreased automatically, so that the standby router takes the active role.

R2

```
Int e0/0
Standby 1 track s1/0 150
```

Verification :

R2#show standby brief

P indicates configured to preempt.

Interface	Grp	Prio	P	State	Active addr	Standby addr	Group addr
Et0/0	1	200	P	Active	local	10.1.1.2	10.1.1.10

The output displays that this is the standby router and the active router is (10.1.1.1), the virtual ip is (10.1.1.10) and this router configured to preempt.

From PC1:

```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

Tracing route to 20.1.1.1 over a maximum of 30 hops

1	2 ms	3 ms	3 ms	10.1.1.1
2	28 ms	28 ms	28 ms	20.1.1.1

Traceroute command displays that the packet reaches 10.1.1.1 (active router) and reaches 20.1.1.1.

Shut down the interface s1/0 on router 2 and traceroute .

R2#show standby brief

P indicates configured to preempt.

Interface	Grp	Prio	P	State	Active addr	Standby addr	Group addr
Et0/0	1	50	P	Standby	10.1.1.2	local	10.1.1.10

The priority is decreased to 50 from 200 automatically, and the router becomes standby for the HSRP group.

From PC1:

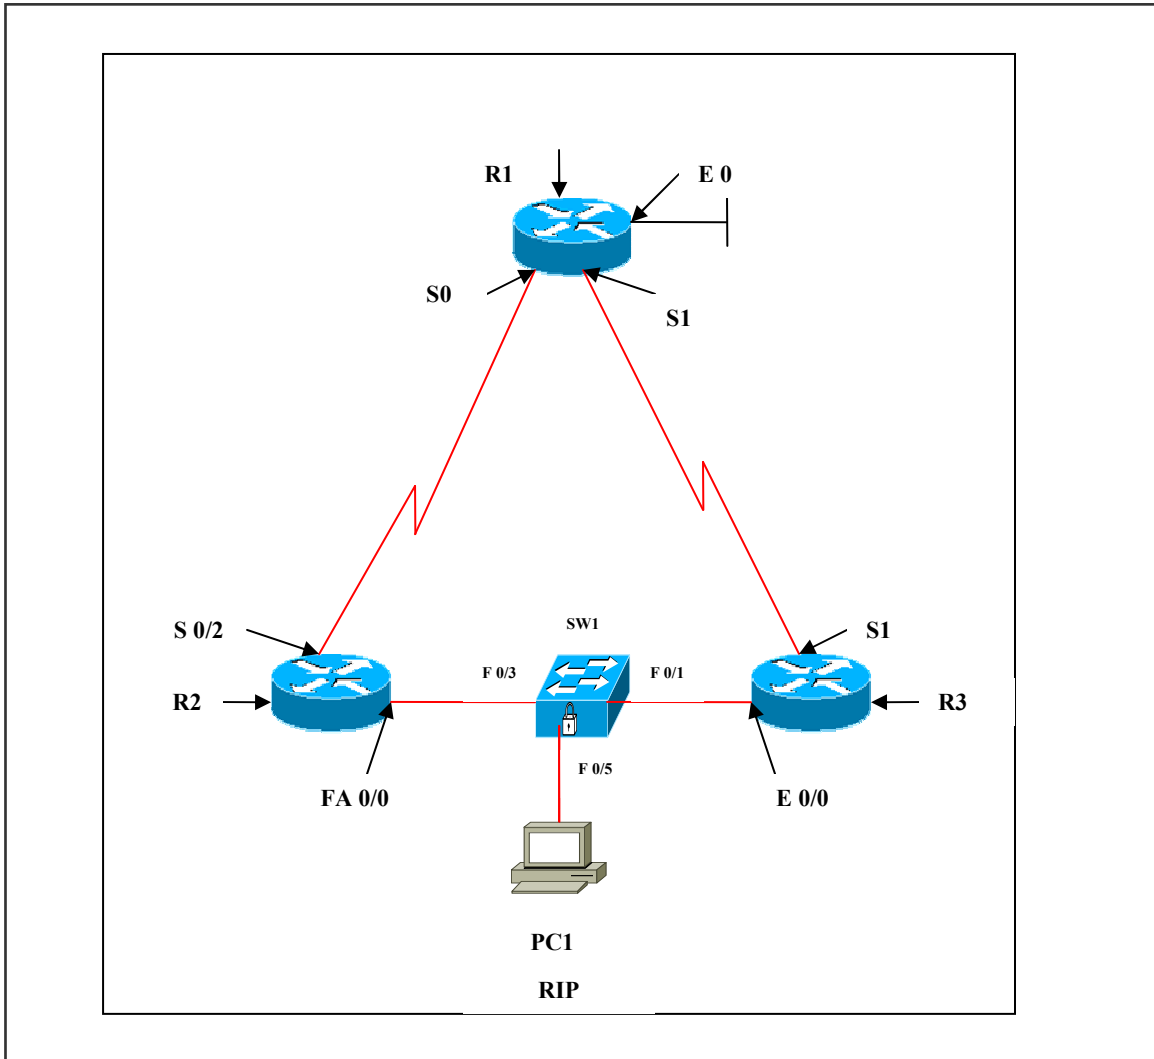
```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

Tracing route to 20.1.1.1 over a maximum of 30 hops

1	1 ms	1 ms	1 ms	10.1.1.2
2	27 ms	27 ms	28 ms	20.1.1.1

The traceroute command displays that the packets are sent via 10.1.1.2 (standby becomes active) because of higher priority value.

Lab 11 – Configuring VRRP



Interface IP Address Configuration :

R1

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
S 1	2.2.2.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0/2	3.3.3.1	255.0.0.0
E 0/0	10.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	2.2.2.1	255.0.0.0
E 0/0	10.1.1.2	255.0.0.0

SW1

Ports	VLAN Assigned	Connected To
FA 0/1	VLAN 1	R3 (10.1.1.2)
FA 0/3	VLAN 1	R2 (10.1.1.1)
FA 0/5	VLAN 1	PC 1 (10.1.1.3)

Task 1

Configure routing protocol (RIP) on R1, R2, and R3.

Configure ports fa0/3, fa0/1 & fa0/5 as access ports on SW1

Configure VRRP group 1 on R2 & R3, using the virtual ip address 10.1.1.10.

R2	R3
Int fa 0/0 Vrrp 1 ip 10.1.1.10 Vrrp 1 priority 200 Vrrp 1 timers advertise 4 Vrrp 1 preempt	Int e 0/0 Vrrp 1 ip 10.1.1.10 Vrrp 1 priority 100 Vrrp 1 timers learn Vrrp 1 preempt

Verification :

R2#show vrrp brief

```
Interface      Grp Pri Time Own Pre State  Master addr  Group addr
FastEthernet0/0  1  200 12218   Y  Master  10.1.1.1    10.1.1.10
```

The output displays that this router is master and virtual ip address is 10.1.1.10

R3#show vrrp brief

Interface	Grp	Pri	Time	Own	Pre	State	Master addr	Group addr
Ethernet0/0	1	100	12609		Y	Backup	10.1.1.1	10.1.1.10

The output displays that this router is backup router

From PC1:

```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

```
Tracing route to 20.1.1.1 over a maximum of 30 hops
```

```
 1      1 ms      1 ms      1 ms  10.1.1.1
 2     27 ms     27 ms     27 ms  20.1.1.1
```

Trace complete.

When packets sent to network 20.1.1.1 from PC1 (10.1.1.3), the packet first reaches 10.1.1.1 (master) and finally reaches the destination.

Shut down the interface Fa0/0 on router 2 and traceroute .

R3#debug vrrp packets

```
*Mar 1 05:01:40.790: %VRRP-6-STATECHANGE: Et0/0 Grp 1 state Backup -> Master
*Mar 1 05:01:40.790: VRRP: Grp 1 sending Advertisement checksum 6FF1
```

The output displays the transition of backup to master on R3.

R3#show vrrp brief

Interface	Grp	Pri	Time	Own	Pre	State	Master addr	Group addr
Ethernet0/0	1	100	3609		Y	Master	10.1.1.2	10.1.1.10

The output displays that R3 is master now.

From PC1:

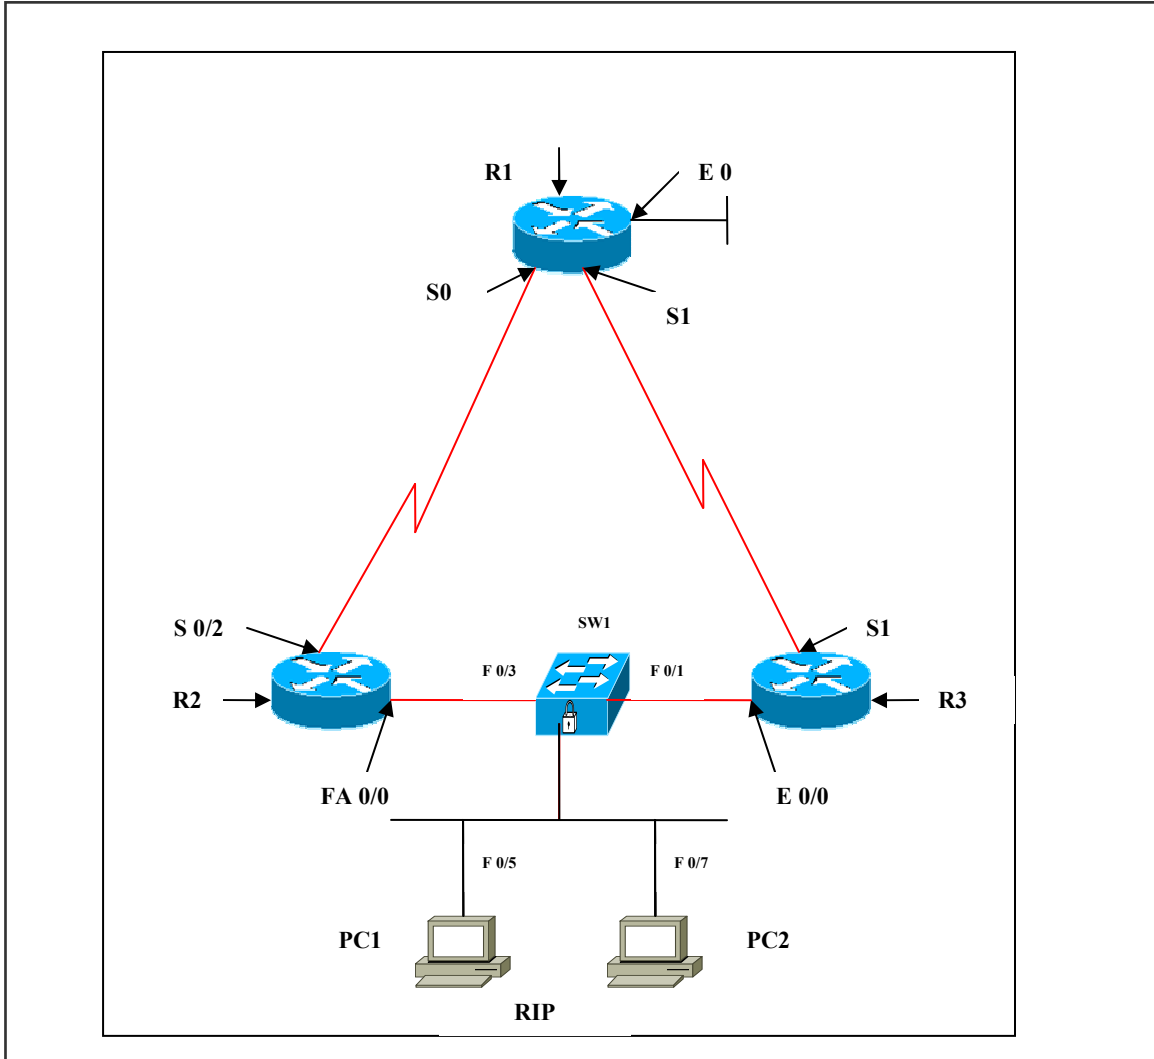
```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

```
Tracing route to 20.1.1.1 over a maximum of 30 hops
```

```
 1      1 ms      1 ms      1 ms  10.1.1.2
 2     28 ms     28 ms     28 ms  20.1.1.1
```

When traceroute from PC1 to 20.1.1.1, the output displays that packet is reaching 20.1.1.1 via 10.1.1.2.

Lab 12 – Configuring GLBP



Interface IP Address Configuration :

R1

Interface	IP Address	Subnet Mask
S 0	3.3.3.2	255.0.0.0
S 1	2.2.2.2	255.0.0.0
E 0	20.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0/2	3.3.3.1	255.0.0.0
FA 0/0	10.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 1	2.2.2.1	255.0.0.0
E 0/0	10.1.1.2	255.0.0.0

SW1

Ports	VLAN Assigned	Connected To
FA 0/1	VLAN 1	R3 (10.1.1.2)
FA 0/3	VLAN 1	R2 (10.1.1.1)
FA 0/5	VLAN 1	PC 1 (10.1.1.3)
FA 0/7	VLAN 1	PC 2 (10.1.1.4)

Task 1

Configure routing protocol (RIP) on R1, R2, and R3.
Configure ports fa0/3, fa0/1, fa0/5 and fa0/7 as access ports on SW1
Configure GLBP.

R2	R3
Int fa 0/0 Glbp 1 ip 10.1.1.10 Glbp 1 priority 200 Glbp 1 timers msec 250 msec 750 Glbp 1 preempt	Int e 0/0 Glbp 1 ip 10.1.1.10 Glbp 1 priority 100 Glbp 1 timers msec 250 msec 750

Verification:

R2#show glbp brief

Interface	Grp	Fwd	Pri	State	Address	Active router	Standby route
Fa0/0	1	-	200	Active	10.1.1.10	local	10.1.1.2
Fa0/0	1	1	7	Active	0007.b400.0101	local	-
Fa0/0	1	2	7	Listen	0007.b400.0102	10.1.1.2	-

When the PC's send traffic to 20.0.0.0 network, the traffic is send via the active router which is R2. If R2 is busy in sending the traffic then R3 takes the active state and R2 is in the listening state. Thus load balancing is achieved.

R2#show glbp

FastEthernet0/0 - Group 1

State is Active

2 state changes, last state change 00:11:30

Virtual IP address is 10.1.1.10

Hello time 250 msec, hold time 750 msec

Next hello sent in 0.000 secs

Redirect time 600 sec, forwarder time-out 14400 sec

Preemption enabled, min delay 0 sec

Active is local

Standby is 10.1.1.2, priority 100 (expires in 0.530 sec)

Priority 200 (configured)

Weighting 100 (default 100), thresholds: lower 1, upper 100

Load balancing: round-robin

Group members:

0001.4289.a241 (10.1.1.2)

0006.534b.7090 (10.1.1.1) local

There are 2 forwarders (1 active)

Forwarder 1

State is Active

1 state change, last state change 00:11:20

MAC address is 0007.b400.0101 (default)

Owner ID is 0006.534b.7090

Redirection enabled

Preemption enabled, min delay 30 sec

Active is local, weighting 100

Forwarder 2

State is Listen

MAC address is 0007.b400.0102 (learnt)

Owner ID is 0001.4289.a241

Redirection enabled, 599.800 sec remaining (maximum 600 sec)

Time to live: 14399.800 sec (maximum 14400 sec)

Preemption enabled, min delay 30 sec

Active is 10.1.1.2 (primary), weighting 100 (expires in 0.546 sec)

The output displays that the active router takes its default MAC address, whereas the second forwarder learns the MAC address from the default gateway (active forwarding router) (i.e. R2). Load balancing is achieved in round-robin algorithm.

How to verify:

Traceroute from PC1 to 20.0.0.0 network, the packet is send via R2.

Traceroute from PC2 to 20.0.0.0 network, the packet is send via R2.

If R2 is busy then the packet is send via R3, which is verified from the above output (traceroute 20.1.1.1 from PC2)

From PC1:

```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

```
Tracing route to 20.1.1.1 over a maximum of 30 hops
```

```
 1  <10 ms  <10 ms  <10 ms  10.1.1.1
 2  27 ms   27 ms   27 ms   20.1.1.1
```

The first packet reaches 20.1.1.1 via 10.1.1.1

From PC2:

```
C:\Documents and Settings\Administrator>tracert 20.1.1.1
```

```
Tracing route to 20.1.1.1 over a maximum of 30 hops
```

```
 1  1 ms    1 ms    1 ms    10.1.1.2
 2  28 ms   27 ms   27 ms   20.1.1.1
```

The first packet reaches 20.1.1.1 via 10.1.1.2

PAPER 3

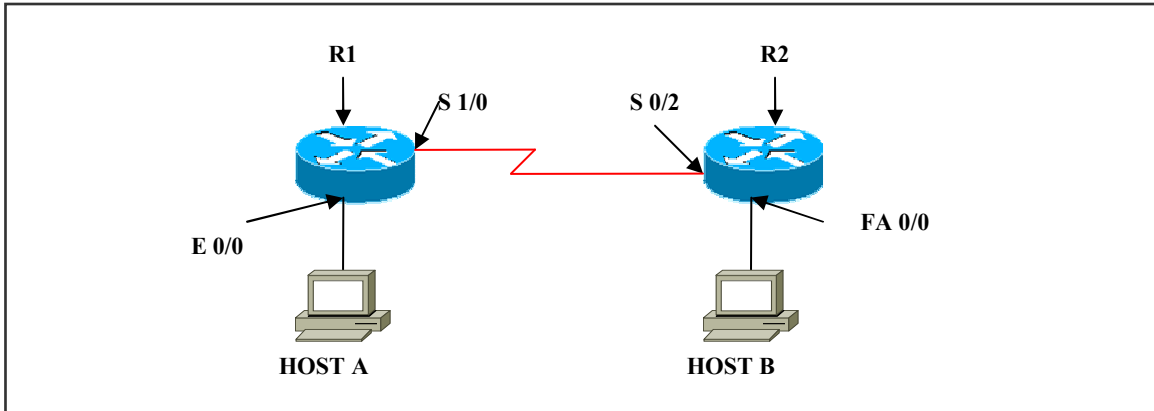
IMPLEMENTING SECURE CONVERGED WIDE AREA NETWORKS

ISCW (642–825)

ISCW LAB INDEX

1. CONFIGURE SITE-TO-SITE IPSEC VPN
2. CONFIGURE SITE-TO-SITE IPSEC VPN USING SDM.
3. CONFIGURE SPLIT TUNNELLING
4. CONFIGURE GRE TUNNEL (POINT-TO-POINT)
5. CONFIGURE GRE TUNNELLING USING THREE ROUTERS WITH NO ROUTING IN THE MIDDLE ROUTER
6. CONFIGURE GRE OVER IPSEC
7. CONFIGURE GRE OVER IPSEC SITE-TO-SITE TUNNEL USING SDM
8. CONFIGURE CISCO VPN CLIENT (PC) / REMOTE ACCESS VPN
9. CONFIGURE CISCO EASY VPN SERVER AND CLIENT (PC)
10. CONFIGURE CISCO EASY VPN SERVER AND CLIENT (ROUTER)
11. CONFIGURE FRAME MODE MPLS
12. CONFIGURING SSH SERVER FOR SECURE MANAGEMENT AND REPORTING
13. CONFIGURING SYSLOG LOGGING
14. CONFIGURATION OF SNMP
15. CONFIGURATION OF NTPv3
16. CONFIGURING AAA ON CISCO ROUTERS
17. DISABLING UNUSED CISCO ROUTERS USING NETWORK SERVICES AND INTERFACES
18. SECURITY CISCO ROUTER INSTALLATION AND ADMINISTRATIVE ACCESS

Lab 1 – Configure Site-to-Site IPSEC VPN



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.1	255.0.0.0
E 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0/2	2.2.2.2	255.0.0.0
Fa 0/0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure the ISAKMP policy required to establish on IKE tunnel.

Define the IPsec transform-set.

Create crypto ACL to define which traffic should be sent through the IPsec tunnel.

Create crypto map that maps the previously configured parameters and defines IPsec peer device.

Apply the crypto map to the outgoing interface of the VPN device.

R1	R2
Crypto isakmp enable	Crypto isakmp enable
Crypto isakmp policy 20	Crypto isakmp policy 15
Encryption 3des	Encryption 3des
Hash md5	Hash md5
Authentication pre-share	Authentication pre-share
Group1	Group1
Crypto isakmp key cisco123 address 2.2.2.2	Crypto isakmp key cisco123 address 2.2.2.1
Crypto ipsec transform-set set1 esp-des	Crypto ipsec transform-set set1 esp-des
Access-list 101 permit ip	Access-list 101 permit ip
10.0.0.0 0.255.255.255	20.0.0.0 0.255.255.255
20.0.0.0 0.255.255.255	10.0.0.0 0.255.255.255
Crypto map map1 10 ipsec-isakmp	Crypto map map1 10 ipsec-isakmp
Set peer 2.2.2.2	Set peer 2.2.2.1
Set transform-set set1	Set transform-set set1
Match address 101	Match address 101
Int s1/0	Int s0/2
Crypto map map1	Crypto map map1
Ip route 20.0.0.0 255.0.0.0 2.2.2.2	Ip route 10.0.0.0 255.0.0.0 2.2.2.2

Verification:

R1#show crypto isakmp sa

```
dst          src          state        conn-id  slot
2.2.2.2     2.2.2.1     QM_IDLE     1        0
```

The output displays the IKE tunnel established between src and dst. With the state displayed as QM-IDLE and a connection-id, if nothing of the above is displayed then the IKE phase I has not established.

R1#show crypto isakmp policy

```
Protection suite of priority 10
  encryption algorithm: Three key triple DES
  hash algorithm:      Message Digest 5
  authentication method: Pre-Shared Key
  Diffie-Hellman group: #1 (768 bit)
  lifetime:           86400 seconds, no volume limit
```

The output displays all the policies defined and also the default policy set.

R1#show crypto isakmp key

```
Keyring      Hostname/Address      Preshared Key
default      2.2.2.2              cisco123
```

The output displays the pre-shared key defined manually.

R1#show crypto map

```
Crypto Map "map1" 10 ipsec-isakmp
  Peer = 2.2.2.2
  Extended IP access list 101
  access-list 101 permit ip 10.0.0.0 0.255.255.255 20.0.0.0 0.255.255.255
  Current peer: 2.2.2.2
  Security association lifetime: 4608000 kilobytes/3600 seconds
  PFS (Y/N): N
  Transform sets={
    set1,
  }
  Interfaces using crypto map map1:
    Serial1/0
```

The output displays the crypto map configured and also SA lifetime is displayed.

R1#show crypto ipsec sa

```
interface: Serial1/0
  Crypto map tag: map1, local addr. 2.2.2.1

protected vrf:
local ident (addr/mask/prot/port): (10.0.0.0/255.0.0.0/0/0)
remote ident (addr/mask/prot/port): (20.0.0.0/255.0.0.0/0/0)
current_peer: 2.2.2.2:500
  PERMIT, flags={origin_is_acl,}
  #pkts encaps: 103, #pkts encrypt: 103, #pkts digest 0
```



```
#pkts decaps: 103, #pkts decrypt: 103, #pkts verify 0
#pkts compressed: 0, #pkts decompressed: 0
#pkts not compressed: 0, #pkts compr. failed: 0
#pkts not decompressed: 0, #pkts decompress failed: 0
#send errors 1, #recv errors 0
local crypto endpt.: 2.2.2.1, remote crypto endpt.: 2.2.2.2
path mtu 1500, media mtu 1500
current outbound spi: 163B5574
inbound esp sas:
  spi: 0xECF19512(3975255314)
  transform: esp-des,
  in use settings = {Tunnel, }
  slot: 0, conn id: 2000, flow_id: 1, crypto map: map1
  sa timing: remaining key lifetime (k/sec): (4412874/2809)
  IV size: 8 bytes
  replay detection support: N
outbound esp sas:
  spi: 0x163B5574(372987252)
  transform: esp-des,
  in use settings = {Tunnel, }
  slot: 0, conn id: 2001, flow_id: 2, crypto map: map1
  sa timing: remaining key lifetime (k/sec): (4412874/2809)
  IV size: 8 bytes
  replay detection support: N
```

The output displays the packets encrypted or decrypted.
Before verifying this command ping to the destination, i.e., once the interesting traffic is sent the SA is formed and then secured.

R1#show crypto ipsec transform-set

```
Transform set set1: { esp-des }
will negotiate = { Tunnel, },
```

The output displays the transform-set.

Lab 2 – Configure IPSEC Site-to-Site VPN Using SDM

Configure IPsec side-to-side VPN using SDM (Security Device Manager).

- SDM is an easy-to-use internet browser-based device management tool that is embedded within Cisco IOS 800 – 3800 series router at no cost.
- SDM simplifies router and security configuration through the use of intelligent wizards to enable customers and partners to quickly and easily deploy, configure, and monitor Cisco router.

Navigations

From the desktop, start the cisco SDM launcher software.
Click configure icon from the main window.
Click VPN icon to open VPN page.
Choose side-to-side VPN wizard from the list.
Click launch the selected task button.
Window will open to choose wizard mode.
Choose step-by-step setup.
Choose the outside interface towards IPsec peer.
Specify the IP address of the peer.
Choose the authentication method and specify the key.
Click next button to proceed.

Set IKE policies by clicking add button and specify the parameters:

IKE proposal priority – 2
Encryption algorithm – 3des.
HMAC – sha
IKE authentication method
Diffie-Hellman group – 1
IKE lifetime

Click next button to proceed..

Set transform-set by clicking add button and specify the parameters:

Transform set name – set 1
Encryption algorithm – esp-des
HMAC
Mode of operation – tunnel
Optional compression


Click next to proceed

- Click create / select an access-list for IPSec traffic radio button.
- Click create a new rule (ACL) and select option.
- Give the access rule a name and click add button.
 - At the end of step-by-step setup the wizard presents a summary of the configured parameters.
 - Click finish button to complete the configuration.

Verify

- Click “test tunnel” button to run a test to determine the configuration correctness of the tunnel.
- Click “monitor icon” – the screen will display all IPSec tunnels, parameters and status.
- Same with “VPN status” icon & “IPSec tunnels”.

Lab 3 – Configure Split Tunneling



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.1	255.0.0.0
E 1/0	10.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0/2	3.3.3.2	255.0.0.0

Fa 0/0	20.1.1.1	255.0.0.0
--------	----------	-----------

Lab Objective:

Task 1

Configure routing (EIGRP 10) on R1, R2, and R3.

Configure IPsec VPN only on R1 and R2.

No IPsec VPN configuration on R3.

R1	R2
Crypto isakmp enable	Crypto isakmp enable
Crypto isakmp policy 10 Encryption 3des Hash md5 Authentication pre-share Group1	Crypto isakmp policy 15 Encryption 3des Hash md5 Authentication pre-share Group1
Crypto isakmp key cisco123 address 3.3.3.2	Crypto isakmp key cisco123 address 2.2.2.1
Crypto ipsec transform-set set1 esp-des	Crypto ipsec transform-set set1 esp-des
Access-list 101 permit ip 10.0.0.0 0.255.255.255 20.0.0.0 0.255.255.255	Access-list 101 permit ip 20.0.0.0 0.255.255.255 10.0.0.0 0.255.255.255
Crypto map map1 10 ipsec-isakmp Set peer 3.3.3.2 Set transform-set set1 Match address 101	Crypto map map1 10 ipsec-isakmp Set peer 2.2.2.1 Set transform-set set1 Match address 101
Int s1/0 Crypto map map1	Int s0/2 Crypto map map1

Verification:

R1#show crypto isakmp sa

```
dst      src      state      conn-id  slot
3.3.3.2  2.2.2.1  QM_IDLE   1        0
```

The output displays current IKE SA's. QM_IDLE status indicates an active IKE SA.

R1#show crypto ipsec sa

interface: Serial1/0

Crypto map tag: map1, local addr. 2.2.2.1

protected vrf:

local ident (addr/mask/prot/port): (10.0.0.0/255.0.0.0/0/0)

remote ident (addr/mask/prot/port): (20.0.0.0/255.0.0.0/0/0)

current_peer: 3.3.3.2:500

PERMIT, flags={origin_is_acl,}

#pkts encaps: 99, #pkts encrypt: 99, #pkts digest 0

#pkts decaps: 99, #pkts decrypt: 99, #pkts verify 0

#pkts compressed: 0, #pkts decompressed: 0

#pkts not compressed: 0, #pkts compr. failed: 0


#pkts not decompressed: 0, #pkts decompress failed: 0

#send errors 1, #recv errors 0

----output omitted----

The output displays current settings used by current SA's. Non-zero encryption and decryption statistics can indicate a working set of IPsec SA's.

Lab 4 – Configure GRE Tunnel (Point-to-point)



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.1	255.0.0.0
E 1/0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
Fa 0/0	20.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure static route for reachability to the destination ip address for both R1 and R2.
 Configure tunneling by creating interface tunnel 0 on both R1 & R2.
 Assign virtual IP address to this interface tunnel 0 on both R1 & R2.

R1	R2
<pre>Ip route 20.0.0.0 255.0.0.0 2.2.2.2 Int tunnel 0 Ip address 30.1.1.1 255.0.0.0 Tunnel source s1/0 Tunnel destination 2.2.2.2 Tunnel mode gre ip</pre>	<pre>Ip route 10.0.0.0 255.0.0.0 2.2.2.1 Int tunnel 0 Ip address 30.1.1.2 255.0.0.0 Tunnel source s0 Tunnel destination 2.2.2.1 Tunnel mode gre ip</pre>

Verification:

R1#show ip int brief

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	10.1.1.1	YES	manual	up	up
Serial1/0	2.2.2.1	YES	manual	up	up
Tunnel0	30.1.1.1	YES	manual	up	up

The output displays int tunnel 0 status is up and protocol status is also up, which indicates that GRE tunnel configuration is successful.

R1#ping 30.1.1.2

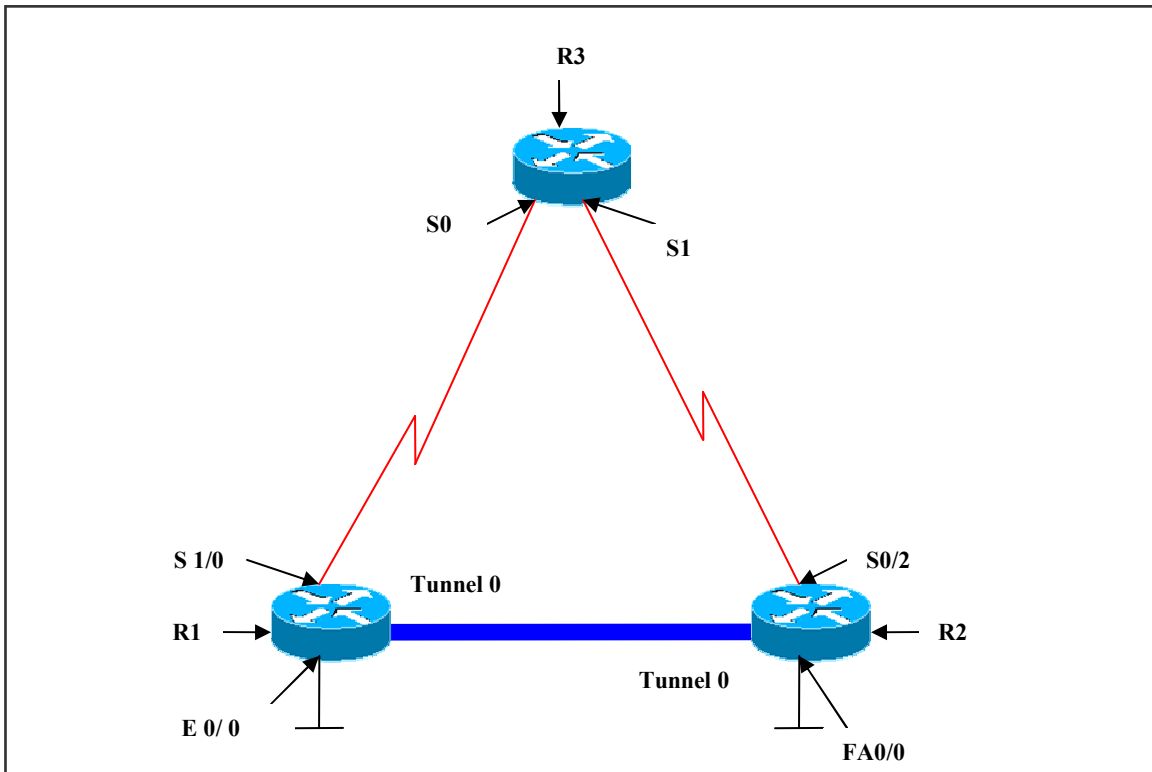
Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 30.1.1.2, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 36/39/44 ms

Lab 5 – GRE Tunneling Using Three Routers With no Routing in the Middle Router



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.1	255.0.0.0
E 0/0	10.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0/2	3.3.3.2	255.0.0.0

Fa 0/0	20.1.1.1	255.0.0.0
--------	----------	-----------

Lab Objective:

Task

Create interface tunnel 0 on R1 & R2.
 Verify connectivity.
 Configure OSPF routing protocol on R1 & R2 only.
 Verify if routes are visible in the routing table of R1 & R2.

<p>R1</p> <p>Ip route 3.0.0.0 255.0.0.0 2.2.2.2</p> <p>Int tunnel 0 Ip address 30.1.1.1 255.0.0.0 Tunnel source s1/0 Tunnel destination 3.3.3.2 Tunnel mode gre ip</p>	<p>R2</p> <p>Ip route 2.0.0.0 255.0.0.0 3.3.3.1</p> <p>Int tunnel 0 Ip address 30.1.1.2 255.0.0.0 Tunnel source s0/2 Tunnel destination 2.2.2.1 Tunnel mode gre ip</p>
---	---

Verification:

R1#show ip int brief

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	10.1.1.1	YES	manual	up	up
Serial1/0	2.2.2.1	YES	manual	up	up
Tunnel0	30.1.1.1	YES	manual	up	up

The output displays tunnel 0 is up

Task

Configure OSPF routing protocol on R1 & R2 only.
 Verify if routes are visible in the routing table of R1 & R2.

<p>R1</p> <p>Router ospf 1 Network 10.0.0.0 0.255.255.255 area 0 Network 30.0.0.0 0.255.255.255 area 0</p>	<p>R2</p> <p>Router ospf 1 Network 20.0.0.0 0.255.255.255 area 0 Network 30.0.0.0 0.255.255.255 area 0</p>
---	---

Verification:

R1#show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
30.1.1.2	0	FULL/ -	00:00:37	30.1.1.2	Tunnel0

The output displays tunnel ip address (30.1.1.2) as the neighbor-id

R1#show ip route

```
C 2.0.0.0/8 is directly connected, Serial1/0
S 3.0.0.0/8 [1/0] via 2.2.2.2
O 20.0.0.0/8 [110/11112] via 30.1.1.2, 00:15:35, Tunnel0
C 10.0.0.0/8 is directly connected, Ethernet0/0
C 30.0.0.0/8 is directly connected, Tunnel0
```

The output displays 'O' (OSPF) route for network 20.0.0.0 carrying via the tunnel IP address (30.1.1.2)

This indicates that routes are traveling via the gre tunnel.

Though, there is another router in between R1 & R2, the tunnel appears as a point-to-point link.

R1#traceroute 20.1.1.1

Type escape sequence to abort.

Tracing the route to 20.1.1.1

```
1 30.1.1.2 40 msec 40 msec *
```

The output displays the trace as 1 hop because of the gre tunnel.

Lab 6 – Configuring GRE Over IPSEC

(Scenario Based On Lab 5)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
S 1/0	2.2.2.1	255.0.0.0
E 0/0	10.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
S 0	2.2.2.2	255.0.0.0
S 1	3.3.3.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
S 0/2	3.3.3.2	255.0.0.0
Fa 0/0	20.1.1.1	255.0.0.0

Lab Objective:

Task

Create interface tunnel 0 on R1 and R2
Verify connectivity
Configure OSPF on R1 & R2 only.
Verify if the routes are traveling via the tunnel.

R1	R2
Ip route 3.0.0.0 255.0.0.0 2.2.2.2	Ip route 2.0.0.0 255.0.0.0 3.3.3.1
Int tunnel 0 Ip address 30.1.1.1 255.0.0.0 Tunnel source s1/0 Tunnel destination 3.3.3.2 Tunnel mode gre ip	Int tunnel 0 Ip address 30.1.1.2 255.0.0.0 Tunnel source s0/2 Tunnel destination 2.2.2.1 Tunnel mode gre ip

<pre>Router ospf 1 Network 10.0.0.0 0.255.255.255 area 0 Network 30.0.0.0 0.255.255.255 area 0</pre>	<pre>Router ospf 1 Network 20.0.0.0 0.255.255.255 area 0 Network 30.0.0.0 0.255.255.255 area 0</pre>
--	--

Task

Configure IPSec from R1 to R2 on the GRE tunnel.

<p>R1</p> <pre>Crypto isakmp enable Crypto isakmp policy 10 Encryption 3des Hash md5 Authentication pre-share Group1 Crypto isakmp key cisco123 address 3.3.3.2 Crypto ipsec transform-set set1 esp-des Access-list 101 permit ip 10.0.0.0 0.255.255.255 20.0.0.0 0.255.255.255 Crypto map map1 10 ipsec-isakmp Set peer 3.3.3.2 Set transform-set set1 Match address 101 Int s1/0 Crypto map map1 Int tunnel 0 Crypto map map1</pre>	<p>R2</p> <pre>Crypto isakmp enable Crypto isakmp policy 20 Encryption 3des Hash md5 Authentication pre-share Group1 Crypto isakmp key cisco123 address 2.2.2.1 Crypto ipsec transform-set set1 esp-des Access-list 101 permit ip 20.0.0.0 0.255.255.255 10.0.0.0 0.255.255.255 Crypto map map1 10 ipsec-isakmp Set peer 2.2.2.1 Set transform-set set1 Match address 101 Int s0/2 Crypto map map1 Int tunnel 0 Crypto map map1</pre>
---	---

Verification:

R1#show crypto isakmp sa

dst	src	state	conn-id	slot
3.3.3.2	2.2.2.1	QM_IDLE	1	0

The output displays the current IKE session and QM_IDLE indicates that the IKE is active.

R1#show crypto ipsec sa

interface: Serial1/0

Crypto map tag: **map1**, local addr. **2.2.2.1**

protected vrf:

local ident (addr/mask/prot/port): (10.0.0.0/255.0.0.0/0/0)

remote ident (addr/mask/prot/port): (20.0.0.0/255.0.0.0/0/0)

current_peer: **3.3.3.2**:500

PERMIT, flags={origin_is_acl,}

#pkts encaps: 99, #pkts encrypt: 99, #pkts digest 0

#pkts decaps: 99, #pkts decrypt: 99, #pkts verify 0

#pkts compressed: 0, #pkts decompressed: 0

#pkts not compressed: 0, #pkts compr. failed: 0

#pkts not decompressed: 0, #pkts decompress failed: 0

#send errors 1, #recv errors 0

---output omitted---

interface: Tunnel0

Crypto map tag: **map1**, local addr. **2.2.2.1**

protected vrf:

local ident (addr/mask/prot/port): (10.0.0.0/255.0.0.0/0/0)

remote ident (addr/mask/prot/port): (20.0.0.0/255.0.0.0/0/0)

current_peer: **3.3.3.2**:500

PERMIT, flags={origin_is_acl,}

#pkts encaps: 99, #pkts encrypt: 99, #pkts digest 0

#pkts decaps: 99, #pkts decrypt: 99, #pkts verify 0

#pkts compressed: 0, #pkts decompressed: 0

#pkts not compressed: 0, #pkts compr. failed: 0

#pkts not decompressed: 0, #pkts decompress failed: 0

#send errors 1, #recv errors 0

---output omitted---

The output displays current settings used by current SA's. Non-zero encryption and decryption statistics can indicate a working set of IPsec SA's.

Lab 7 – Configuring GRE Over IPSEC Site-to-Site Tunnel Using SDM


Configure GRE over IPsec side-to-side tunnel using SDM

- Click configure icon to enter configuration page.
- Click VPN icon.
- Choose site-to-site VPN wizard
- Click create site-to-site VPN tab.
- Click create secure GRE tunnel (GRE over IPsec) radio button.
- Click launch the selected task.

GRE tunnel information

- Specify GRE tunnel source IP address and destination IP address.
- Define the IP address & subnet mask that are applied to virtual point-to-point link.
- Click next
- Optionally we can create second GRE tunnel and click next.
- IPsec – specific parameters :
 - Click preshared keys authentication method radio button.
 - Specify preshared key and click next
- IKE proposals
 - Click add button and create custom IKE policy & click next.
- Transform-set
 - Click add button and specify parameters and click next.
- Select the routing protocol
 - Select OSPF routing protocol radio button.
 - Define router OSPF process ID & area number for tunnel.
 - Define one or more local subnets to be advertised to OSPF neighbors.
- At the end, the wizard will present a summary of the configured parameters and click finish to complete the configuration.
- Verification
 - Click test tunnel button and also click monitor icon to display the status of the tunnel.

Lab 8 – Configure Cisco VPN Client (Remote Access VPN)



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	100.100.100.100	255.255.255.0
Fa 0/0	20.1.1.1	255.0.0.0
PC 1	20.1.1.20	255.0.0.0

Lab Objective:

Task

Configure R1 with the VPN server configuration.

Install Cisco VPN client software on the PC.

Create a loopback 0 (100.100.100.100) and try sending traffic to this address from the PC and verify if the VPN tunnel is established or not.

R1

```
aaa new-model
aaa authentication login list1 local
aaa authorization network list2 local
```

```
Username user1 password user1
```

```
Crypto isakmp policy 10
Encryption 3des
Hash md5
```



```
Authentication pre-share
Group 2

IP local pool p1 30.1.1.1 30.1.1.100

Crypto isakmp client configuration group group1
Key cisco123
Pool p1

crypto ipsec transform-set set1 esp-3des esp-md5-hmac

Crypto dynamic-map dmap1 10
Set transform-set set1
Reverse-route

Crypto map map1 10 ipsec-isakmp dynamic map1
Crypto map map1 client configuration address respond
Crypto map map1 client authentication list list1
Crypto map map1 isakmp authorization list list2

Int fa0/0
Crypto map map1

Int loopback 0
IP address 100.100.100.100 255.0.0.0
```

PC client S/W installation :

- Install a Cisco VPN client on the remote user PC.
- Start → programs → Cisco systems VPN client → click VPN client.
- VPN client application starts.
- Click the “new” icon in the toolbar.
- Enter a name for the new connection enter field.
- Enter description of this connection in the description field.
- Enter the hostname or IP address of the remote VPN device (server) (20.1.1.1) that we want to access.
- Under the authentication tab, select the group authentication radio button.
- In the name field, enter the name of the IPsec group (group1) to which you belong.
- In the password field, enter the password (cisco123) for IPsec group.
- Verify password in the confirm password field.
- Save the connection entry by clicking the save button.

- Before we connect to the server from the client, send traffic through the path where tunnel is established
 - Therefore, ping from PC (20.1.1.20) to loopback 0 (100.100.100.100)
 - PC > ping 100.100.100.100 -t **(the output display that reply received from the address 100.100.100.100).**
 - Verify by clicking connect on the VPN client application. The VPN client window prompts for username and password
 - In the user name field enter username (user1)
 - In the password field enter password (user1)
 - As soon as you enter the above details the connection is established.

Verification:

R1#show crypto isakmp sa

dst	src	state	conn-id	slot
20.1.1.1	20.1.1.20	QM_IDLE	2	0

The output displays quick mode state and a connection id that indicates that tunnel is established

R1#show crypto ipsec sa

interface: FastEthernet0/0

Crypto map tag: map1, local addr. 20.1.1.1

protected vrf:

local ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)

remote ident (addr/mask/prot/port): (30.1.1.1/255.255.255.255/0/0)

current_peer: 20.1.1.20:500

PERMIT, flags={}

#pkts encaps: 110, #pkts encrypt: 110, #pkts digest 110

#pkts decaps: 153, #pkts decrypt: 153, #pkts verify 153

#pkts compressed: 0, #pkts decompressed: 0

#pkts not compressed: 0, #pkts compr. failed: 0

#pkts not decompressed: 0, #pkts decompress failed: 0

#send errors 0, #recv errors 0

The output displays that packets passing the tunnel are encrypted and also decrypted. This indicates the tunnel created is secure, thus giving access to the remote clients to the server on the internet securely.

Lab 9 – Configure Cisco Easy VPN

(Scenario Based On Lab 8)

Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
Loopback 0	100.100.100.100	255.255.255.0
Fa 0/0	20.1.1.1	255.0.0.0

Lab Objective:

Task

Configure R1 as easy VPN server.

Do not configure authentication and username and password.

Create loopback address to send traffic from the client PC to verify the tunnel.

R1

```
aaa new-model
aaa authorization network list2 local

Crypto isakmp policy 10
Encryption 3des
Hash md5
Authentication pre-share
Group 2

IP local pool p1 30.1.1.1 30.1.1.100

Crypto isakmp client configuration group group1
Key cisco123
Pool p1

crypto ipsec transform-set set1 esp-3des esp-md5-hmac

Crypto dynamic-map dmap1 10
Set transform-set set1
```

Reverse-route

```
Crypto map map1 10 ipsec-isakmp dynamic map1  
Crypto map map1 client configuration address respond  
Crypto map map1 isakmp authorization list list2
```

```
Int fa0/0  
Crypto map map1
```

```
Int loopback 0  
IP address 100.100.100.100 255.0.0.0
```

PC : Easy VPN client

- Install a Cisco VPN client on the PC.
- Start → programs → Cisco systems VPN client → click VPN client.
- VPN client application starts.
- Click the “new” icon in the toolbar.
- Enter name and description for the connection entry.
- Enter the hostname or IP address of the server.
- Under the authentication tab : enter the name of the IPSec group (group1) and password for the group (cisco123).
- Save the connection entry
 - Before we connect to the server from the client PC, send traffic through the path where tunnel is established.
 - Therefore, ping from PC (20.1.1.20) to loopback 0 (100.100.100.100)
 - PC > ping 100.100.100.100 -t
 - The output displays that replies are received from the address
- Verify by clicking connect on the application. It does not ask for username or password.

Verification :

R1#show crypto isakmp sa

dst	src	state	conn-id	slot
20.1.1.1	20.1.1.20	QM_IDLE	2	0

The output displays quick mode state and a connection id that indicates that tunnel is established

R1#show crypto ipsec sa

interface: FastEthernet0/0

Crypto map tag: map1, local addr. 20.1.1.1

protected vrf:

local ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)

remote ident (addr/mask/prot/port): (30.1.1.1/255.255.255.255/0/0)

current_peer: 20.1.1.20:500

PERMIT, flags={}

#pkts encaps: 110, #pkts encrypt: 110, #pkts digest 110

#pkts decaps: 153, #pkts decrypt: 153, #pkts verify 153

#pkts compressed: 0, #pkts decompressed: 0

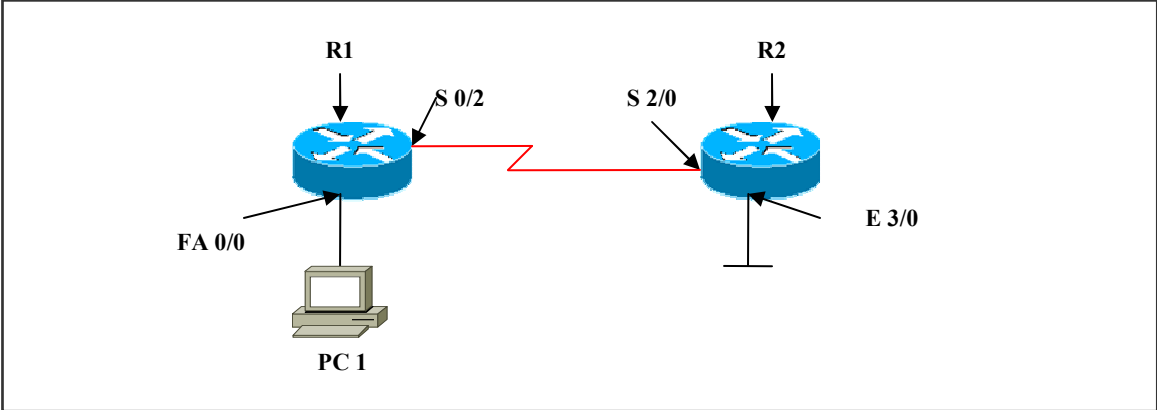
#pkts not compressed: 0, #pkts compr. failed: 0

#pkts not decompressed: 0, #pkts decompress failed: 0

#send errors 0, #recv errors 0

The output displays that packets passing the tunnel are encrypted and also decrypted. This indicates the tunnel created is secure, thus giving access to the remote clients to the server on the internet securely.

Lab 10 – Configure Easy-VPN Server and Client as Router



Interface IP Address Configuration

R2

Interface	IP Address	Subnet Mask
S 2/0	1.1.1.2	255.0.0.0
E 3/0	20.1.1.1	255.0.0.0

R1

Interface	IP Address	Subnet Mask
S 0/2	1.1.1.1	255.0.0.0
Fa 0/0	10.1.1.1	255.0.0.0
PC 1	10.1.1.20	255.0.0.0

Lab Objective:

Task

Configure the client in network-extension mode.
 Create reverse-route on the server and a static route in client to reach server.
 Do not telnet until the VPN tunnel is established.

```

R2

aaa new-model
aaa authentication login xyz none
    
```

```
aaa authorization network lauthor local
```

```
Crypto isakmp policy 10  
Encryption 3des  
Hash md5  
Authentication pre-share  
Group 2
```

```
IP local pool p1 30.1.1.1 30.1.1.100
```

```
Crypto isakmp client configuration group group1  
Key cisco123  
Pool p1
```

```
crypto ipsec transform-set set1 esp-3des esp-md5-hmac
```

```
Crypto dynamic-map dmap1 10  
Set transform-set set1  
Reverse-route
```

```
Crypto map map1 10 ipsec-isakmp dynamic map1  
Crypto map map1 client configuration address respond  
Crypto map map1 isakmp authorization list lauthor
```

```
Line vty 0 4  
Login authentication xyz
```

```
Int s2/0  
Crypto map map1
```

R1

```
Crypto ipsec client ezvpn vpn1  
Group group1 key cisco123  
Peer 1.1.1.2  
Connect auto  
Mode network-extension
```

```
Int fa0/0  
Crypto ipsec client ezvpn vpn1 inside
```

```
Int s0/2
Crypto ipsec client ezvpn vpn1 outside

Ip route 20.0.0.0 255.0.0.0 1.1.1.2
```

Verification :

R2#show crypto isakmp sa

```
dst          src          state        conn-id  slot
1.1.1.2     1.1.1.1     QM_IDLE      1        0
```

The output displays a connection-id and quick mode state denoting SA is created

R2#show crypto ipsec sa

```
interface: Serial2/0
  Crypto map tag: map1, local addr. 1.1.1.2

  protected vrf:
  local ident (addr/mask/prot/port): (0.0.0.0/0.0.0.0/0/0)
  remote ident (addr/mask/prot/port): (10.0.0.0/255.0.0.0/0/0)
  current_peer: 1.1.1.1:500
    PERMIT, flags={}
    #pkts encaps: 50, #pkts encrypt: 50, #pkts digest 50
    #pkts decaps: 51, #pkts decrypt: 51, #pkts verify 51
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
    #pkts not decompressed: 0, #pkts decompress failed: 0
    #send errors 0, #recv errors 0
```

The output displays packets being encrypted and decrypted.

R1#show crypto ipsec client ezvpn

Easy VPN Remote Phase: 2

```
Tunnel name : vpn1
Inside interface list: FastEthernet0/0,
Outside interface: Serial0/2
Current State: IPSEC_ACTIVE
Last Event: SOCKET_UP
```


The output displays current state for IPSec as active that indicates the tunnel is established.

If mode client configured on the client side, the client does nat translations.

R1#show ip nat translations

Pro	Inside global	Inside local	Outside local	Outside global
icmp	30.1.1.1:512	20.1.1.20:512	10.1.1.1:512	10.1.1.1:512

The output displays the client doing nat translations. This happens only if the client is configured in 'client mode'.

R2#show ip route

Gateway of last resort is not set

```
C 1.0.0.0/8 is directly connected, Serial2/0
C 10.0.0.0/8 is directly connected, Ethernet3/0
 30.0.0.0/32 is subnetted, 1 subnets
S 30.1.1.1 [1/0] via 1.1.1.2
```

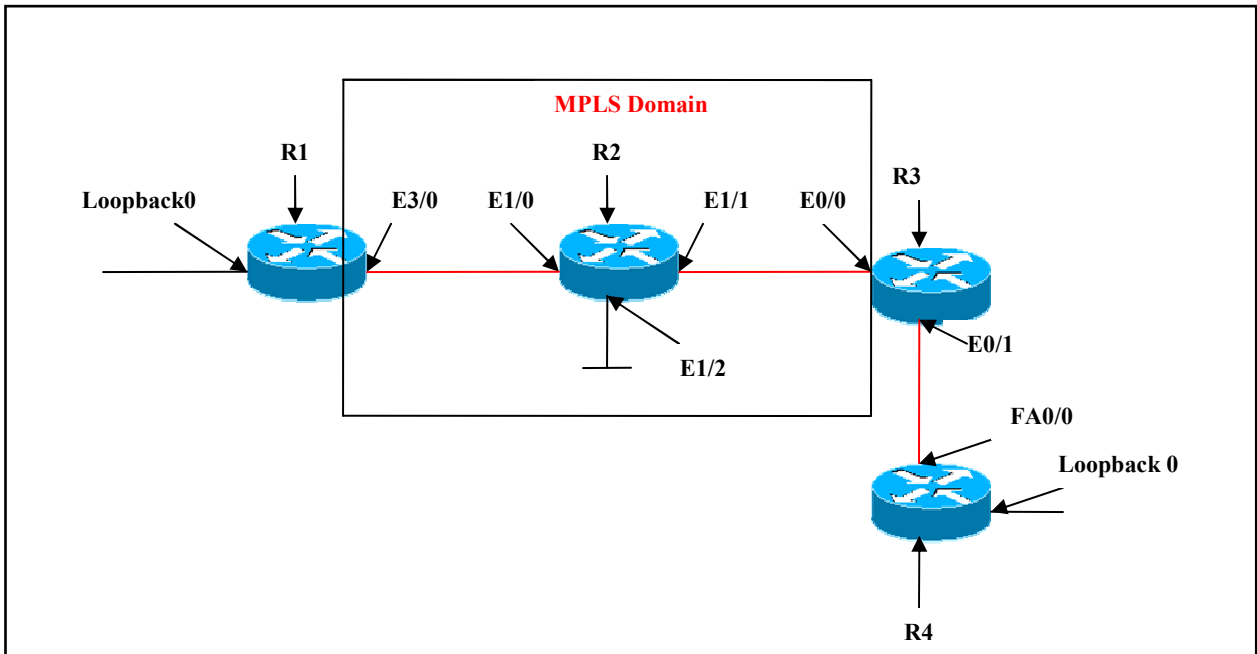
The static route for 30.0.0.0 network is automatically created because of the reverse-route configured in the server.

R2#show ip local pool

Pool	Begin	End	Free	In use
p1	30.1.1.1	30.1.1.100	100	1

The output shows the ip address in the pool.

Lab 11– Configure Frame Mode MPLS



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	1.1.1.1	255.0.0.0
Loopback 0	10.1.1.1	255.0.0.0

R2

Interface	IP Address	Subnet Mask
E 1/0	1.1.1.2	255.0.0.0
E 1/1	2.2.2.1	255.0.0.0
E 1/2	20.1.1.1	255.0.0.0

R3

Interface	IP Address	Subnet Mask
E 0/0	2.2.2.2	255.0.0.0
E 0/1	3.3.3.1	255.0.0.0

R4

Interface	IP Address	Subnet Mask
FA 0/0	3.3.3.2	255.0.0.0
Loopback 0	30.1.1.1	255.0.0.0

Lab Objective:**Task**

Configure OSPF in Area 0 in the MPLS domain as per the scenario and EIGRP AS 200 on R1 (Loopback 0) and EIGRP AS 100 on R3 (E0/1), R4 (FA0/0, Loopback 0).

Mutually redistribute these two routing protocols.

Configure MPLS on R1 (E 3/0), R2 (E1/0, E1/1, E 1/2) and R3 (E0/0).

Enable CEF on routers configured in MPLS domain.

R1 Ip cef Interface e3/0 Mpls ip Mpls label protocol ldp Mpls mtu 1512	R3 Ip cef Interface e0/0 Mpls ip Mpls label protocol ldp Mpls mtu 1512
R2 Ip cef Interface e1/0 Mpls ip Mpls label protocol ldp Mpls mtu 1512 Interface e1/1 Mpls ip Mpls label protocol ldp Mpls mtu 1512 Interface e1/2 Mpls ip	

```
Mpls label protocol ldp
Mpls mtu 1512
```

Verification :

R1#show mpls ldp neighbor

```
Peer LDP Ident: 20.1.1.1:0; Local LDP Ident 10.1.1.1:0
TCP connection: 20.1.1.1.11043 - 10.1.1.1.646
State: Oper; Msgs sent/rcvd: 99/99; Downstream
Up time: 01:18:58
LDP discovery sources:
  Ethernet3/0, Src IP addr: 1.1.1.2
Addresses bound to peer LDP Ident:
  1.1.1.2    2.2.2.1    20.1.1.1
```

The output displays the neighbor for R1.

R1#show mpls forwarding-table

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Pop tag	2.0.0.0/8	0	Et3/0	1.1.1.2
17	Pop tag	20.0.0.0/8	0	Et3/0	1.1.1.2
18	16	3.0.0.0/8	0	Et3/0	1.1.1.2
19	18	30.0.0.0/8	0	Et3/0	1.1.1.2

The output displays the local tags attached to the router, outgoing tags and the outgoing interface.

R1#show mpls ldp bindings

```
tib entry: 1.0.0.0/8, rev 2
  local binding: tag: imp-null
  remote binding: tsr: 20.1.1.1:0, tag: imp-null
tib entry: 2.0.0.0/8, rev 6
  local binding: tag: 16
  remote binding: tsr: 20.1.1.1:0, tag: imp-null
tib entry: 3.0.0.0/8, rev 10
  local binding: tag: 18
  remote binding: tsr: 20.1.1.1:0, tag: 16
tib entry: 10.0.0.0/8, rev 4
  local binding: tag: imp-null
  remote binding: tsr: 20.1.1.1:0, tag: 17
tib entry: 20.0.0.0/8, rev 8
  local binding: tag: 17
  remote binding: tsr: 20.1.1.1:0, tag: imp-null
```

tib entry: 30.0.0.0/8, rev 12
local binding: tag: 19
remote binding: tsr: 20.1.1.1:0, tag: 18

The output displays local bindings of the tag to the router and also the remote bindings of the same tag by its neighbor.

R2#sh mpls interfaces

Interface	IP	Tunnel	Operational
Ethernet1/0	Yes (ldp)	No	Yes
Ethernet1/1	Yes (ldp)	No	Yes
Ethernet1/2	Yes (ldp)	No	Yes

The output displays the MPLS configured interfaces on the router.

R2#sh mpls label range

Downstream Generic label region: Min/Max label: 16/100000


The output displays the range for the labels from 16, as 1 -15 are reserved.

R2#show ip cef

Prefix	Next Hop	Interface
0.0.0.0/0	drop	Null0 (default route handler entry)
0.0.0.0/32	receive	
1.0.0.0/8	attached	Ethernet1/0
1.0.0.0/32	receive	
1.1.1.1/32	1.1.1.1	Ethernet1/0
1.1.1.2/32	receive	
1.255.255.255/32	receive	
2.0.0.0/8	attached	Ethernet1/1
2.0.0.0/32	receive	
2.2.2.1/32	receive	
2.2.2.2/32	2.2.2.2	Ethernet1/1
2.255.255.255/32	receive	
3.0.0.0/8	2.2.2.2	Ethernet1/1
10.0.0.0/8	1.1.1.1	Ethernet1/0
20.0.0.0/8	attached	Ethernet1/2
20.0.0.0/32	receive	
20.1.1.1/32	receive	
20.255.255.255/32	receive	
30.0.0.0/8	2.2.2.2	Ethernet1/1
224.0.0.0/4	drop	
224.0.0.0/24	receive	
255.255.255.255/32	receive	

The output displays the summary of the FIB table.

Lab 12 – Configure an SSH Server for Secure Management and Reporting



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 1/0	10.1.1.1	255.0.0.0

Lab Objective:

Task

Configure R1 as SSH server.
Install SSH client software on the PC.

R1

```
Username user1 password user1
Enable secret Cisco
aaa new-model
aaa authentication login lauth local
Ip domain-name netmetrics
Crypto key generate rsa 512
```

```
Line vty 0 4
Login authentication lauth
Transport input ssh
```

```
Ip ssh time-out 120
Ip ssh authentication-retries 3
```

Next step is install putty software, i.e. it is a third party tool to log into the router.

Page 262 of 315

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com

All contents are copyright @ 2007-2010

All rights reserved.

- Click run and install on the PC.
- Putty configuration window opens: → click session → the window prompts for basic options for your putty session.
- Specify the destination you want to connect to : -
 - Hostname (or IP address) : 10.1.1.1
 - Port : 22
- Connection type – select the radio button for SSH.
- Click open → software window prompts for username & password.
- When all details are specified, you get access to the router.

Verification :

R-1(config) #do show crypto key mypubkey rsa

```
% Key pair was generated at: 00:21:16 UTC JUL 7 2007
Key name: R-1.netmetrics
Usage: General Purpose Key
Key is not exportable.
Key Data:
 305C300D 06092A86 4886F70D 01010105 00034B00 30480241 00B2E7D3
1328D75A
 EF058A59 E6A4D4A1 44015A01 10A0B0B9 6B286D32 B889182C 5DFDAC8F
1B289436
 D08768DD D9B0D192 24B94D14 5D0F077E 478AD8EB 6026D789 FB020301 0001
% Key pair was generated at: 00:21:19 UTC JUL 7 2007
Key name: R-1.netmetrics.server
Usage: Encryption Key
Key is exportable.
Key Data:
 307C300D 06092A86 4886F70D 01010105 00036B00 30680261 00D11809
646CA0C6
 10B53FF6 1C372194 ABBC2720 8BFCFB5F 95B7BF71 0BD4B5DF B11BFB66
E9A4BC92
 1A835176 79F97BF8 4A59E21F 5A0DD904 67D9184F F513FFC5 9E279965
9EF0483D
 51242BDC 2DA4F53C 00105C2C 0389F9E1 1994DB91 3EEC6BE2 AD020301 0001
```

The output displays the generated key.

From PC1:

```
C:\Documents and Settings>telnet 10.1.1.1
Connecting To 10.1.1.1...Could not open connection to the host, on port 23: Conn
ect failed
```

If user wants to telnet to the router, access is denied as it is configured as SSH .

R-1#show users

Line	User	Host(s)	Idle	Location
* 0 con 0		idle	00:00:00	
130 vty 0	user1	idle	00:00:21	10.1.1.2

The output displays the user (10.1.1.2) via SSH.

Lab 13 – Configure Syslog Logging

(Scenario Based On Lab 12)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 1/0	10.1.1.1	255.0.0.0
PC 1	10.1.1.2	255.0.0.0

PC 1

IP Address	Subnet Mask
10.1.1.2	255.0.0.0

Lab Objective:

Task

Configure R1 to send log messages to the Syslog server.
Install Syslog server software on the PC.

R1

Logging on
Logging host 10.1.1.2
Logging trap debugging

PC 1: -

Install kiwi syslog Daemon 8.2.18 installation on PC1.
Click run → agree the license agreement → select install kiwi syslog Daemon as an application → click next → click install → select the checkbox “run kiwi syslog Daemon 8.2.18 and finally click finish.

Verification:

Logged Messages.

Date	Time	Priority	Hostname	Message
07-07-2007	15:37:58	Local7.Debug	10.1.1.1	213: ACK 53359694 WIN 64996
07-07-2007	15:37:58	Local7.Debug	10.1.1.1	212: *July 1 01:07:43.859: tcp130: I ESTAB 10.1.1.4:1041 10.1.1.1:22 seq 2008571168
07-07-2007	15:37:57	Local7.Debug	10.1.1.1	211: DATA 20 ACK 2008571168 PSH WIN 3673
07-07-2007	15:37:57	Local7.Debug	10.1.1.1	210: *July 1 01:07:43.859: tcp130: O ESTAB 10.1.1.4:22 10.1.1.1:1041 seq 53359674
07-07-2007	15:37:57	Local7.Debug	10.1.1.1	209: DATA 20 ACK 2008571168 PSH WIN 3673
07-07-2007	15:37:57	Local7.Debug	10.1.1.1	208: *July 1 01:07:43.823: tcp130: O ESTAB 10.1.1.4:22 10.1.1.1:1041 seq 53359654
07-07-2007	15:37:57	Local7.Debug	10.1.1.1	207: DATA 20 ACK 53359654 PSH WIN 65036

The output displays the logged messages in the Syslog server. Any configuration made via console is logged onto the syslog server. Any unauthorized access, can also be logged and viewed on the syslog server.

R-1#show logging

Syslog logging: enabled (0 messages dropped, 2 messages rate-limited, 0 flushes, 0 overruns, xml disabled)

Console logging: level debugging, 215 messages logged, xml disabled

Monitor logging: level debugging, 0 messages logged, xml disabled

Buffer logging: disabled, xml disabled

Logging Exception size (4096 bytes)


Count and timestamp logging messages: disabled

Trap logging: level debugging, 222 message lines logged

Logging to 10.1.1.2, 123 message lines logged, xml disabled

The output displays that the syslog logging is enabled , ip address of the syslog server and the number of messages logged.

Lab 14 – Configure SNMP



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	50.1.1.1	255.0.0.0

PC 1

IP Address	Subnet Mask
50.1.1.2	255.0.0.0

Lab Objective:

Task

Configure SNMP server.

Install network management tool (solar winds.net) on the PC.

R1

```
Access-list 10 permit 50.1.1.2
Access-list 10 deny any any
Snmp-server community public ro 10
Snmp-server communit private rw 10
```

PC 1 :

-When you install solar winds.net network management tool on the pc, an IP network browser is also installed.

-Double click IP network browser → SNMP tool opens :

Page 267 of 315

NETMETRIC-SOLUTIONS


www.netmetric-solutions.com

All contents are copyright @ 2007-2010

All rights reserved.

- Enter hostname/ IP address: 50.1.1.1.
- Click scan device.
- The tool scans for the device with the IP address.
- You can see R1 with the IP address 50.1.1.1 and all the options and complete detailed information about the router.
- As SNMP is a network management tool, you can download the running configuration or startup-configuration:
 - Click node → click tools → click view cisco config.
- Prompts for router/switch:
 - Enter IP address of router (50.1.1.1)
 - Enter community string
- Select private as it is in read-write feature.
- Click downloads.
- Prompts for type of configuration file:
 - Click either running-config or startup-config accordingly
- Click ok.
- The tool copies the file using cisco-config-MIB, from R1 to PC1 through the SNMP tool.

Lab 15 – Configuring NTP



Interface IP Address Configuration

R1 (MASTER)

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

R2 (CLIENT)

Interface	IP Address	Subnet Mask
E 1/0	10.1.1.2	255.0.0.0

Lab Objective:

Task

Set the clock to local current time on the master router (R1).
Configure R1 as NTP master and R2 as NTP client.

R1	R2
Clock set 16:10:30 07 July 2007 Clock timezone India +5 Ntp master 5 Ntp authentication-key 1 md5 cisco123 Ntp peer 10.1.1.2 key 1 Int e3/0 Ntp broadcast	Ntp authentication - key 1 md5 cisco123 Ntp trusted – key 1 Ntp server 10.1.1.1 key 1 Int e1/0 Ntp broadcast client

Verification:

R2#show clock

19:10:28.355 UTC Sat Jul 7 2007

The output displays current time of R2 before NTP configuration.
Manually change the clock and check if the time is synchronizing :

R1#show clock

07:22:08.231 UTC Wed Sep 19 2007


Verify that the clock has synchronized according to the server time :

R2#show clock

07:21:46.775 UTC Wed Sep 19 2007

The output displays that the client has synchronized with the server time.

Lab 16– Configuring AAA on Cisco Routers



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

PC 1

IP Address	Subnet Mask
10.1.1.10	255.0.0.0

Lab Objective:

Task

Configure the AAA server on R1.

Configure AAA login authentication.

Install Cisco secure ACS 4.0 version on PC 1 and create users1 user2 and user3

R1

```
aaa new-model
Tacacs-server host 10.1.1.10 single-connection
Tacacs-server key cisco123
aaa authentication login default group tacacs+ local
aaa authentication login lauth group tacacs+
username user1 password user1
```

```
Line vty 0 4
Login authentication lauth
```

When the tool is installed on PC 1, it prompts for passwords, specify the same password configured as the tacacs-server key on the router (cisco123).

Follow the steps in configuring users and other parameters :

User setup

Add user → username : user 1 & password : user 1

User account is created for user1 and click submit

Click list all users.

-The output displays all the users and their status

Network configuration

Add AAA client

Client : Router

AAA client

IP address : 10.1.1.1

Key : Cisco123

Authenticate using : TACACS+

Click submit + apply button to save the entry

Add AAA server

Server : PC

AAA server

IP add : 10.1.1.10

Key : cisco123

AAA server type : TACACS+

Traffic type : inbound / outbound

Click submit + apply

Verification :

Telnet from PC 1 (10.1.1.10) to router R1:

The router prompts for username and password:

Specify the username and password created on the TACACS+ Server.

Authentication is approved and gains access to R1.

R1 # debug aaa authentication

Jul 7 20:29:49.139: AAA: parse name=tty0 idb type=-1 tty=-1

*Jul 7 20:29:49.139: AAA: name=tty0 flags=0x11 type=4 shelf=0 slot=0 adapter=0 port=0 channel=0

*Jul 7 20:29:49.139: AAA/MEMORY: create_user (0x65477DC0) user='raduser1' ruser='NULL' ds=0 port='tty0' rem_addr='async' authen_type=ASCII service=ENABLE priv=15 initial_task_id='0', vrf=(id=0)

*Jul 7 20:29:49.139: AAA/AUTHEN/START (2180557774): port='tty0' list="" action=LOGIN service=ENABLE

*Jul 7 20:29:49.143: AAA/AUTHEN/START (2180557774): console enable - default to enable password (if any)

*Jul 7 20:29:49.143: AAA/AUTHEN/START (2180557774): Method=ENABLE

*Jul 7 20:29:49.143: AAA/AUTHEN(2180557774): can't find any passwords
*Jul 7 20:29:49.143: AAA/AUTHEN(2180557774): Status=ERROR
*Jul 7 20:29:49.143: AAA/AUTHEN/START (2180557774): Method=NONE
*Jul 7 20:29:49.143: AAA/AUTHEN(2180557774): Status=PASS
The output displays the status **PASS** indicates successful authentication.

Task 2

(Scenario Based On Task 1)

Configure the AAA server on R1.

Configure AAA login authentication.

Install Cisco secure ACS 4.0 version on PC 1 and create users1 user2 and user3

R1

```
aaa new-model
Radius-server host 10.1.1.10
Radius-server key cisco123
aaa authentication login r1 group radius local
aaa authentication login default local group radius

Line vty 0 4
Login authentication r1
```

Install the Cisco Secure ACS on the PC and complete the parameters.

User setup

Add user → username : user 2 & password : user 2

username : user 3 & password : user 3

password authentication → specify ACS Internal Database

User account is created for user 2, user 3 and click submit

Click list all users.

-The output displays all the users and their status

Network Management :-

Add AAA client

Client : Router

AAA client

IP address : 10.1.1.1

Key : Cisco123

Authenticate using : RADIUS (Cisco IOS)

Click submit + apply button to save the entry

Add AAA server

Server : PC

AAA server

IP add : 10.1.1.10

Key : cisco123

AAA server type : RADIUS
Traffic type : inbound / outbound
Click submit + apply

Verification:

Telnet from PC 1 (10.1.1.10) to router R1:
The router prompts for username and password:
Specify the username and password created on the RADIUS
Server.
Authentication is approved and gains access to R1.

Router# debug aaa authentication

```
*Jul 7 20:12:53.355: AAA: parse name=tty0 idb type=-1 tty=-1
*Jul 7 20:12:53.355: AAA: name=tty0 flags=0x11 type=4 shelf=0 slot=0 adapter=0
port=0 channel=0
*Jul 7 20:12:53.359: AAA/MEMORY: create_user (0x65477DC0) user='raduser1' ruser
='NULL' ds0=0 port='tty0' rem_addr='async' authen_type=ASCII service=ENABLE priv
=15 initial_task_id='0', vrf= (id=0)
*Jul 7 20:12:53.359: AAA/AUTHEN/START (1211612866): port='tty0' list="" action=
LOGIN service=ENABLE
*Jul 7 20:12:53.359: AAA/AUTHEN/START (1211612866): console enable - default to
enable password (if any)
*Jul 7 20:12:53.359: AAA/AUTHEN/START (1211612866): Method=ENABLE
*Jul 7 20:12:53.359: AAA/AUTHEN(1211612866): can't find any passwords
*Jul 7 20:12:53.359: AAA/AUTHEN(1211612866): Status=ERROR
*Jul 7 20:12:53.359: AAA/AUTHEN/START (1211612866): Method=NONE
*Jul 7 20:12:53.359: AAA/AUTHEN(1211612866): Status=PASS
```

The output displays the status **PASS** indicates successful authentication.

Lab 17– DISABLING UNUSED CISCO ROUTERS USING NETWORK SERVICES AND INTERFACES

LOCKING DOWN ROUTERS WITH AUTO SECURE:

Router#**auto secure**

--- AutoSecure Configuration ---

*** AutoSecure configuration enhances the security of the router, but it will not make it absolutely resistant to all security attacks ***

AutoSecure will modify the configuration of your device. All configuration changes will be shown.

At any prompt you may enter '?' for help.
Use ctrl-c to abort this session at any prompt.

Gathering information about the router for AutoSecure

Is this router connected to internet? [no]: y

Enter the number of interfaces facing the internet [1]: 1

Interface	IP-Address	OK?	Method	Status	Protocol
Serial2/0	1.1.1.1	YES	manual	up	down
Ethernet3/0	20.1.1.1	YES	manual	up	up

Enter the interface name that is facing the internet: Ethernet3/0

Securing Management plane services...

Disabling service finger
Disabling service pad
Disabling udp & tcp small servers
Enabling service password encryption
Enabling service tcp-keepalives-in
Enabling service tcp-keepalives-out
Disabling the cdp protocol
Disabling the bootp server
Disabling the http server
Disabling the finger service
Disabling source routing
Disabling gratuitous arp

Here is a sample Security Banner to be shown at every access to device. Modify it to suit your enterprise requirements.

Authorized Access only

This system is the property of So-&-So-Enterprise.

Enter the security banner {Put the banner between k and k, where k is any character}:

% This system is the property of Netmetric Solutions.
Please Handle With Care %

Enable secret is either not configured or is the same as enable password

Enter the new enable secret: netmetrics

Confirm the enable secret : netmetrics

Enter the new enable password: solutions

Confirm the enable password: solutions

Configuration of local user database

Enter the username: User1

Enter the password: password

Confirm the password: password

Configuring AAA local authentication

Configuring Console, Aux and VTY lines for local authentication, exec-timeout, and transport

Securing device against Login Attacks

Configure the following parameters

Blocking Period when Login Attack detected: 300

Maximum Login failures with the device: 3

Maximum time period for crossing the failed login attempts: 60

Configuring interface specific AutoSecure services

Disabling the following ip services on all interfaces:

no ip redirects

no ip proxy-arp

no ip unreachable

no ip directed-broadcast

no ip mask-reply

Disabling mop on Ethernet interfaces

Securing Forwarding plane services...

Enabling CEF (This might impact the memory requirements for your platform)

Enabling unicast rpf on all interfaces connected
to internet

Tcp intercept feature is used prevent tcp syn attack
on the servers in the network. Create autosec_tcp_intercept_list
to form the list of servers to which the tcp traffic is to
be observed

Enable tcp intercept feature? [yes/no]: y

This is the configuration generated:

```
no service finger
no service pad
no service udp-small-servers
no service tcp-small-servers
service password-encryption
service tcp-keepalives-in
service tcp-keepalives-out
no cdp run
no ip bootp server
no ip http server
no ip finger
no ip source-route
no ip gratuitous-arps
no ip identd
banner motd ^C This system is the property of Netmetric Solutions.
Please Handle With Care ^C
security passwords min-length 6
security authentication failure rate 10 log
enable secret 5 $1$TnMh$9BZLJz5BhTyu9wjQJ9DXF/
enable password 7 105D061510031B040217
username User1 password 7 01030717481C091D25
aaa new-model
aaa authentication login local_auth local
line con 0
login authentication local_auth
exec-timeout 5 0
transport output telnet
line aux 0
login authentication local_auth
exec-timeout 10 0
transport output telnet
```

```
line vty 0 4
 login authentication local_auth
 transport input telnet
 login block-for 300 attempts 3 within 60
 service timestamps debug datetime msec localtime show-timezone
 service timestamps log datetime msec localtime show-timezone
 logging facility local2
 logging trap debugging
 service sequence-numbers
 logging console critical
 logging buffered
 interface Serial2/0
  no ip redirects
  no ip proxy-arp
  no ip unreachablees
  no ip directed-broadcast
  no ip mask-reply
 interface Serial2/1
  no ip redirects
  no ip proxy-arp
  no ip unreachablees
  no ip directed-broadcast
  no ip mask-reply
 interface Serial2/2
  no ip redirects
  no ip proxy-arp
  no ip unreachablees
  no ip directed-broadcast
  no ip mask-reply
 interface Serial2/3
  no ip redirects
  no ip proxy-arp
  no ip unreachablees
  no ip directed-broadcast
  no ip mask-reply
 interface Ethernet3/0
  no ip redirects
  no ip proxy-arp
  no ip unreachablees
  no ip directed-broadcast
  no ip mask-reply
  no mop enabled
 interface TokenRing3/0
  no ip redirects
  no ip proxy-arp
```

```
no ip unreachable
no ip directed-broadcast
no ip mask-reply
ip cef
access-list 100 permit udp any any eq bootpc
interface Ethernet3/0
 ip verify unicast source reachable-via rx allow-default 100
ip tcp intercept list autosec_tcp_intercept_list
ip tcp intercept drop-mode random
ip tcp intercept watch-timeout 15
ip tcp intercept connection-timeout 3600
ip tcp intercept max-incomplete low 450
ip tcp intercept max-incomplete high 550
!
end
```

```
Apply this configuration to running-config? [yes]: y
```

Router#sh flash

System flash directory:

File Length Name/status

1 26749788 c3640-js-mz.124-7.bin

2 944 pre autosec.cfg

[26750860 bytes used, 6279280 available, 33030140 total]

32768K bytes of processor board System flash (Read/Write)

Lab 18 – SECURITY CISCO ROUTER INSTALLATION AND ADMINISTRATIVE ACCESS

Page 279 of 315

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com

All contents are copyright @ 2007-2010

All rights reserved.

1. CONFIGURING ROUTER PASSWORDS:

- Every router needs a locally configured router for privilege access.
- Passwords are maintained on an AAA server.
- Different ways to gain administrative access to the router are :
 1. Console port
 2. Telnet
 3. Secure shell (SSH)
 4. Simple Network Management Protocol (SNMP)
 5. Cisco Security Device Manager (SDM) access using HTTP and HTTPS
- Passwords can be 1 to 25 characters in length.
- Passwords can include alphanumeric characters, uppercase and lowercase characters, symbols and spaces.
- Passwords cannot have a number as the first character.
- Password-leading spaces are ignored, but all spaces after the first character are not ignored.
- Best practice is to frequently change the passwords.

COMMANDS:

- (config) # enable secret **ccnp**
 1. Is used to enter the enable mode or privilege exec mode.
 2. This uses a one-way encryption hash based on MD5.
- (config) # enable password **ccnp**
 1. This is also used to enter the enable mode or privilege mode, but if “enable secret **ccnp** “configured will override the “enable password **ccnp** “.
 2. By default this is not encrypted in the router configuration.
 3. The virtual terminal password is not encrypted.

2. CONFIGURE THE LINE-LEVEL PASSWORD:

Configuration: The commands are same for line auxillary 0 and line vty 0 4

R1

```
Line console 0
Login
Password ccnp
```

- CONSOLE PORT :
 1. If a router is configured with the “*no service password-recovery*” , all access to the ROM monitor (ROMMON) is disabled.
- VTY LINES :
 1. You must configure a vty password before attempting to access the router using telnet.
 2. If you fail to set an enable password for the router, you will not be able to access privileged-exec mode using Telnet.
 3. Allow Telnet access from specific hosts only.
 4. You must configure passwords for all the vty lines on the router.
- Auxillary Lines :
 1. If you wish to turn off the EXEC process for the aux port, use the “*no exec*” command within the auxillary line configuration mode.

Configuration for line auxillary line 0 :

R1

```
Line aux 0
Modem input
Speed 9600
Transport input all
Flowcontrol hardware
Login
Password ccnp
```

3. PASSWORD MINIMUM LENGTH ENFORCEMENT:

- Cisco IOS software release 12.3(1) and later allows to set the minimum character length for all router passwords.

- It is recommended that you set your minimum password length to at least 10 characters.

Command:

R1

Config terminal

Enter configuration commands, one per line. End with CNTL/Z.

Security passwords min-length 10

Enable secret ccnp

% Password too short - must be at least 10 characters. Password configuration failed

Enable secret netmetrics

- If the password is not meeting the specified characters mentioned then an error message is displayed on the console as shown in the above configuration.

4. ENCRYPTING PASSWORDS:

- A PROPRIETARY Cisco algorithm based on Vigenere cipher (indicated by the number 7 when viewed in the configuration) allows the “*service password-encryption*” command to encrypt all passwords (except the previously encrypted enable secret passwords).
- When you remove the “*service password-encryption*” command with the “*no*” from, this does not decrypt the passwords.

R1

Service password-encryption

5. ENHANCED USERNAME PASSWORD SECURITY:

Command:

1. R1 (config) # username ccnp password 0 ccnp

0 Specifies an UNENCRYPTED password will follow

OR

R1 (config) # username ccnp password 7 kfhkjfhkrhfr

7 allows you to enter the ciphertext computed by the service password-encryption command.

```
2. R1 (config) # username ccnp secret 0 ccnp
```

0 Indicates that the following clear text password is to be hashed using MD5.

OR

```
R1 (config) # username ccnp secret 5 fhdsjhfsdfkjsdkfskfh
```

5 indicates that the following encrypted secret password was hashed using MD5

The jumbled word followed by the number 5 should be copied from the running configuration. This is the encrypted password from the enable secret command.

6. SETTING A LOGIN FAILURE:

- Cisco IOS Software releases 12.3(1) supports to configure the number of allowable unsuccessful login attempts by using the “*security authentication failure rate*” from the global configuration mode.
- By default, router allows 10 login failures before initiating a 15-second delay.
- Generates a syslog message when rate is exceeded.

Command:

```
R1
```

```
Security authentication failure rate 10 threshold-rate log
```

- Threshold-rate is the number of allowable unsuccessful login attempts. The default is 10 and the range is from 2 to 1024.
- The **log** keyword is required. Results in a generated syslog event.

R1

```
Config terminal
Enter configuration commands, one per line. End with CNTL/Z.
aaa new-model
aaa authentication login local_auth local
username user1 password cisco
security authentication failure rate 2 log

line console 0
login authentication local_auth
```

- To verify the above configuration, exit from the router and try to log in again.
- Login with incorrect password twice and try again for the third time.
- When the number of failed login attempts reaches the configured rate , two events occur:
 1. An error message is sent by the router.
 2. A 15-second delay timer starts. After the 15-second delay has passes, the user may continue to attempt to log in to the router.

Verification:

User Access Verification

```
Username: user1
Password: (incorrect password)
% Authentication failed
Username: user1
Password: (incorrect password)
% Authentication failed
```

```
*Mar 2 17:32:02.939: %LOGIN-3-TOOMANY_AUTHFAILS: Too many Login
Authentication failures have occurred in the last one minute on the line 0.
```

7. SETTING A LOGIN FAILURE BLOCKING PERIOD:

- With this login enhancement command available in Cisco IOS software release 12.3(4) T and later, the router will not accept any additional login connections for a “quiet period”, if the configured number of connection attempts fail within a specified time period.
- But, Hosts that are permitted by a predefined ACL are excluded from the quiet period by the global config command “*login quiet-mode access-class*”.
- Mitigates DoS and break-in attacks.

- All login attempts made via Telnet, Secure Shell (SSH) and HTTP are denied during the quiet period.

Command:

R1

```
login block-for seconds attempts tries within seconds
```

- Seconds: specifies the duration of time, or quiet period, during which login attempts are denied.
- Attempts: maximum number of failed login attempts that triggers the quiet period.
- Within: duration of time in seconds during which the allowed number of failed login attempts must be made before the quiet period is triggered.

R1

```
username user1 password user1
```

```
Enable secret cisco
```

```
Login block-for 30 attempts 4 within 20
```

```
Line vty 0 4
```

```
Login local
```

Verify the above configuration:

- Telnet from R2 (1.1.1.2) to R1(1.1.1.1)
- The router R1 prompts for username and password.
- Try logging with the incorrect password for 4 times as the attempt mentioned is 4.
- After the 4 unsuccessful attempts access to the router is denied and the router is in a quiet period for 30 seconds.
- Can be verified by the following ;

R1#show login

A default login delay of 1 second is applied.

No Quiet-Mode access list has been configured.

Router enabled to watch for login Attacks.

If more than 4 login failures occur in 20 seconds or less, logins will be disabled for 30 seconds.

Router presently in Quiet-Mode.

Will remain in Quiet-Mode for 18 seconds.

Denying logins from all sources

(18 seconds indicates the remaining time left form the 30 seconds to come out of the quiet period.)

R1#show login failures

Total failed logins: 28

Detailed information about last 50 failures

Username	SourceIPAddr	IPort	Count	TimeStamp
user1	1.1.1.2	23	8	19:07:58 UTC Sat Jul 7 2007
user2	1.1.1.2	23	3	19:06:15 UTC Sat Jul 7 2007
user3	1.1.1.2	23	1	19:04:40 UTC Sat Jul 7 2007

- The output displays number of users tried to login unsuccessfully via Telnet.

System logging messages for a quiet period :

- login on-success : Generated for successful login
- login on-failure : Generated for failed login requests.

8. EXCLUDING ADDRESSES FROM LOGIN BLOCKING:

- In Cisco IOS software release 12.3(4) T, the IOS router will use the configured ACL to permit login attempts when the router switches to quiet mode.

Command:

```
R1 (config) # login quiet-mode access-class {acl-name | acl-number}
```

- Configure an ACL permitting network 20.0.0.0 to R1.
- Configure login block-for 30 seconds

R1

```
Username user1 password user1
Enable secret cisco
Login block-for 30 attempts 4 within 20

Line vty 0 4
Login local

Access-list 1 permit 20.0.0.0 0.255.255.255
Login quiet-mode access-class 1
```

Verify the above configuration:

- Telnet from R2 (1.1.1.2) to R1 (1.1.1.1)
- R1 prompts for username and password.
- Try logging with incorrect password for 4 times as the attempts mentioned above is 4.
- After the 4 unsuccessful attempts access to the router is denied and the router is in a quiet period for 30 seconds.
- But with the login quiet-mode access-class command users from network 20.0.0.0 can try logging into the router even if the router is in the quiet period.
- Can be verified by the following: Try Telnet from the PC (20.1.1.20) to the router R1 and access is permitted though the router is in the quiet period.

9. SETTING A LOGIN DELAY:

- A Cisco IOS device can accept login connections such as Telnet, SSH and HTTP as fast as they can be processed.
- The **login delay** command introduces a uniform delay between successive login attempts.
- The delay occurs for all login attempts (failed and successful attempts).
- Secure the device from dictionary attacks, which are an attempt to gain username and password access to your device.
- The command was introduced in Cisco IOS software release 12.3(4)T.
- If not set, a default delay of one second is enforced.

10. SETTING TIMEOUTS:

- By default, an administrative interface stays active for 10 minutes after the last session activity.
- After that the interface times out and logs out of the session.

- Recommended to tune these timers for extra safety when an administrator walks away from an active console session.
- **Do not set the exec-timeout value to 0 as it indicates that there will be no timeout and the session will stay active for unlimited time.**

Command:

R1 (config-line) # exec-timeout minutes [seconds]

- Minutes: specifies the number of minutes the session will be terminated.

11. SETTING MULTIPLE PRIVILEGE LEVELS

- Cisco routers enable you to configure various privilege levels for your administrators.
- Different passwords can be configured to control who has access to the various privilege levels.
- Three types of levels :
 1. Level 0 : predefined for user-level access privileges.
 2. Level 2 to 14 : customized for user-level privileges.
 3. Level 15 | : predefined for enable mode.

Command:

R1 (config) # privilege mode {level level command | reset command}

- Mode : specifies the configuration mode.
- Level : enables setting a privilege level .
- Command : sets the command to which privilege level is associated.
- Reset : command resets the privilege level command.

Scenario:

Assign “ping” and “show” command to the privilege level 2 and establish “cisco” as the secret password for the users to enter the privilege level 2.

R1

Config t

Enter configuration commands, one per line. End with CNTL/Z.

Privilege exec level 2 ping

Privilege exec level 2 show

Enable secret level 2 cisco

Verify the above configuration :

- Using the enable (level) command router will prompt for password to enter into the privilege level 2.
- The users at this level are only restricted for exec commands only and not allowed access to the configuration mode as the mode specified is only exec mode.

Verified by the following:

```
R1>enable 2
Password:cisco
R1#show privilege
```

Current privilege level is 2

```
R1#config t
^
```

% Invalid input detected at '^' marker.

- This error message indicates that the user is restricted to only exec mode and not other modes.

12. CONFIGURE BANNER MESSAGES:

- Banner messages specify what the proper use of the system is.
- Specifies that the system is being monitored.
- Specifies that privacy should not be expected when using this system.

Command :

```
R1 (config) # banner {exec | incoming | login | motd | slip-ppp}
                %message%
```

Where the character (%) mentioned before the start of the message and end of message should be the same and must not be in the message, otherwise the message ends where the character (%) is seen in the line.

Example:

```
R1 (config) # banner motd % This device is netmetric property. Please handle with care.
%
```

13. CONFIGURING ROLE-BASED CLI :

- The role-based CLI access feature allows you to define “views” which are a set of operational commands and configuration capabilities that provide selective or partial access to Cisco IOS EXEC and configuration mode commands.
- Views restrict user access to Cisco IOS CLI and configuration information.
- View can define which commands are accepted and what configuration information is visible.
- Cisco IOS software release 12.3(11)T can also specify an interface or a group of interfaces to a view, allowing access based on specified interfaces.
- Access to the view is protected with a password.
- To simplify view management, views can be grouped to superviews to create large sets of commands and interfaces.
- Root view is the highest administrative view and creating and modifying a view or superview is possible only from root view.
- CLI views require AAA new-model.
- A maximum of 15 CLI views can exist in addition to the root view.

Scenario:

- Enable aaa new-model and create a view.
- Specify the mode in which the specified command exists.

R1

```

aaa new-model
exit

enable view
configure terminal
parser view view1
secret 0 cisco
command exec include show version
command exec include configure terminal
command exec include all show ip
exit

```

Verify the above configuration:

- If the user wants to enter into the root view.

```

R1 (config) # aaa new-model
R1 (config) # exit
R1 # enable view view1
The router prompts for password

```

Password: cisco

R1 #

- If the user wants to view the available commands in the view.

R1#?

Exec commands:

<1-99> Session number to resume
configure Enter configuration mode0
enable Turn on privileged commands
exit Exit from the EXEC
show Show running system information

- The output displays available commands in the exec mode .

R1#show ?

flash: display information about flash: file system
ip IP information
parser Display parser information
slot0: display information about slot0: file system
slot1: display information about slot1: file system
version System hardware and software status

- The output displays configured keywords ip and version apart form parser which is always available.

R1#show ip ?

accounting The active IP accounting database
aliases IP alias table
arp IP ARP table
as-path-access-list List AS path access lists
bgp BGP information
cache IP fast-switching route cache
casa display casa information
cef Cisco Express Forwarding
ddns Dynamic DNS
dfp DFP information
extcommunity-list List extended-community list
--More--

- The output displays all the sub-options available in the view.

Note :

- Role-based CLI facilitates the concept of grouping CLI views into view supersets , called superviews.
- A superview consists of one or more CLI views.
- CLI view can be shared among multiple superviews.
- Each superview has a password .
- If a superview is deleted , all CLI views associated with that deleted superview will not be deleted.

Configuration:

R1

```
aaa new-model
exit
```

```
enable view
```

```
configure terminal
parser view view1
secret 0 cisco
command exec include show version
command exec include configure terminal
command exec include all show ip
exit
```

```
parser view view2
secret 0 ccnp
command exec include show flash
command exec include ping
exit
```

```
parser view superview1
password 0 ccnp1
view view1
view view2
```

Verify:

R1 # show parser view [all]

The output will display all the CLI views configured on the router.

14. SECURE CONFIGURATION FILES :

- The Cisco resilient configuration feature enables a router to secure and maintain a working copy of the running image and configuration so that those files can withstand malicious attempts to erase the contents of persistent storage in NVRAM and flash.
- This set of image and router running configuration is referred to as the primary bootset.
- This feature is available only on platforms that support a PCMCIA card.

Command:

R1

```
secure boot-image
secure boot-config
```

- The above configuration can be verified :

R1 # show secure bootset

- The output displays the status of the configuration resilience and the primary bootset filename.

Secure configuration files recovery :

- Use the reload command in the privilege mode to restart it and interrupt the boot sequence to enter the ROMMON mode .
- In the ROMMON , use the dir and boot commands to view the contents of the file system and select a secure image to boot the router from .

Command:

```
rommon 1 > dir slot0:
rommon 2 > boot slot0:c3745-js2-mz
```

- If the startup configuration was deleted, the router will prompt for interactive configuration input.
- You should decline to enter an interactive configuration session in setup mode if you secured the configuration file.
- Use the secure boot-config restore command to recover the secured startup configuration and save it under a specified filename.
- Finally copy the recovered file to the running configuration to resume normal operations.

R1

```
secure boot-config restore slot0:rescue  
copy slot0:rescue running-config
```

- Restores the secure configuration to a filename.

PAPER 4


OPTIMIZING CONVERGED CISCO NETWORKS

ONT (642–845)

ONT LAB INDEX

- 1. CONFIGURING LEGACY CUSTOM QUEUEING.**
- 2. CONFIGURING LEGACY PRIORITY QUEUEING**
- 3. CONFIGURING MODULAR QOS CLI WITH NBAR**
- 4. CONFIGURING MQC WITHOUT NBAR**
- 5. CONFIGURING MQC LOW LATENCY QUEUEING WITH NBAR**
- 6. CONFIGURING MQC LLQ WITHOUT NBAR**
- 7. CONFIGURING LEGACY GENERIC TRAFFIC SHAPING**
- 8. CONFIGURING LEGACY COMMITTED ACCESS RATE**
- 9. CONFIGURING MQC POLICING**

Lab 1 – Configuring Legacy Custom Queueing



Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure Custom Queueing on R1 so that traffic leaving e3/0 is guaranteed 50 % of bandwidth for HTTP, 20 % of bandwidth for SMTP, 10 % bandwidth for TELNET and 20 % bandwidth for other traffic.

Create custom queue list 1

Assign the traffic according to the bandwidth availability and allocate byte counts.

Apply the custom queue list to the e3/0 interface on R1.

R1

```
Queue-list 1 protocol ip 1 tcp www
Queue-list 1 protocol ip 2 tcp smtp
Queue-list 1 protocol ip 3 tcp telnet
Queue-list 1 default 4
```

```
Queue-list 1 queue 1 byte-count 5000
Queue-list 1 queue 2 byte-count 2000
Queue-list 1 queue 3 byte-count 1000
Queue-list 1 queue 4 byte-count 2000
```

```
Int e3/0
Custom-queue-list 1
```

Verification:

R1#show queueing custom

Current custom queue configuration:

```
List Queue Args
1 4 default
1 1 protocol ip tcp port www
1 2 protocol ip tcp port smtp
1 3 protocol ip tcp port telnet
1 1 byte-count 5000
1 2 byte-count 2000
1 3 byte-count 1000
1 4 byte-count 2000
```

The output displays the current custom queue configuration.

R1#show interfaces e3/0

```
Ethernet3/0 is up, line protocol is up
Hardware is AmdP2, address is 000b.5f56.fb31 (bia 000b.5f56.fb31)
Internet address is 10.1.1.1/8
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
  reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:02:15, output 00:00:07, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: custom-list 1
Output queues: (queue #: size/max/drops)
  0: 0/20/0 1: 0/20/0 2: 0/20/0 3: 0/20/0 4: 0/20/0
  5: 0/20/0 6: 0/20/0 7: 0/20/0 8: 0/20/0 9: 0/20/0
 10: 0/20/0 11: 0/20/0 12: 0/20/0 13: 0/20/0 14: 0/20/0
 15: 0/20/0 16: 0/20/0
```

As the custom-queue is applied to this interface, we can see from the output that the queuing strategy used is custom-queue.

By default 16 queues can be assigned in one custom queue.

Lab 2 – Configuring Legacy Priority Queuing

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure legacy priority on R1 so that traffic leaving e3/0 interface is prioritized accordingly.

Assign the traffic priority

- High – http
- Medium – telnet
- Normal – IP
- Low - other

Apply the priority list to the Ethernet interface

R1

```
Priority-list 1 protocol ip high tcp www
Priority-list 1 protocol ip medium tcp telnet
Priority-list 1 protocol ip normal
Priority-list 1 default
```

```
Int e3/0
Priority-group 1
```

Verification :

Router#show queuing priority

Current DLCI priority queue configuration:

Current priority queue configuration:

```
List Queue Args
1 low default
1 medium protocol ip tcp port telnet
1 normal protocol ip
1 high protocol ip tcp port www
```

The output displays the current priority queue configuration.

Router#show interfaces e3/0

```
Ethernet3/0 is up, line protocol is up
Hardware is AmdP2, address is 000b.5f56.fb31 (bia 000b.5f56.fb31)
Internet address is 10.1.1.1/8
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:39, output 00:00:09, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: priority-list 1
Output queue (queue priority: size/max/drops):
high: 0/20/0, medium: 0/40/0, normal: 0/60/0, low: 0/80/0
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
```

The output displays the queuing strategy for that particular interface.

Router#show queuing interface e3/0

Interface Ethernet3/0 queuing strategy: **priority**

```
Output queue utilization (queue/count)
high/182 medium/0 normal/4784 low/413
```

The output displays the created priority-list and the priorities.

Lab 3 – Configuring Modular QOS CLI With NBAR

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure the MQC on R1 so that traffic leaving ethernet 3/0 interface is guaranteed the following amount of bandwidth, i.e., http 50%, SMTP 20 %, telnet 10%, and other 20 %.

Create class-map and assign the traffic to the class-map.

Create policy-map and assign the class-map --- certain amount of bandwidth.

Apply the policy to the interface.

R1

Ip cef

Class-map match-all map1

Match protocol http

Class-map match-all map2

Match protocol smtp

Class-map match-all map3

Match protocol telnet

Policy-map pmap

Class map1

Bandwidth percent 50

Class map2

Bandwidth percent 20

Class map3

Bandwidth percent 10

Class class-default

Bandwidth percent 20

```
Int e3/0
Max-reserved-bandwidth 100
Service-policy output pmap
```

Note: Note that the bandwidth value to be reserved on the ethernet interface is set to 100%.

Verification:

Router#show policy-map interface ethernet 3/0

```
Ethernet3/0
Service-policy output: pmap
Class-map: map1 (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: protocol http
  Queueing
    Output Queue: Conversation 265
    Bandwidth 50 (%)
Class-map: map2 (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: protocol smtp
  Queueing
    Output Queue: Conversation 266
    Bandwidth 20 (%)
Class-map: map3 (match-all)
  207 packets, 17433 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: protocol telnet
  Queueing
    Output Queue: Conversation 267
    Bandwidth 10 (%)
Class-map: class-default (match-any)
  30 packets, 2897 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any
  Queueing
    Output Queue: Conversation 268
    Bandwidth 20 (%)
```

The output displays all the class-maps created.

Lab 4 – Configuring MQC without NBAR

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure MQC on R1 so that traffic leaving its ethernet 3/0 interface is guaranteed the amount of bandwidth configured.

Create class-map for traffic classification.

Traffic classified using ACL.

Create policy-map to apply QOS features to the class-map

Apply to the interface.

R1

Ip cef

```
Class-map match-all map1
Match access-group name map1
Class-map match-all map2
Match access-group name map2
Class-map match-all map3
Match access-group name map3
```

```
Policy-map pmap
Class map1
Bandwidth percent 50
Class map2
Bandwidth percent 20
```

```
Class map3
Bandwidth percent 10
Class class-default
Bandwidth percent 20
```

```
Int e3/0
Max-reserved-bandwidth 100
Service-policy output pmap
```

```
Ip access-list extended map1
permit tcp any any eq www
permit tcp any eq www any
```

```
Ip access-list extended map2
permit tcp any any eq smtp
permit tcp any eq smtp any
```

```
Ip access-list extended map3
permit tcp any any eq telnet
permit tcp any eq telnet any
```

Verification:

Router#show policy-map interface e3/0

```
Ethernet3/0
Service-policy output: pmap
```

```
Class-map: map1 (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: access-group name map1
  Queueing
    Output Queue: Conversation 265
    Bandwidth 50 (%)
Class-map: map2 (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: access-group name map2
  Queueing
    Output Queue: Conversation 266
    Bandwidth 20 (%)
```



```
Class-map: map3 (match-all)
  17 packets, 1113 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: access-group name map3
  Queueing
    Output Queue: Conversation 267
    Bandwidth 10 (%)
Class-map: class-default (match-any)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any
  Queueing
    Output Queue: Conversation 268
    Bandwidth 20 (%)
```

The output displays the classmap's created.

Router#show queueing interface e3/0

```
Interface Ethernet3/0 queueing strategy: fair
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: Class-based queueing
Output queue: 0/1000/64/0 (size/max total/threshold/drops)
Conversations 0/1/256 (active/max active/max total)
Reserved Conversations 4/4 (allocated/max allocated)
Available Bandwidth 0 kilobits/sec
```

Lab 5 – Configuring MQC Low Latency Queueing with NBAR

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure MQC on R1 so that telnet traffic up to 500 kbps is sent first.

Create class-map and policy-map.

Apply the policy to the ethernet interface

R1

Ip cef

Class-map match-all map1
Match protocol telnet

Policy-map pmap
Class map1
Priority 500

Int e3/0
Service-policy output pmap

Verification:

Router#show policy-map interface e3/0

Ethernet3/0

Service-policy output: pmap

Class-map: map1 (match-all)
0 packets, 0 bytes
5 minute offered rate 0 bps, drop rate 0 bps
Match: protocol telnet
Queueing
Strict Priority
Output Queue: Conversation 264
Bandwidth 500 (kbps) Burst 12500 (Bytes)
(pkts matched/bytes matched) 0/0
(total drops/bytes drops) 0/0

Class-map: class-default (match-any)
5 packets, 570 bytes
5 minute offered rate 0 bps, drop rate 0 bps
Match: any

The output displays policy-map created and the class-maps within the policy-map.

Lab 6 – Configuring MQC LLO without NBAR

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure MQC on R1 so that all telnet traffic upto 500 kbps is sent first without NBAR.

Create class-map and policy-map.

Apply the policy to the ethernet interface

R1

```
Ip cef
```

```
Class-map match-all map1  
Match access-group name group1
```

```
Policy-map pmap  
Class map1  
Priority 500
```

```
Ip access-list extended group1  
Permit tcp any any eq telnet  
Permit tcp any eq telnet any
```

```
Int e3/0  
Service-policy output pmap
```

Verification:

Router#show policy-map interface e3/0

Ethernet3/0

Service-policy output: pmap

Class-map: map1 (match-all)

0 packets, 0 bytes

5 minute offered rate 0 bps, drop rate 0 bps

Match: access-group name group1

Queueing

Strict Priority

Output Queue: Conversation 264

Bandwidth 500 (kbps) Burst 12500 (Bytes)

(pkts matched/bytes matched) 0/0

(total drops/bytes drops) 0/0

Class-map: class-default (match-any)

3 packets, 180 bytes

5 minute offered rate 0 bps, drop rate 0 bps

Match: any

The output displays the class-map created in the policy-map that is applied on this interface.

Lab 7 – Configure Legacy Generic Traffic Shaping

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure legacy generic traffic shaping on R1 to limit the output rate on the ethernet interface to 720 kbps.

Use a committed bursts value of 60 kbps.

R1

```
Int e3/0
Traffic-shape rate 720000 60000 0 1500
```

Where:

- 720000 – traffic bit rate in bits per second.
- 60000 – bits per interval sustained.
- 0 -- bits per interval excess in first interval.
- 1500 -- set buffer limit.

Verification:

Router#show traffic-shape

```
Interface Et3/0
  Access Target Byte Sustain Excess Interval Increment Adapt
VC List Rate Limit bits/int bits/int (ms) (bytes) Active
- 720000 7500 60000 0 83 7500 -
```

The output displays the configured routes to shape the traffic.

Router#show traffic-shape statistics

```
Acc. Queue Packets Bytes Packets Bytes Shaping
I/F List Depth Delayed Delayed Active
Et3/0 0 31 3210 0 0 no
```

The output displays all the packets delayed and bytes delayed

Router#show traffic-shape queue

```
Traffic queued in shaping queue on Ethernet3/0
Queueing strategy: weighted fair
Queueing Stats: 0/1500/64/0 (size/max total/threshold/drops)
Conversations 0/0/64 (active/max active/max total)
Reserved Conversations 0/0 (allocated/max allocated)
Available Bandwidth 720 kilobits/sec
```

The output displays the available bandwidth that is configured.

Lab 8 – Configuring Legacy Committed Access-Rate

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure legacy committed access rate on R1 to limit the input rate on e3/0 to 720 kbps. All other traffic above this rate to be dropped.

R1

```
Int e3/0
Rate-limit input 720000 15000 15000 confirm-action transmit exceed-action drop
```

Where:

720000 – bits per second (rate)
15000 -- normal burst size in bytes
15000 -- maximum burst in bytes.

Router#show interfaces ethernet 3/0 rate-limit

Ethernet3/0

Input

matches: all traffic

params: 720000 bps, 15000 limit, 15000 extended limit

conformed 0 packets, 0 bytes; action: transmit

exceeded 0 packets, 0 bytes; action: drop

last packet: 19565252ms ago, current burst: 0 bytes

last cleared 00:00:24 ago, conformed 0 bps, exceeded 0 bps

The output displays the configured rates to this particular interface e3/0 on R1.

Lab 9 – Configuring MQC Policing

(Scenario Based On Lab 1)
Interface IP Address Configuration

R1

Interface	IP Address	Subnet Mask
E 3/0	10.1.1.1	255.0.0.0

Lab Objective:

Task 1

Configure MQC policing on R1 to limit the input rate on the ethernet interface to 720 kbps. All traffic above this rate is dropped.

R1

```
Policy-map pmap
Class class-default
Police cir 720000 bc 15000 bc 15000
Confirm-action transmit
Exceed-action drop

Int e3/0
Service-policy input pmap
```

Where:

BC → confirm burst
Burst of 15000 in bytes

Be → excess burst
Burst of 15000 in bytes

Verification:

Router#show policy-map interface e3/0

Ethernet3/0

Service-policy input: pmap

Class-map: class-default (match-any)

3 packets, 276 bytes

5 minute offered rate 0 bps, drop rate 0 bps

Match: any

police:

cir 720000 bps, bc 15000 bytes

conformed 3 packets, 276 bytes; actions:

transmit

exceeded 0 packets, 0 bytes; actions:

drop

conformed 0 bps, exceed 0 bps

The output displays the configured traffic policing parameters on this particular interface.

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com

TRAINING TOMMORROW'S PROFESSIONAL TODAY
INFO@NETMETRIC-SOLUTIONS.COM

Page 315 of 315

NETMETRIC-SOLUTIONS

www.netmetric-solutions.com

All contents are copyright @ 2007-2010

All rights reserved.